



Faculteit Wetenschappen
Departement Fysica

A Decision Support System for the Assessment and Stenting of Tracheal Stenosis

Een Beslissingsondersteunend Systeem voor de Evaluatie en het Stenten van Tracheale Stenose

Proefschrift voorgelegd tot het behalen van de graad van

Doctor in de Wetenschappen

aan de Universiteit Antwerpen, te verdedigen door

Rômulo PINHO

Promotor

Prof. Dr. Jan Sijbers

Doctoral committee

Chairman

Prof. Dr. Joris Dirckx

Members

Prof. Dr. Jan Sijbers (Thesis supervisor)

Prof. Dr. Paul Scheunders

External members

Prof. Dr. Rudi Deklerck - Vrije Universiteit Brussel

Prof. Dr. Bram van Ginneken - Radboud Universiteit Nijmegen

Prof. Dr. Kurt Tournoy - Universiteit Gent

All the quotes reproduced in this thesis have been taken from the Web. Their authenticity (apart from some cartoon quotes which I know by heart) has not been verified.

Contents

Acknowledgements	xiii
Summary	xvii
1 Introduction	1
1.1 The Proposed System	4
1.2 Thesis Organization	7
I Preliminaries	9
2 Clinical Background	11
2.1 Introduction	12
2.2 Anatomy of the Lower Airways	12
2.3 Anatomy of the Trachea	14
2.4 Tracheal Stenosis	14
2.5 Tracheal Stents	16
3 Related Work	21
3.1 Introduction	22
3.2 Traditional Methods for Airway Assessment	22
3.2.1 Rigid Bronchoscopy	22
3.2.2 Flexible Bronchoscopy	23
3.3 Computer-Aided Methods	24

3.3.1	Manual Methods	25
3.3.2	Semi-automatic Methods	27
II	A System for Assessment and Stenting of Tracheal Stenosis	31
4	Out-of-core Image Processing	33
4.1	Introduction	34
4.2	<i>WindowCache</i>	37
4.2.1	The Cache Buffer	37
4.2.2	Prefetching	38
4.2.3	Local Prefetching	40
4.3	Experiments	48
4.4	Results	49
4.4.1	Raster Scan Algorithms	49
4.4.2	Convolution Algorithms	52
4.4.3	Propagative Algorithms	53
4.5	Discussion	55
4.6	Conclusion	56
5	Airway Tree Segmentation	57
5.1	Introduction	58
5.2	Method	60
5.2.1	Seed Point Selection	65
5.3	Experiments	69
5.4	Results	70
5.5	Discussion	71
5.6	Conclusions	76
6	Estimation of Healthy Tracheas	79
6.1	Introduction	80
6.2	Training Set	81
6.3	Establishment of Correspondences	82

6.4	Active Shape Model	83
6.5	ASM Registration	85
6.5.1	Initialization	86
6.5.2	Rigid Alignment	87
6.5.3	Pose and Shape Registration	87
6.6	Experiments	95
6.6.1	Quantitative Comparison	95
6.6.2	Qualitative Comparison	98
6.7	Results	99
6.7.1	Quantitative Comparison	99
6.7.2	Qualitative Comparison	101
6.8	Discussion	102
7	Segmentation of Tracheal Stenosis	107
7.1	Introduction	108
7.2	Active Contour Models	108
7.3	Segmentation of the Stenotic Trachea	110
7.3.1	Initialization	110
7.3.2	External Force	111
7.3.3	Internal Forces	112
7.3.4	Surface Manipulation	115
7.4	Discussion	116
8	Assessment and Stenting of Tracheal Stenosis	119
8.1	Introduction	120
8.2	Assessment of Stenosis	120
8.3	Prediction of Stents	122
8.4	Experiments	123
8.4.1	Simulation Data	124
8.4.2	Clinical Data	126
8.5	Results	127
8.5.1	Simulation Data	127
8.5.2	Clinical Data	133
8.6	Discussion	139

8.6.1	Fixed Landmarks	139
8.6.2	Execution Times	139
III	Conclusions	145
9	Conclusion	147
10	Future Work	151
10.1	Cache and Prefetching	152
10.2	Airway Tree Segmentation	152
10.3	<i>Fixed Landmarks</i>	152
10.4	Functional Stent Analysis	154
	Contributions	155
	Bibliography	159
	Nederlandse Samenvatting	171
	Agradecimientos	175

List of Figures

2.1	Airway structure (Ikeda, 1974).	13
2.2	Different views of the trachea and its composing structures. (a) Illustration of the trachea and main bronchi. (b) Schematic view of an axial cross-section.	15
2.3	(a) A schematic view of stenosis (by courtesy of Dr. P. De- laere, Center for Larynx, Trachea & Hypopharynx Recon- struction, Katholieke Universiteit Leuven, Belgium). (b) A real case showing one sagittal slice of a chest CT scan. . . .	17
2.4	Typical tracheal stents. Top row, from left to right: Mont- gomery T stent, Dumon stent, Polyflex stent and Noppen stent. Bottom row, from left to right: Ultraflex stent, Aero stent, Eco Nanjing stent, Hanaro stent and Tae woong stent. (Freitag, 2010, copyright: ERS, 2010)	18
3.1	(a) Dr. Killian and the rigid bronchoscope. (b) Flexible bronchoscopy.	23
4.1	Typical image processing pipeline.	34
4.2	Illustration of the cache structure for a 3D image volume. (a) A 3D image with its dimensions. (b) Virtual subdivision of the image. Each block is referenced by a unique identifier. (c) Cache (or image) block and its dimensions.	38

4.3	Sequence diagram of the prefetching process. The main program synchronously requests pixels, and the cache system asynchronously decides when to load another image block from disk.	40
4.4	Different sliding window configurations. The crisper areas of the pictures represent the image regions actually present in main memory.	42
4.5	Prefetching window for sequential access patterns.	45
4.6	Situation when the mask reaches the border between two blocks: (a) the mask at the beginning of the central block; (b) the mask at the border; (c) the cache slides and the last block is prefetched; (d) the mask remains at the border. The arrows indicate the traversal pattern.	46
4.7	<i>WindowCache</i> when the convolution mask reaches the border between two blocks: (a) the mask is in the central block; (b) the mask invades next block, but no prefetch is triggered; (c) the mask exceeds the extra space \mathbf{S} and a prefetch is triggered; (d) the mask is completely inside the central block. The arrows indicate the traversal pattern.	47
4.8	Prefetching window for 6-connected, propagative access patterns.	48
5.1	Steps of the basic algorithm.	61
5.2	Diameter distributions of the airway branches (Tschirren et al., 2005)	62
5.3	(a) Schematic view of the proposed skeleton estimation. (b) An example with real data.	64
5.4	Avoiding leaks.	65
5.5	Mislead of the detection of the trachea.	67

-
- 5.6 Example of the trachea detection algorithm for CASE24 in the image database. (a) The slice of the image chosen as the beginning of the trachea (top) and a schematic view of the parameters of the tracheal contour used by the detection algorithm (bottom). (b) 3D views of the selected tracheal contour (in red). 73
- 5.7 Segmentations with ROIs coloured per level (top) and respective 3D reconstructions (bottom) for CASE34 (left) and CASE39 (right), respectively. 75
- 5.8 Positioning ROIs from level L to $L+1$. (a) Uncovered regions (in red) if the direction of ROI_{L+1} differs from that of ROI_L . (b) Zoomed area showing the proposed solution, i.e., move ROI_{L+1} "backwards" along the corresponding centre line of ROI_L so as to cover the longest perpendicular distance from the top of ROI_{L+1} to the indicated intersection. 77
- 6.1 Example of the process to transform the shapes of the training set from a centre line to a boundary representation and to establish correspondences between them. (a) Centre line \mathcal{C}_i , obtained directly from the image segmentation, displayed with contours \mathcal{B}'_{ik} . (b) Centre line \mathcal{C}'_i , after arc-length parametrization of \mathcal{C}_i , also shown with contours \mathcal{B}'_{ik} . (c) Centre line \mathcal{C}''_i and contours \mathcal{B}''_{ik} , after b-spline interpolation of \mathcal{C}'_i and \mathcal{B}'_{ik} , respectively. (d) Smooth surface \mathcal{X}_i , in wireframe mode. (e) Smooth surface \mathcal{X}_i , in shaded mode. The lines crossing the contours of (a), (b), and (c) indicate the position at which the sampling of the boundary starts (the region of the carina corresponds to the bottom contour). 84
- 6.2 From (a) to (f), the first 6 modes of variation. Each figure shows the average shape (middle), $-3\sqrt{\lambda_i}$ (left), and $+3\sqrt{\lambda_i}$ (right) along the mode. 86
- 6.3 Huber, Talwar, Cauchy, and Gaussian weighting functions. 89

-
- 6.4 The *Fixed Landmarks* concept: (a) $\hat{\mathbf{x}}^{(k)}$ (dashed) is attracted by high gradients of the stenotic trachea in the image (thick lines) – the arrows indicate the landmark displacements $d\mathbf{y}^{(k+1)}$. (b) the set of landmarks that remain fixed, $\hat{\mathbf{x}}^{(k)}|_{\mathcal{L}''}$, because the displacement was longer than d . (c) $\hat{\mathbf{x}}^{(k+1)}$ – the fixed landmarks keep the shape far from the stenotic area. 93
- 6.5 Examples of the *Fixed Landmarks* method in action with simulation data (top) and clinical data (bottom). The green shape is the shape generated by the model at one iteration and the red dots indicate the landmarks of the candidate shape, which match high gradients in the image. From left to right, the figures show the trachea with stenosis, two non-consecutive steps of the rigid alignment and two non-consecutive steps of the *Fixed Landmarks*. 94
- 6.6 Application to build phantoms of tracheal stenosis. 96
- 6.7 Per-group μ_{\max_g} (a), μ_{mean_g} (b), and respective standard deviation bars for each method across the whole set of leave-one-out tests. Along the horizontal axis, the number of phantoms in the test group, S_g , is shown in parentheses. 100
- 6.8 Shape estimation for phantom II-4A of one healthy trachea from the simulation experiments, using, from left to right, the *StandardLS*, *Gauss WLS*, *Huber WLS*, *Surface Extrapolation*, and *Fixed Landmarks* methods. The dashed, outermost silhouettes represent the healthy trachea used to build the phantom. 102
- 6.9 Results of the estimation of the healthy trachea with the *Fixed Landmarks* for patients 1, 2, and 3, from left to right. The estimated trachea, in green, is shown in the CT scan of the patient, overlaid on their segmented stenotic trachea. 103

6.10	<i>Gauss WLS</i> (a), <i>Surface Extrapolation</i> (b), and <i>StandardLS</i> (c) registration methods applied to the CT scan of Patient 1. The estimated trachea, in green, is placed in the image, overlaid on the segmented stenotic trachea. With <i>Surface Extrapolation</i> , the estimated surface failed to match the lower part of the trachea. With other methods, the surface was too narrow or deformed (compare with Figure 6.9(a)).	104
7.1	Distance map I_D used in the establishment of the external force of the ACM.	111
7.2	The elastic force acting on a point \mathbf{x}_{v_j} of the surface. (a) The general case, in which the point is moved to a central location relative to its One-ring. (b) When \mathbf{x}_{v_j} is at one of the open ends of the cylindrical surface.	114
7.3	The general case of the application of the bending force, so that sharp angles are removed from the surface.	115
7.4	Illustration of a difficult case (Patient 4) due to severe stenosis (a) and how the ACM managed to reconstruct the missing part of the tracheal wall (b).	117
8.1	Flowchart of the proposed method.	123
8.2	Evaluation of the proposed method with simulation data. The graphs show the cross-sectional area profiles and the parameters of stenosis for the estimated and the ground truth versions of the phantom tracheas.	129
8.2	(Cont.) Evaluation of the proposed method with simulation data. The graphs show the cross-sectional area profiles and the parameters of stenosis for the estimated and the ground truth versions of the phantom tracheas.	130
8.3	The graphs show, for each phantom, the minimum, 1 st quartile, median, 3 rd quartile, maximum, and average errors of the estimated parameters of the 72 simulated categories of stenosis relative to the ground truth (the maximum values are shown at the top of the graph).	131

8.4	Average estimation error per simulated category of stenosis, along with their respective standard deviations.	132
8.5	The graphs show, for each phantom, the percentiles and average errors of the degree of narrowing of the 72 simulated categories of stenosis relative to the ground truth when using the <i>Talwar WLS</i> and the <i>Fixed Landmarks</i> methods.	133
8.6	GUI of the application developed for manual/visual assessment of stenosis. The expert marks on the patient's CT image the parameters of the stenosis (start, end, point of maximum degree) and chooses the degree of narrowing from ranges of values. Zooming and panning enhances the visualization.	134
8.7	Stenosis quantification for the experiments with clinical data. The graphs show the cross-sectional area profiles and the parameters of stenosis for the narrowed and estimated healthy versions of the patients' tracheas.	137
8.7	(Cont.) Stenosis quantification for the experiments with clinical data. The graphs show the cross-sectional area profiles and the parameters of stenosis for the narrowed and the estimated healthy versions of the patients' tracheas.	138
8.8	Visualization of the stents computed for the patients in the experiments with clinical data.	140
8.8	(cont.) Visualization of the stents computed for the patients in the experiments with clinical data.	141
8.9	Variation in the results obtained from the simulation experiments when varying the parameter d of the <i>Fixed Landmarks</i> method. The graphs show how the error in the computation of the degree of narrowing of the stenosis increased when $d \neq 1\text{mm}$	143

List of Tables

4.1	Cache sizes and number of blocks in the experiments with sequential traversal patterns.	50
4.2	Statistics computed for the sequential access pattern.	51
4.3	Statistics computed for the blur filter. FPS_s refers to the frame rate using the sequential caches.	53
4.4	Statistics computed for the region growing algorithm.	54
5.1	Evaluation measures for the twenty cases in the test set.	72
5.2	Execution times of the algorithms applied to the test set.	74
6.1	Categories of stenosis based on location and length.	97
6.2	Categories of stenosis based on its degree of narrowing.	97
8.1	Categories of stenosis based on location and length.	125
8.2	Categories of stenosis based on its degree of narrowing.	125
8.3	Samples of the manual/visual measurements performed by 3 experts in the pulmonology field. All values are in millimetres and indicate the position along the axial axis where the corresponding parameters of the stenosis were marked by the experts, relative to the start of the patients' CT scan.	135
8.4	Variability in experts' manual/visual assessment of stenosis, for the 9 patients used in the experiments. The table shows the standard deviation (in millimetres) observed when the experts manually determined the parameters of the stenoses.	136

8.5 Approximate execution times of each step of the proposed method, measured for one example case. 142

Acknowledgements

Although contrary to the prerequisites, it is often said that a PhD is never done alone. Indeed, the list of people that were part of my PhD endeavour is not short. I will start with those that were closer to me at all times.

I could not be happier to have such wonderful and loving parents, Cleber and Madalena, who regardless of the anguish of having their only son living abroad for almost 5 years¹, supported me a lot. All their sacrifice to give me good education is being rewarded now, and my achievement is no less than theirs too. Concerning this achievement, there is an alleged Swedish proverb that says, “shared sorrow is half sorrow and shared joy is double joy”. My PhD would have been meaningless had I not have the love of my life, Adriana, to share it with. She is part of every word of this text, of every line of code I had to write, of every extra hour of work. She was there to cherish the good moments and to cheer up the not so good ones, and I am sure she is as happy for me now as I was for her when she defended her PhD.

Still with proverbs, my grandmother Dimas repeatedly said two things during my PhD that I will never forget: “It is very important to work on something that you like.” and, “I always pray that you will be a winner”. Winner or not, and despite the physical distance, I know she was always by my side, and so was I by hers. My dear mother-in-law, Regina, is someone that was always by side as well, and I am sure that she supported me as much as if I were her own son.

¹Blessed is he who invented the webcam, right mom?

But when you are a son, grandson, or son-in-law living in a different country, while taking time to adapt, it is your job that fulfils you most of the time. A good topic to work on and a good supervisor to interact with are therefore important parts of the process. My supervisor was Prof. Dr. Jan Sijbers. I believe I will not exaggerate if I say that Jan is unanimously held in high regard by all his doctoral students. He was always available for discussions, even by chat on Sunday nights, and often in a good mood. He blindly relied in my capacity to work remotely, allowing me to finish my thesis from outside Belgium, for which I was very grateful. For these and other reasons, he proved to be not only a great supervisor, but a great person too.

Still, a good supervisor is just part of the equation. A friendly environment to work in is of course complementary, and I also found it in the VisionLab. I found it, for example, in the cheerfulness of my good friend Sander van der Maar (who was always generous to have me as a guest in the several times I needed to travel to Belgium), in the friendliness of Ruben Capiiau and the many photography chats we had, in the helpfulness of Toon Huysmans and Gert van Gompel, in the badminton matches with Mai Zhenhua, in the friendly smile of Juntu Jaber, in the enriching project discussions with Joost Batenburg, etc. All of them, even those who came and went, contributed in their own way to making the VisionLab a very nice place in which to work.

I was also very fortunate to meet Dr. Kurt Tournoy during the course of the development of my thesis. In a moment where I seemed to be blocked due to the lack of images to test my method on, Dr. Tournoy not only appeared as the provider of such images, but was always willing to collaborate, being an important source of medical feedback.

But before coming to Belgium and to the VisionLab, I needed to quit my job in Brazil. Quitting a job is never an easy task, especially one that you like. However, when one has such a nice boss as Mr. Silvio Pereira from TV Globo, the strain of the moment in fact turns into unmeasurable support. It is not by chance that Silvio is nowadays one of my dearest friends.

Dear is also Prof. Dr. Alexandre Plastino, exemplary professor from

my undergraduate period, at the Universidade Fereal Fluminense (UFF). Alexandre showed immense willingness when I asked him reference letters for the PhD post. It is not by chance either that he is one of the big names at UFF.

Other people have indirectly taken part in this PhD endeavour. People such as Sabine Covent, who hosted my wife and I in her B&B during the first weeks in Belgium and gave us endless help (and even a bed) while we were settling down. Lastly, Armand and Marília Janssens, more than simply nice landlords, were people of unlimited generosity, who we will carry in our hearts wherever we are.

Despite being a rather rational person, I firmly believe that there are moments in life that transcend the explanations given by maths, physics, statistics, logic, chance, luck, merit, etc. The roadmap of this PhD endeavour has eventually been beyond any plan or strategy, something that only a superior force is able to control. In ways I will certainly never be able to comprehend, this force paved the (sometimes tortuous) way that led me to this very special moment and put in it the people and things that helped me making this PhD dream come true. I therefore offer to You, my God, two humble, but sincere words:

THANK YOU.

Summary

Tracheal stenosis is an unnatural narrowing of the trachea that, despite being relatively rare, can be life threatening. Possible choices of treatment include surgery and the use of stents. Whichever approach is taken, an important diagnostic step is the correct assessment of the stricture to determine the point where it starts, where it ends, and the degree of narrowing. Such assessment is to date very operator dependent and despite advances in the imaging and computing fields, it remains mostly a manual operation.

This thesis sets forth a decision support system that proposes a method for automatic assessment of tracheal stenosis and prediction of stent length and diameter from chest CT scans. The main idea behind the proposed method is to estimate the shape of the trachea of a patient as if stenosis were not present. This shape can be used the by the physician for surgery planning and is the basis for the automatic assessment of the stenosis and the prediction of patient-specific stents.

Several challenges were solved in order to achieve the system's goal. The first was how to efficiently manipulate large medical image files. With current technology, CT scans of the chest can generate files whose sizes are in the order of gigabytes. When these files need to be manipulated in current desktop computers, memory restrictions often arise. This thesis proposes an *out-of-core* image processing technique that keeps only parts of the large file in memory. Such technique, named *WindowCache*, employs a cache and prefetching strategy that optimizes the file traversal using a

priori knowledge about the data access pattern of the image processing algorithm.

The second challenge was the automatic segmentation of healthy tracheas from the CT images, an important pre-preprocessing step of the proposed method. More than solving this problem, this thesis presents a method for the segmentation of the entire airway tree, which comprises the trachea and bronchi. The proposed method, based on region growing, automatically identifies the start and end of the trachea and introduces a new technique to avoid the appearance of leaks in the segmentation, which is a classical problem of region growing.

The estimation of the patient's healthy trachea is achieved thanks to the construction of a statistical shape model of healthy tracheas. Such a model, as described by Cootes et al. (1995) captures the geometric variation present in a training set of shapes and is capable of generating new shape instances that are constrained by the acquired variation. In other words, the model proposed here can only generate healthy tracheas. The objective is then to register this model to the CT image of a patient with tracheal stenosis. Although the shape resulting from the registration never contains local geometric distortions typical of stenosis, the narrowed areas of the patient's trachea can globally influence the registration algorithm. This thesis further proposes a new registration method called *Fixed Landmarks* to more robustly avoid these narrowed areas and generate a plausible estimation of the patient's healthy trachea.

In the following step, the proposed system automatically segments the narrowed trachea from the image, using the estimated healthy trachea as the starting point. The segmentation is based on the classic Active Contour Model by Kass et al. (1988), with internal and external forces tailored for the task. The obtained healthy and the narrowed versions of the trachea are finally used in an automatic procedure that computes the parameters of the stenosis (start, end, degree of narrowing) and uses these parameters to predict the length and diameter of a stent to be used in that patient.

Extensive sets of experiments were carried out to validate all proposed

methods. For the *out-of-core* technique, different algorithms and cache structures were tested on large files. The airway segmentation method was tested on a database of 40 image volumes, subdivided into training and testing sets. For the statistical shape model, 38 healthy tracheas were used, corresponding to a considerable amount of shape variation. Finally, the assessment of stenoses and the prediction of stents were validated with a large set of simulation experiments and with a retrospective study with 11 CT scans of 9 patients. The results showed that the proposed system effectively reduces the observed operator dependency in the assessment of tracheal stenosis and that it is fast and accurate enough to be used in the clinical setting.

Chapter 1

Introduction

“What’s up Doc?”

Bugs Bunny

“The scientist is not a person who gives the right answers,
he’s one who asks the right questions.”

Claude Lévi-Strauss

Le Cru et le cuit, 1964

Tracheal stenosis is an unnatural narrowing of the trachea with traumatic, neoplastic, or idiopathic causes that, despite being relatively rare, can be life threatening (Spittle and McCluskey, 2000). Until now, tracheal resection surgery remains the preferred choice of treatment. In this modality, the narrowed part of the trachea is removed and the ends are sutured together (Elliott et al., 2003; Grillo et al., 1995). When the stenosis is too long or when the patient status does not permit a surgical procedure, stent implants can render a successful solution to the stricture (Venuta et al., 2004; Saito and Imamura, 2005).

Stents are tubular structures, currently made of silicone or metallic alloys, whose aim is to return normal breathing function to the patient, by pushing the walls of the stenotic area back to their original place. Since they are usually implanted with bronchoscopes, they reduce the surgical risk to the patient. Case studies of stent implants have widely been covered in the literature (Ching-Yang et al., 2007; Mandour et al., 2003; Miyazawa et al., 2004), and they are especially valuable as a palliative measure when the stenosis appears as a result of malignancy (Chin et al., 2008). Stents have also been used in tracheal reconstruction surgery in order to aid the support and recovery of the reconstructed areas (Stamenkovic et al., 2007).

Despite the well established use of tracheal stents, a pre-condition to a successful treatment is that the size and diameter of the tubes must be correctly estimated, otherwise problems such as stent migration and improper mucus clearance may occur (Venuta et al., 2004; Saito and Imamura, 2005). Stents are currently available in lengths varying by centimetres and diameters by millimetres. Prasad et al. (2002), for instance, reported a case in which the stent needed to be removed for being too narrow for the patient and suggested that a variety of stents should be ordered in advance to prepare for various contingencies.

The choice of the stent dimensions is a direct result of the assessment of tracheal stenosis. Hence, when assessing the strictures, it is important to carefully determine their location, length, and degree of narrowing (Boiselle et al., 2008). The correct assessment, in conjunction with other diagnostic tools, also aids in determining if the patient is a candidate for surgery or stent implant.

Traditionally, stenoses have been assessed with rigid or flexible bronchoscopy (Carretta et al., 2006; Nouraei et al., 2006). These methods, however, are invasive and require patient sedation. They also depend on the expertise of the specialist in charge and may not even serve their purpose if the stricture is too narrow to allow the passage of the bronchoscope itself.

With advances in the imaging field, image analysis has increasingly been used as an alternative or complementary tool for the assessment of tracheobronchial stenosis, and computer tomography (CT) is often the imaging modality of choice. A variety of methods has been proposed to assess stenosis using CT image analysis and processing. These include evaluation of multi-slice CT scans and 3D reconstructions (Lee et al., 2005) and virtual bronchoscopy (Shitrit et al., 2005; Hoppe et al., 2004), in which a virtual camera moves along the centre line of a 3D reconstruction of the airways. In fact, Shitrit et al. (2005) showed significant correlation between the use of virtual bronchoscopy and flexible bronchoscopy in grading tracheobronchial stenosis.

Evaluations of the cross-sectional diameter or area of tracheas segmented from CT images, performed using the centre line of the segmented tracheas, have also been reported (Graham et al., 2000; Sorantin et al., 2002; Kiesler et al., 2007; Triglia et al., 2002; Valdés-Cristerna et al., 2000). In this approach, the profile of the laringotracheal diameter or area is plotted on a graph. This graph eases the visualization of the decrease in the diameter of the trachea and aids the operator in determining the start and end points of the stenosis.

Although potentially valuable for the assessment of stenosis, the use of the methods above in the prediction of stent parameters is not straightforward. This is because neither of them can easily indicate how much the stent should push the narrowed walls of the trachea, since they do not give a clear indication of the tracheal geometry if stenosis were not present. The decision of which stent dimensions to use is therefore still very operator dependent.

Ballester et al. (2009) recently proposed an integrated system for tracheal surgery planning and choice of stents. Via a graphical user interface,

the operator can see the segmented trachea together with the original image and may choose a stent from a pre-built database. However, despite being a powerful visualization tool, the choice of stent and its appropriate positioning in the image are completely manual.

Stenosis assessment and stent choice have also been a topic of interest in the cardiovascular community. F. Valencia et al. (2002) proposed a graphical model for vascular stent pose simulation. The model represents a deformable cylindrical surface that responds to an energy minimization function controlling its bending and how it is attracted to artery boundaries in image data. van Bemmelen et al. (2004) employed level-sets for the segmentation of the stenotic artery and manually chose a healthy section of the artery as a reference to compute the level of stenosis.

Concomitantly with developments in the search for treatments for tracheal stenosis, Active Contour Models (ACMs, or snakes) and Active Shape Models (ASMs) have been important tools in computer aided diagnoses. The former is used to detect contours in an image by minimising an energy function controlling the bending and stretching of the contour and how it is attracted by image features (Kass et al., 1988). The latter type of models capture statistical shape variations from a set of training shapes and can be registered to clinical image data by simple adjustment of their parameters. In addition, a distinctive characteristic of ASMs is that they only generate shapes that resemble those in the training set. This makes them more robust to noise and occlusion in comparison with other deformable models and free-form registration techniques (Cootes et al., 1995).

1.1 The Proposed System

Having the above concepts in mind, this thesis sets forth a decision support system that proposes a method for assessment of tracheal stenosis and prediction of stent length and diameter. The main idea behind this method is to estimate the shape of the trachea of a patient as if stenosis were not present. To this end, an ASM built solely with healthy tracheas is registered to the CT scan of that patient. The expected result is that the registration

yields a shape that matches the healthy areas of the trachea, while at the same time avoids the influence of the narrowed regions. Since, as mentioned above, an ASM only generates shapes that resemble those in its training set, the local geometric variations of the strictures, typical of stenosis, are avoided.

Despite the motivation to use an ASM with healthy tracheas in order to avoid the local variations of the stenosis, the strictures may have global influence on the registration of the ASM to the image. Since variations in calibre are usually present in the model, the strictures can eventually make the resulting shape globally much narrower than desired. For this reason, a novel *Robust ASM Fitting* procedure, called *Fixed Landmarks*, is proposed. The idea behind the *Fixed Landmarks* is to reinforce the resistance of the registration to the attraction incurred by the narrowed parts of the trachea. This is achieved by keeping the parts of the estimated trachea corresponding to regions with stenosis fixed at each iteration, relative to their location at the previous iteration. In this way, not only are the narrowed regions avoided, but the fixed areas incur an extra repelling force that will aid in the estimation of the healthy tracheal wall over those regions.

When the ASM registration is finished, it is thus expected that the resulting shape be estimation of the patient's trachea as if stenosis were not present. In order to assess the extent and degree of stenosis, however, it is still necessary to segment the narrowed trachea from the image. This is achieved using an ACM tailored for this purpose. Using the estimated healthy trachea as a starting point, the ACM will deform it until the boundary of the narrowed trachea is completely detected. The internal and external forces of the ACM were therefore specifically designed to let the model yield the best possible contour delineation with the lowest possible shape deformation.

Once healthy and narrowed versions of the trachea are obtained, the eventual assessment of the stenosis and prediction of the stent is straightforward. The proposed method simply compares the cross sectional areas of the two surfaces along their centre lines and detects significant reductions. These determine the start and end points of the stenosis and its degree of narrowing. Such parameters are then used in another algorithm, which

yields the stent to be chosen for the patient.

A couple of extra challenges needed to be solved before producing the desired results. The first was how to efficiently manipulate large medical image files. With current technology, and depending on the scanning resolution, CT image volumes of the chest can be stored in files whose sizes are in the order of gigabytes. When these files needed to be manipulated by the various algorithms used in this work, memory restrictions often arose. In particular, many of the methods proposed depend on pre-processing tasks that involve basic image processing algorithms, such as edge detection, image arithmetic, region growing, etc. As a result, this work also proposes an efficient *out-of-core* strategy to manipulate these large files in image processing tasks. The idea is to keep most of the file in disk (hence the term *out-of-core*) and take advantage of the access pattern of the algorithm through the image when loading data into main memory. In this way, an efficient cache and prefetching strategy is built in order to keep the relevant data in memory and load parts of the image before they are requested by the application.

Finally, in order to build the ASM, samples of healthy tracheas are necessary. These are usually obtained through segmentation and a comprehensive database of chest CT scans of healthy subjects was used in this work. Although only the trachea needs to be segmented for the presented application, this thesis proposes a method to segment the entire intrathoracic airway tree. This tree is the hierarchical tubular structure below the vocal cords, comprising trachea and bronchi, that leads the air into the lungs. The proposed method can easily be configured to segment only the trachea and generates the segmented surface along with its centre line in one go. These centre lines are also useful for the construction of the ASM.

An extensive set of simulations for all proposed methods was carried out in order to perform their validation. For the *out-of-core* technique, different algorithms and cache structures were tested on large files. The airway segmentation method was tested on a database of 40 image volumes, subdivided into training and testing sets. For the ASM, 38 healthy tracheas were used, corresponding to a considerable amount of shape variation. Finally, the assessment of stenosis and prediction of stents was validated with a

large set of simulation experiments and with a retrospective study with 11 CT scans of 9 patients.

1.2 Thesis Organization

To present the proposed methods and algorithms, this thesis has been subdivided as follows:

- Part I introduces the reader to the anatomy of the trachea and of the airways. It also discusses the main characteristics of tracheal stenosis and of airway stents. Finally, traditional and computer-aided methods for the assessment and stenting of tracheal stenosis are reviewed.
- Part II presents all the algorithms proposed in this thesis. Chapter 4 describes the *out-of-core* technique and how the proposed cache and prefetching strategies can be applied to a range of image processing algorithms. Chapter 5 presents the region growing algorithm used in the segmentation of the intrathoracic airway tree. Next, Chapter 6 describes the algorithm used to estimate the patient's trachea if stenosis were not present. It also describes the proposed *Fixed Landmarks* approach to avoid the influence of the narrowed regions of the trachea in the ASM registration. A comparison between the *Fixed Landmarks* and other techniques is also shown. In Chapter 7, the narrowed tracheas are segmented with a method that uses the aforementioned tailor-made ACM. Chapter 8 finally presents how the stenosis and stents are computed from the healthy and narrowed versions of the trachea. An extensive set of experiments with simulation and clinical data was carried out to validate the proposed method.
- Part III terminates this thesis with a summary and directions for future work, in Chapter 9 and Chapter 10, respectively.

Part I

Preliminaries

Chapter 2

Clinical Background

Parts of this chapter appear in:

R. Pinho, K. G. Tournoy, and J. Sijbers. Computer-aided assessment and stenting of tracheal stenosis. In Ayman El-Baz, editor, *Computer Aided Diagnosis of Lung Imaging*, volume (in press). American Scientific Press, 2011.

2.1 Introduction

In order to fully understand the purpose of the airway related algorithms presented in this thesis, the reader should have some familiarity with the anatomy the airways, especially of the trachea. It is also important to understand the main characteristics of tracheal stenosis and stents, including the challenges involved in their choice. This chapter thus provides the reader with an overview of all these topics.

2.2 Anatomy of the Lower Airways

The lower (or infraglottic) airways correspond to the tracheobronchial tree below the vocal cords down to the alveoli. It starts to form around the fourth week of gestation. In a normal individual, the trachea enters the thorax at 1 to 3 centimetres above the sternal notch. It bifurcates at the level of the sternal angle into the right and left main bronchi. The right bronchus is wider and shorter and tends to be more vertically placed relative to the left bronchus. They bifurcate to form the lobar (or secondary) bronchi. Three lobar bronchi appear in the right lung, whereas only two appear in the left lung. Each of these bronchi further subdivides into the tertiary (or segmental) bronchi, from where another 20 to 25 generations of progressively smaller branches are found (Chmura et al., 2008). The airway geometry can be described by a recursively inter-related series of branch lengths and angles (Hegedus; Sauret et al., 2002). Figure 2.1 illustrates the airway structure with the segmental labelling proposed by Ikeda (1974).




Figure 2.1: Airway structure (Ikeda, 1974).

2.3 Anatomy of the Trachea

The trachea is the primary tube of the tracheobronchial tree. It functions as a conduit for ventilation and clearance of secretions.

The shape of the healthy trachea may present considerable variation. With age, it gradually goes from funnel-shaped in infancy to ovoid in adulthood. The radial shape is normally ovoid anteroposteriorly, but can also appear ovoid in the sagittal direction, be circular or, less frequently, slightly triangular. Unique and unexplained distortions may occur as well (Grillo et al., 1995). The length goes from 9 to 12 centimetres in adults, measured from the lower border of cricoid cartilage down to the top of the carinal spur, while the average laterolateral diameter is between 1.3 and 2.2 centimetres (Webb et al., 2000; Grillo et al., 1995).

The trachea is comprised of 18 to 22 horseshoe-shaped cartilaginous open-circles anteriorly located, whose ends are posteriorly connected by a soft membrane. It bifurcates approximately at the sternal angle into the left and right main bronchi, and the carina is the ridge separating their openings. On CT, the tracheal wall is usually visible as a 1- to 3-millimetre soft tissue stripe, delineated internally by air in the tracheal lumen and externally by mediastinal fat, the lungs, larger vessels, and the oesophagus. The posterior tracheal wall appears thinner and more variable in contour, since it lacks cartilage. It can appear convex, concave, or flat. During expiration, CT typically shows significant anterior bulging of the posterior membrane, but the anterior and lateral walls change little (Webb et al., 2000). Figure 2.2 provides illustrations of the tracheal anatomy.

2.4 Tracheal Stenosis

Tracheobronchial stenosis is defined as focal or diffuse narrowing of the tracheal lumen (Figure 2.3). By far, the most common cause of tracheal stenosis is trauma, specifically internal trauma, usually occasioned by any type of intubation (Grillo, 2004; Boiselle et al., 2008; Lee and Boiselle, 2008). There are other potential causes of tracheal stenosis, including inflammatory dis-

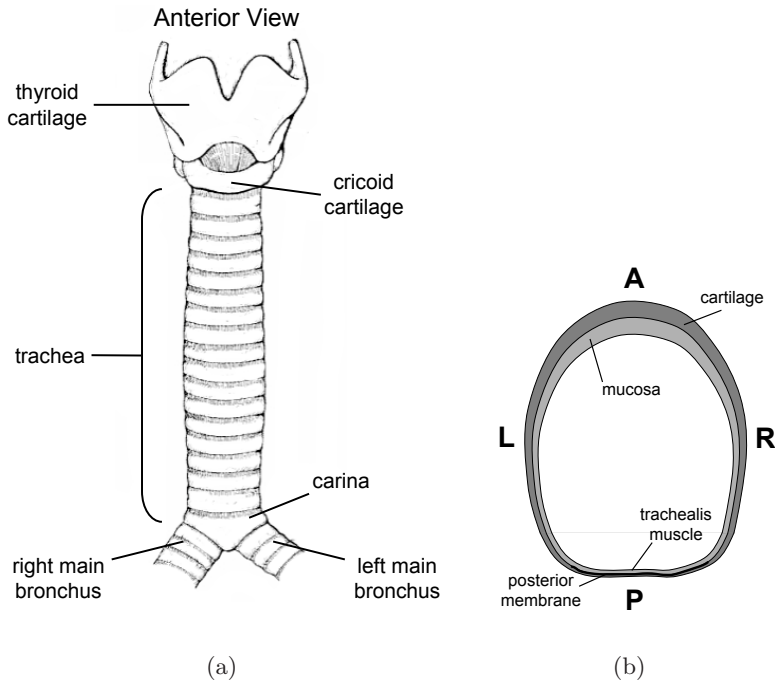


Figure 2.2: Different views of the trachea and its composing structures. (a) Illustration of the trachea and main bronchi. (b) Schematic view of an axial cross-section.

eases, benign or malignant neoplastic conditions, and extrinsic pressure. In addition, idiopathic narrowing of the trachea is not completely uncommon, predominantly appearing in women (Grillo et al., 1995).

Typical symptoms of tracheal stenosis include shortness of breath, coughing, and stridor. The narrowing may be seen in various different shapes. It typically extends from 1.5 to 2.5 centimetres in length (Webb et al., 2000), but cases of long-segment tracheal stenosis have also been reported (Elliott et al., 2003). It has also been reported that the shape of the stricture is usually associated with its causes (Lee and Boiselle, 2008; Webb et al.,


2000).

2.5 Tracheal Stents

The medical term “stent” was introduced by Charles R. Stent, a British dentist who developed a device that supported the healing of gum grafts. The term has since been used to refer to any device designed to maintain the integrity of hollow tubular structures (Saito and Imamura, 2005; Freitag, 2000). Tracheal stents are tubular structures, currently made of silicone or metallic alloys, whose aim is to return normal breathing function to the patient by pushing the narrowed regions of the trachea (Ching-Yang et al., 2007; Mandour et al., 2003; Miyazawa et al., 2004). Examples of the different materials, shapes, and sizes of coated and uncoated stents are shown in Figure 2.4

Although airway resection and reconstruction is most often the preferred therapy for both benign and malignant lesions, a variety of factors, including long stenosis, failed previous repair, metastatic or unresectable malignancy, or even patient refusal may dictate non-surgical management. In these cases, bronchoscopic treatment is indicated, and the use of stents is also an option. There are three indications for the use of stents: support of weakened tracheal walls due to tracheal malacia, re-establishment of the full tracheal lumen in case of extrinsic pressure, and sealing of fistulas towards the oesophagus or pleural cavity (Freitag, 2010; Venuta et al., 2004; Saito and Imamura, 2005; Stephens and Wood, 2000; Ching-Yang et al., 2007; Mandour et al., 2003; Miyazawa et al., 2004).

Good results from airway stenting largely depend on patient selection, but satisfactory to excellent results are generally achieved in most patients undergoing airway stenting (Saito and Imamura, 2005). Possibly, the largest study on the use of tracheal stents was that performed by Wood et al. (2003). They reported the results obtained in 631 interventional bronchoscopy procedures carried out in 332 patients during the course of ten years (May, 1992 to December, 2001). Out of the 631, 309 stenting procedures (49%) were conducted in 143 patients (43%). A complication



(a)



(b)

Figure 2.3: (a) A schematic view of stenosis (by courtesy of Dr. P. Delaere, Center for Larynx, Trachea & Hypopharynx Reconstruction, Katholieke Universiteit Leuven, Belgium). (b) A real case showing one sagittal slice of a chest CT scan.



Figure 2.4: Typical tracheal stents. Top row, from left to right: Montgomery T stent, Dumon stent, Polyflex stent and Noppen stent. Bottom row, from left to right: Ultraflex stent, Aero stent, Eco Nanjing stent, Hanaro stent and Tae woong stent. (Freitag, 2010, copyright: ERS, 2010)

rate of 42% was reported, including stent migration, occlusion of the stent by secretions, and partial occlusion by granulations. However, 95% of the patients noted significant improvement of their symptoms after the interventions. Yet, complete patient follow-up was not possible in most cases. According to the authors, it is difficult to determine, without 100% follow-up at the endoscopy center, whether patients failing to return are because of excellent results, poor results with an unwillingness to return for further intervention, or because of death from progression of the underlying disease. In another, recent study, covered self-expandable metallic stents had their effectiveness evaluated in a total of 17 patients. In 21 stenting procedures, the complication rate was 75%, deeming such types of stents to be abandoned in the clinical routine of the authors of the study (Dooms et al., 2009).

Stents have also been used in tracheal reconstruction surgeries to aid the support and recovery of the reconstructed areas (Stamenkovic et al., 2007). In addition, stents are valuable complementary tools to other healing techniques and can provide prolonged palliation. More important, it is the only bronchoscopic technique to provide a solution for extrinsic compression (Wood et al., 2003).

The size of the stent should be adjusted according to the normal diameter proximal and distal to the area of stenosis. A narrow stent risks displacement and will be separated from the tracheal wall with failure of incorporation into the mucosa. A wider stent may lead to wall necrosis or excessive elongation with proximal or distal airway obstruction. Stents that are too long, in turn, can impede mucus clearance. Stent choice thus depends on the characteristics of the stenosis and on a physician's training and experience. Devices such as the AeroSizer[®] (Merit Medical Systems, Inc., USA) have even been created to aid in the stent sizing process. A rule of thumb is that the stent should extend 0.5cm at each end of the stenosis (Freitag, 2010; Venuta et al., 2004; Graham et al., 2000; Ching-Yang et al., 2007; Grillo, 2004; Wood et al., 2003; Mostafa, 2003). Lee et al. (2010) reported that a stent alert card detailing the type and dimensions of the stent and its location in the tracheobronchial tree should be given to the patient. It should also indicate the appropriate size of endotracheal tube

to be used if emergent intubation is required with the stent in site.

In the clinical routine experienced during the development of this thesis, the healthy diameter of a patient's trachea is guessed from the images available. Image sources include bronchoscopy, CT, and 3D reconstructions. The start and end points and the degree of narrowing of the stenosis are then visually estimated from this guess. In general, stents can be the treatment of choice for strictures with a degree of narrowing between 25% and 75%.

Chapter 3

Related Work

Parts of this chapter appear in:

R. Pinho, K. G. Tournoy, and J. Sijbers. Computer-aided assessment and stenting of tracheal stenosis. In Ayman El-Baz, editor, *Computer Aided Diagnosis of Lung Imaging*, volume (in press). American Scientific Press, 2011.

3.1 Introduction

Since the introduction of the rigid bronchoscope in 1897, the technology in the field of airway imaging has greatly evolved. The flexible bronchoscope immediately gained popularity and became the most common method for airway visualization. With the advent of CT and other imaging modalities, complementary tools have emerged to reduce the subjectivity and operator dependency of airway assessment. In the case of stent choice, in particular, the new techniques have proven to be an important non-invasive alternative to the traditional bronchoscopic methods. This chapter reviews the main characteristics of the traditional rigid and flexible bronchoscopies and presents a short survey on the existing image-based and computer-aided techniques for airway assessment and, especially, the treatment of tracheal stenosis.

3.2 Traditional Methods for Airway Assessment

3.2.1 Rigid Bronchoscopy

The technique and clinical application of bronchoscopy was introduced by the German doctor Gustav Killian in 1897, when he first employed an endoscope to examine the airways. His tool was in fact an oesophagoscope, which he used to remove a piece of pork bone stuck in the right main bronchus of a 63-year-old man (Becker and Marsh, 2000). Figure 3.1(a) shows a picture of Dr. Killian and his pioneering technique.




Figure 3.1: (a) Dr. Killian and the rigid bronchoscope. (b) Flexible bronchoscopy. (both figures from the public domain <http://en.wikipedia.org>).

A rigid bronchoscope is thus a straight, hollow, metal tube, through which a rigid fibre-optic device is placed, providing access to the central airways. The tube can only be inserted through the mouth and requires general anaesthesia. It was the standard procedure for airway assessment until the 1970's, when the flexible bronchoscope was introduced. Today, it remains the procedure of choice for the removal of foreign objects from the airways or when a larger biopsy is needed. The rigid bronchoscope is also essential for the insertion of stents, since it provides the necessary working channel for the insertion and removal of the tubes (Prakash, 1999).

3.2.2 Flexible Bronchoscopy

In the late 1960's, Dr. Shigeto Ikeda introduced the flexible bronchoscope, which aimed to reach further segments of the airway and to have a constant light source at the distal tip of the tube to enhance illumination of the visited areas. Later, in the late 1980's, Dr. Ikeda was also the one to introduce the videobronchoscope, which projected magnified images of the airways on a display (Miyazawa, 2000). The flexible bronchoscope became

common since its introduction and today it is the most used invasive technique in the study of the airways (Prakash, 1999). Figure 3.1(b) shows the flexible bronchoscope in use.

The flexible bronchoscope is a thin flexible device that carries an optic fibre to conduct light from an external light source. It also contains an optical system to enable the visualization of the airway structures through an eyepiece or an image display. Auxiliary tools can also be inserted through the tube, commonly used for the collection of samples and therapy. Besides the ability of reaching further areas of the airways, one advantage of the flexible bronchoscope over the rigid bronchoscope is that it generally requires only local anaesthesia, and mild patient sedation if necessary (Prakash, 1999; Miyazawa, 2000).

As an example of the employment of flexible videobronchoscopy in the assessment of the airways, Czaja et al. (2007) used a computer program through which the recorded images could be measured. A probe with known dimensions was used in the calibration of the program and the areas of interest could then be measured by manual delineation.

3.3 Computer-Aided Methods

Despite the fact that the traditional methods above are still the gold standard in the management of tracheal stenosis, they do have limitations. Stephens and Wood (2000), for instance, reported cases in which flexible bronchoscopes were avoided in the evaluation of critically narrowed airways because of the potential for occlusion of the narrowed tracheal lumen. Furthermore, as mentioned in the previous section, both flexible and rigid bronchoscopy often require patient sedation and are invasive procedures by nature.

For these reasons, and followed by advances in imaging technology, non-invasive computer-aided methods for assessment of stenosis have been developed to substitute or complement the traditional techniques. Next, a number of different computer-aided methods are described and discussed.

3.3.1 Manual Methods

In the category of manual methods, the physician directly interacts with the images acquired for the patient evaluation. No computer methods are employed to aid the physician in assessing the stenosis or choosing a stent if one is to be used. Although radiography can still be used as an imaging modality, volume imaging is the preferred choice, with CT being the most often modality used.

Williamson et al. (2010) and McLaughlin et al. (2008) employed a novel imaging technique called anatomic optical coherence tomography (aOCT) in the measurement of the tracheal calibre. Optical coherence tomography (OCT) is a light-based imaging technique in which receptors capture how much the light emitted from a source has been absorbed or scattered inside the analysed tissue. aOCT is a modification of conventional OCT designed to allow macroscopic imaging of hollow organs. Light is emitted from a very thin probe ($\approx 2\text{mm}$) passing through a flexible bronchoscope inside the trachea. The probe rotates at around 2.5Hz, tracing an axial image of the airway. Longitudinal displacement of the probe inside the trachea eventually yields 3D reconstructions. An advantage of aOCT over other imaging techniques is that it can be used for long periods, enabling physiological as well as anatomical analyses.

Callanan et al. (1997) were the first to use magnetic resonance imaging to assess tracheal stenosis following tracheotomy, in a study with 18 patients. The diameter of the trachea was measured in sagittal images at five separate intervals from the cricoid. If stenosis was noted, the narrowing was quantified as the percentage of reduction in the diameter relative to the regions immediately above and below the stenotic segment. The authors claimed that MRI was indeed an interesting imaging modality for the assessment of stenosis, providing good image quality and, compared to CT, with the advantage of not subjecting the patient to ionising radiation.

Despite the apparent advantages of MRI, the superior resolution and contrast of CT images and the much faster scanning times make CT the preferred imaging modality for tracheal and airway analysis. Lee et al. (2005), for example, used CT image analysis in the pre- and post-stent placement

settings. They emphasized how multi-slice CT provides invaluable information in the initial evaluation of patients presenting airway obstruction, being an important complementary tool to bronchoscopy. They further showed that CT aids planning of stent implants by clearly delineating the anatomy, pathology, and severity of the airway obstructions. As an added benefit, CT images show the relationship between the central airways and the adjacent structures, which are not visible with the bronchoscope. In the post-stent placement procedures, multi-slice CT proved to enable accurate detection of stent complications, including migration, size discrepancies, and fracture. The authors also referred to other successful cases reported in the literature.

More recently, Parida and Gupta (2008) demonstrated good correlation in assessment of stenosis between spiral CT with 3D reconstructions (SCT-3DI) and bronchoscopy. In this imaging modality, a 3D model of the trachea is generated from the CT data. This improves over traditional multi-slice CT analysis, providing the physician with volumetric information. The authors also reported that in 16 out of 30 cases, bronchoscopy could not measure the length of the stenosis due to severe narrowing. With SCT-3DI, all cases could be measured.

Another improvement obtained with image analysis is the use of virtual bronchoscopy, in which a virtual camera is controlled by the physician through a 3D reconstruction of the airways obtained from CT. In this way, bronchoscopy is simulated in a completely virtual environment. Such a technique has the advantage of being similar in procedure to rigid or flexible bronchoscopy, with which the physician tends to be familiar, but without being invasive. For this reason, virtual bronchoscopy is often used in training. Ferguson and McLennan (2005) reviewed several application scenarios of virtual bronchoscopy, demonstrating current and future trends for the procedure. They described, for instance, how images captured with the bronchoscope can be correlated with those of virtual bronchoscopy. This augments the physician's view and enables the visualization of structures that would otherwise be hidden when using only bronchoscopy. In airway analysis, Shitrit et al. (2005); Triglia et al. (2002) showed significant correlation between results obtained with virtual bronchoscopy and flexible

bronchoscopy in the assessment of tracheal stenosis. Hoppe et al. (2004) also demonstrated that virtual bronchoscopy can even be used to grade stenosis of the airways down to the segmental levels.

Finally, Ballester et al. (2009) recently proposed an interesting system for tracheal surgery planning and choice of stents. A graphical user interface enables the user to visualize the CT data, 3D reconstructions of the trachea, and volume rendered images altogether. The stents are then chosen from a pre-built database containing CAD models of commercially available tubes. These are manually overlaid on the images so that the physicians can have a better idea of how the stent will interact with the tracheal wall and neighbouring regions. The 3D reconstructions of trachea were also used for biomechanical airflow simulation studies before and after stent insertion.

All these methods have certainly proved the benefits of image analysis in the assessment of the airways and especially of tracheal stenosis. Yet, they are all still very reliant on the expertise of the operator in charge. The next section reviews methods that incorporate sophisticated algorithms aiming to reduce this operator dependency.

3.3.2 Semi-automatic Methods

To the best of our knowledge, all semi-automatic algorithms devoted to the assessment of tracheal stenosis share the same characteristic. Namely, their output is a graph or some other interface to convey to the clinician the variation in the cross-sectional diameter or area of the trachea. What is different between these algorithms is how they achieve such results, which may involve more or less user-intervention.

The first step in this process is the segmentation of the trachea from the CT image data and, possibly, its conversion to a 3D surface. Hertel et al. (2006) developed a system for the segmentation and 3D visualization of the trachea, employing techniques such as live-wire, region growing, and watershed. Their aim was the planning of tracheal resection surgery, and all the measurements related to the assessment of the stenosis were carried out manually. They reported that the average time necessary for the segmentation of the trachea was 35 minutes.

Graham et al. (2000) presented a study with 8 patients in which several image based techniques were employed to assess and measure the airways. They employed the software presented in (D'Souza et al., 1996; Hoffman et al., 1992) to segment the airways and extract their centre lines. In the software, segmentation is semi-automatically performed on a slice-by-slice basis: rays cast from a user-defined point near the barycentre of the airway lumen determine the location of inner and outer airway walls. This information is also used to segment the lumen with region growing.

Triglia et al. (2002) also segmented the trachea on a slice-by-slice basis, but used 2D active contours (Kass et al., 1988) to delineate the axial boundary between the tracheal wall and the lumen. In their approach, user intervention is potentially reduced by selecting only the first and last slices of the trachea, segmenting the first slice, and letting the algorithm use the result of each slice as the starting point for the next. After the segmentation, new contours perpendicular to the surface's centre line are computed and used in the definition of a smooth B-spline surface. The operator then marks on the centre line the start and end of the stenosis, which are used to compute the length of the narrowing. The degree of narrowing is finally computed as the ratio of the cross-sectional area of the selected tracheal segment to the area of a reference healthy section also selected by the operator. The area or diameter profile of the entire trachea can further be displayed on a graph, accompanied by a visualization of the 3D surface.

In (Valdés-Cristerna et al., 2000; Valdés-Cristerna and Yáñez-Suárez, 2003), a 3D contour model was employed in the segmentation of the trachea. This improves over (Triglia et al., 2002) by eliminating the surface generation step, since the 3D model already tends to yield a smooth 3D surface by definition. Triglia et al. (2002) also report difficulties in the slice-by-slice segmentation whenever the narrowing was too severe and the tracheal lumen was barely visible in the image. In these cases, the missing tracheal section had to be reconstructed manually, which can obviously introduce errors. Although the 3D contour model can potentially overcome this problem, Valdés-Cristerna et al. (2000); Valdés-Cristerna and Yáñez-Suárez (2003) did not report any result related to such situation. The tracheal area profile was once again plotted on a graph, but details on this

process were not given either.

In the method presented in (Sorantin et al., 2002, 2006) the cross-sectional area profile of the trachea is obtained by first segmenting the trachea with a region growing method using fuzzy connectedness (Udupa and Samarasekera, 1996). Next, a thinning algorithm (Palágyi and Kuba, 1998) is applied to the segmentation result to extract the trachea's centre line. Both steps may require post-processing by user intervention. As in the other discussed approaches, the tracheal area profile is plotted on a graph, and the start and end of the stenosis are determined by the operator based on this graph and its relationship with the 3D reconstruction of the trachea. The degree of narrowing is chosen as a percentage of the minimum cross-sectional area relative to the area of the trachea at the begin and end points of the stenosis. Kiesler et al. (2007) later used this method for the assessment of 3 patients.

The improvements brought about by all these techniques over the manual methods are undeniable. Nonetheless, the methods above share another common characteristic: the analysis of the result graph and the determination of the parameters of the stenosis are still, to a large extent, operator dependent. Although the visualization of the tracheal area profile on a graph does ease the assessment, it is the operator that decides where the stricture starts and where it ends, and the degree of narrowing is eventually computed from these parameters. In addition, in the case of stent choice, it is not completely clear from the graphs how much the stent should push the tracheal wall, that is, how to decide the appropriate stent diameter to be employed. In the end, it is the physician that still needs to guess the healthy tracheal shape from the available results in order to choose the correct stent.

The next part of this thesis introduces a new method for the assessment of stenosis based on deformable models of shape. The proposed method estimates the shape of the patient's trachea in case stenosis were not present. An added benefit of this method is the automatic prediction of a patient-specific stent.

Part II

A System for
Assessment and Stenting of
Tracheal Stenosis

Chapter 4

Out-of-core Image Processing

“The stars don’t tell the future, Donkey. They tell stories.
That one is Bloodnut, the Flatulent.
You can guess what he’s famous for.”
Shrek

“Research is to see what everybody else has seen,
and to think what nobody else has thought.”
Albert Szent-Gyrgi (1893 – 1986)

Parts of this chapter appear in:

R. Pinho, K. J. Batenburg, and J. Sijbers. Seeing through the window: Pre-fetching strategies for out-of-core image processing algorithms. In *Proceedings of SPIE Medical Imaging*, volume 6919, page 69190D, San Diego, CA, USA, 2008. SPIE.

R. Pinho, K. J. Batenburg, and J. Sijbers. Windowcache: Pre-fetching strategies for out-of-core image processing algorithms. *Image Processing, IEEE Transactions on*, (Under Review), 2010.

4.1 Introduction

The first challenge in the development of this thesis was the manipulation of large CT image files in image processing (IP) applications. A typical scenario in this type of applications is a pipeline of processing units, such as the one depicted in Figure 4.1. The first unit receives as input the original image buffer, generates the result (possibly in a new buffer) and passes it on to the next unit. The second unit receives the buffer, processes it, and hands it over to the next. This process is repeated until the last unit produces the final image.




Figure 4.1: Typical image processing pipeline.

Considering such a pipeline, IP algorithms often require the original image and one or several copies of it to be kept in main memory (RAM). A problem may thus arise if the processed image buffers are too big. With current technology, image acquisition in various fields can yield image files of several gigabytes, or even terabytes. Even optimized IP pipelines that

eliminate intermediate image copies commonly need at least the original image and a copy of it in memory to store the final result.

Although the advent of 64-bit desktop computers has largely extended the theoretical memory size limit, hardware manufactures and operating systems still fall far short of reaching it. One solution to the memory restriction problem is to keep the data in secondary storage during the execution of the algorithms, known as *out-of-core* processing. Secondary storage is virtually unlimited and its cost per megabyte is considerably lower than primary memory, making it a cost-effective solution in terms of expansion. They are used for instance to efficiently render gigantic 3D scenes or objects (Nielsen et al., 2007; Farias and Silva, 2001; Bruckschen et al., 2001) and to provide interactive visualization of large scientific datasets (Wang et al., 2007; Bruckner, 2004). For interesting surveys on the subject, please see (Silva et al.; Vitter, 2001)

The biggest challenge in out-of-core implementations is to overcome the still existing gap in speed between fast primary memories and slow secondary memory access. The slower speeds of disks are mostly evidenced in terms of high latency due to relatively long disk seek times. Another problem with accessing the image on disk is that file layouts are inherently unidimensional. Even a multidimensional structure, such as that of 3D medical image files, is stored as a unidimensional sequence of data. Traversing the image in directions that differ from the file layout causes jumps in the file access and can immensely increase latency. Virtual memory schemes offered by operating systems do not reduce this latency, since they do not take into account the multidimensionality of the data (Caron et al., 1999).

A common feature of out-of-core algorithms in the computer graphics and imaging fields is that they are implemented to match the data access pattern of the application (Nielsen et al., 2007; Bruckschen et al., 2001; Wang et al., 2007; Prohaska et al., 2004). Moreover, they try to reduce disk latency by increasing the size of image blocks read in a single operation. To this end, many researches proposed to preprocess and reorganize the original file in such a way that its new layout matches the access pattern of the processing algorithm and that large image blocks can be read

at once (Kohlmann et al., 2007; R. et al., 2005; Agrawal et al., 2010). Although improvements in file access were shown, the disadvantage of these approaches is the waste of disk space and possible long copying times prior to the execution of the algorithms. Michikawa et al. (2007) proposed a different approach, specific for out-of-core Euclidean distance transforms, in which the original volume is processed block-wise, and an extra step takes care of computing the distance function at the borders between blocks.

Whichever the strategy adopted, the behaviour of out-of-core algorithms intuitively resembles that of a cache. Cucchiara and Piccardi (1998); Cucchiara et al. (2004) indeed showed that basic IP algorithms have a very predictable behaviour and can greatly benefit from CPU-level cache and prefetching strategies due to their high level of temporal and spatial locality. Rhodes et al. (2005a,b) tackled out-of-core volume visualization with a cache which operates on the original file layout and is updated asynchronously according to the access pattern of the visualization algorithm. Their *iterator aware cache* used a similar principle to the *one-block look ahead* strategy presented in Cucchiara et al. (2004) to prefetch the next image block to be read from disk. Lipsa et al. (2007) used another cache structure for visualization, which simulates a block-wise subdivision of the data. However, their cache was updated synchronously, upon a cache miss event.

Despite all the good work presented for visualizations, to the best of our knowledge, no general solution for out-of-core IP has yet been developed, especially in the field of medical imaging. Tian et al. (2008) recently addressed this issue and analysed existing libraries for visualization and medical IP. They proposed a new platform that supports out-of-core processing by design, using clever software design techniques that separate from the implementer of an algorithm the burden of taking care of memory management. However, the authors did not propose any specific out-of-core IP strategy, only giving indications that native operating system resources such as memory mapped files and virtual memory are used. As discussed above, these resources do not take into account the multidimensionality of the data and may not be the most efficient approach.

To fill the existing gap, the present work sets forth *WindowCache*, a

read-write cache with an asynchronous local prefetching strategy based on a sliding window protocol. *WindowCache* operates on the original file layout and takes into account a priori knowledge about the access pattern of the IP algorithm in order to predict future data to be read. The idea is that the window structure and the sliding protocol can be adapted to different types of access patterns. In this paper, three types of algorithms are studied, namely, raster scan algorithms, convolution algorithms, and propagative algorithms, which cover a large proportion of basic IP tasks. A window structure and a prefetching strategy are proposed for each of them. The focus here is on medical imaging processing applications, and an in-depth analysis of *WindowCache* through experiments with large medical image files is performed on machines with restricted memory sizes.

The remainder of this chapter is subdivided as follows. In Section 4.2, *WindowCache* is described, and it is also shown how the cache structure and window sliding protocol can be adapted to different IP algorithms. Section 4.3 is dedicated to the experiments and their results. Section 4.5 discusses the results obtained, and the chapter is concluded in Section 4.6.

4.2 *WindowCache*

4.2.1 The Cache Buffer

WindowCache is internally organized as a set of N_B image blocks. The blocks of the cache have a direct correspondence with a virtual subdivision of the d -dimensional image into regions, where $d > 1$. It simulates a subdivision in which the original image file is reformatted in chunks of equal size and each chunk, or block, is referenced by a unique identifier b_i . Figures 4.2.a, 4.2.b, and 4.2.c show an example of how the subdivision of the image space can be made for an image with $d = 3$.

The blocks of the image are read into the cache buffer upon request. Whenever the buffer is full, *WindowCache* uses the least recently used (LRU) protocol (Hennessy and Patterson, 1990) for block replacement. For efficiency purposes, the copy-back writing protocol is adopted (Hennessy and Patterson, 1990), and the blocks of the cache are organized in a




Figure 4.2: Illustration of the cache structure for a 3D image volume. (a) A 3D image with its dimensions. (b) Virtual subdivision of the image. Each block is referenced by a unique identifier. (c) Cache (or image) block and its dimensions.

balanced binary tree, sorted by their unique identifier. In this way, block searches are executed in $O(\log_2 N)$ time.

4.2.2 Prefetching

Prefetching is part of the process through which the cache buffer is filled. The idea is to pre-load from disk some of the data before they are requested by the main application. The scheme of Figure 4.3 is a generic representation of the prefetching strategy of *WindowCache* and shows how it behaves during the execution of an IP algorithm. There are 4 entities represented in this scheme:

- Main Program: the application that executes the IP algorithm.
- Image: an abstract representation of the image that is being processed. It is abstract in the sense that it is not directly connected to the image's physical representation, in main memory or in disk.
- Cache: the cache system, where the image data is actually stored.
- Parallel Prefetch: controlled by the Cache, it is responsible for updating the cache data.

The Main Program runs the IP algorithm, synchronously requesting pixels from the Image. Since the Image has no actual data, it just forwards the requests to the Cache, which knows if the requested pixel is available in main memory or if it should be read from disk. If the pixel is available, the Cache promptly returns it to the Main Program. Since *WindowCache* is based on a priori knowledge of the access pattern, the Cache determines the best moment to trigger the read of the next image block from disk. If the Cache is full, one of its blocks is replaced with the image block read from the disk, but it is first flushed to disk if modified, according to the copy-back protocol. The (flush, read) pair is executed in the Parallel Prefetch entity, representing an asynchronous thread of execution. As a result, this execution does not affect the Main Program.

The key point of *WindowCache* is that updating the cache is always an asynchronous operation, performed in a separate thread of execution. This reduces the probability that the application will stall while a cache block is being filled with data from disk. In addition, depending on the predictability of the access pattern, a block can be read from disk far ahead of its use, while others that will no longer be used can quickly be discarded.

The prefetching also benefits from the complexity of the algorithm. It means that the longer the Main Program takes to process a pixel, the higher is the probability that a prefetching operation will complete before the associated image block is needed.

To further clarify the concepts presented above, Algorithm 1 details the scheme of Figure 4.3, showing what happens inside the Cache in the event of a pixel request. The Cache first checks if the requested pixel resides in the cache block accessed in the previous request, i.e., the current block (step 2). The block's unique identifier can be computed from the pixel coordinates. If the pixel is in the current block, the Cache just returns it, synchronously. Otherwise, the block corresponding to the pixel, b_i , is searched within the cache buffer. If the block is not found, it is loaded from disk and inserted into the cache buffer. If the buffer already has N_B blocks, the LRU block is replaced. Finally, one or more block prefetches may be triggered (step 15), according to the access pattern, which, once again, is known by the Cache. This step, referred to as local prefetching, will be further detailed

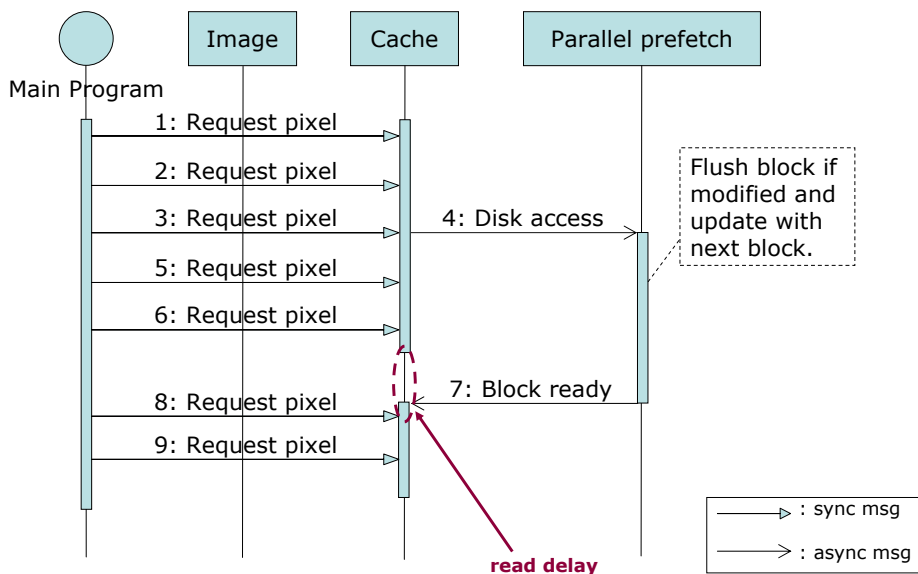


Figure 4.3: Sequence diagram of the prefetching process. The main program synchronously requests pixels, and the cache system asynchronously decides when to load another image block from disk.

in the next section.

4.2.3 Local Prefetching

What is special about *WindowCache* is that it uses a local prefetching strategy that optimizes block searches and the triggering of prefetches. It is local in the sense that the blocks to be prefetched are always connected neighbours of the block currently in use. The prefetching strategy varies according to the access pattern of the IP algorithm and is described in the following.

Algorithm 1 Handling of a pixel request by the cache.

```

function cache::pixel(p)
  1:  $b_i \leftarrow f(p)$  {obtain block identifier from pixel coordinates}
  2: if  $b_i \neq \text{index}(\text{current\_block})$  then
  3:   if  $\neg \text{find}(\text{cache\_buffer}, b_i)$  then
  4:      $\text{load\_from\_disk}(b_i)$ 
  5:     if  $\text{full}(\text{cache\_buffer})$  then
  6:        $b_{lru} \leftarrow \text{lru}(\text{cache\_buffer})$ 
  7:       if  $\text{modified}(b_{lru})$  then
  8:          $\text{flush}(b_{lru})$  {write to disk}
  9:       end if
  10:       $b_{lru} \leftarrow b_i$  {replace LRU block}
  11:     else
  12:        $\text{insert}(\text{cache\_buffer}, b_i)$  {add block to cache}
  13:     end if
  14:   end if
  15:    $\text{prefetch}(p)$  {prefetch blocks}
  16:    $\text{current\_block} \leftarrow b_i$  { $b_i$  becomes the current block}
  17: end if
  18: return  $\text{pixel}(\text{current\_block}, p)$ 
end

```

Sliding Window

The concept of local prefetches behind *WindowCache* is that of a window that slides over the data. In other words, the cache only “sees” a small part of the image. This part consists of the image region in use at one moment and those that are candidates to be accessed in the subsequent steps of the IP algorithm. Figure 4.4 illustrates the concept with a 2-dimensional image and three different window configurations (the crisper areas of the images represent the part of the file that is already in primary memory). The central block of the window in each picture represents the image region in use. The neighbour regions are those already available for



Figure 4.4: Different sliding window configurations. The crisper areas of the pictures represent the image regions actually present in main memory.

the next steps. As the algorithm moves through the image, so does the sliding window. Regions of the image that are no longer covered by the window are discarded and flushed to disk if modified, and new areas of the window are made available to be filled.

The local prefetching uses in-depth knowledge about the access pattern of an algorithm in order to take advantage of the locality of the data access. As such, each type of algorithm has a different window structure and window sliding protocol. The local prefetching optimizes the search for blocks in the cache buffer whenever the requested pixel is not in the central block, since, using the locality principle, there is a high probability that the requested pixel is located in one of the neighbour blocks. In other words, the cache system needs only to search for the pixel within the other blocks of the cache window, whose number of blocks is considerably lower than the total number of blocks in the cache. In addition, the windows are implemented as simple references to the blocks of the cache buffer. As a result, their memory overhead is minimal.

Algorithm 2 reflects the changes in Algorithm 1 above when the local prefetching strategy of *WindowCache* is added to the process. The main difference is the replacement of a search for block b_i within the cache buffer by one within the cache window (step 3), which can be executed in $O(k)$ time ($k \ll N_B$). If b_i is not in the cache window, it means that the window does not cover the region containing the requested pixel and therefore needs to be rebuilt. The rebuild is executed synchronously, but this operation does not necessarily translates into I/O access, since the rebuild will

Algorithm 2 Handling of a pixel request by the cache, with local prefetching.

```

function cache::pixel(p)
  1:  $b_i \leftarrow f(p)$  {obtain block identifier from pixel coordinates}
  2: if  $b_i \neq \text{index}(\text{current\_block})$  then
  3:   if  $\neg \text{find}(\text{cache\_window}, b_i)$  then
  4:     rebuild_window(p) {requested pixel was not in the window}
  5:   else
  6:     local_prefetch(p) {prefetch blocks}
  7:   end if
  8:   current_block  $\leftarrow b_i$  { $b_i$  becomes the current block}
  9: end if
  10: return pixel(current_block, p)
end

```

actually execute steps 3 through 14 of Algorithm 1. The same holds for the prefetches of step 6 of Algorithm 2. The cache window slides and the references within the window are updated. The blocks to be prefetched only issue I/O operations (in this case, asynchronous) if they are not in the cache buffer. Regardless of whether a rebuild or a prefetch is executed, the requested block, b_i , is guaranteed to correspond to the central block of the cache window, and is made the current block before the pixel is returned to the Main Program.

Prefetching Strategies

Next, three types of IP algorithms are presented with their respective cache windows and sliding protocols. These algorithms cover a large proportion of IP tasks, which are frequently used as the basis for other more complex tasks.

Raster Scan Algorithms Various IP algorithms are implemented as a sequential traversal of the image, processing one single pixel at a time.

Examples are thresholding, brightness and contrast adjustment, histogram equalization, image arithmetic, etc (Gonzalez and Woods, 2001). This is the simplest and fastest traversal pattern, especially when it matches the layout of the input file. In this particular situation, native virtual memory and memory mapping schemes offered by operating systems are likely to efficiently traverse the file. However, other traversal patterns that do not match the file layout, but are still sequential (e.g., traverse the columns of a 2D image in row-column order, or the X plane of a 3D image in ZYX order), will not benefit from such schemes.

In this cache strategy, the idea is to prefetch blocks b_{i-1} and b_{i+1} of the image every time block b_i is accessed. The window structure of Figure 4.5, with 3 sequential blocks, is proposed to prefetch the neighbour blocks in an efficient manner by detecting in which the direction the traversal is moving. The central window block, C, is the window area being processed at a certain time, corresponding to block b_i . Neighbour blocks represent antecedent, b_{i-1} , and subsequent, b_{i+1} , image neighbours. When the cache detects that the traversal crossed the border of C towards the window block N, C receives the block of the cache buffer referenced by N; the first window block, P, receives the image block previously referenced by C, and the last window block, N, becomes available. At this moment, a prefetch to update N is triggered, but is satisfied in a separate thread, thus not causing the application to stall. If the algorithm moves backward, a similar arrangement of blocks is performed, and the window block P becomes available. This strategy can also be adapted for accesses in strides, in which the neighbour window blocks correspond to image blocks b_{i-N} and b_{i+N} , with N being the stride length. The proposed sequential window is an extension and a generalization of the *one-block look ahead* strategy proposed in Cucchiara et al. (2004), or the *iterator aware cache* proposed in Rhodes et al. (2005a,b).

Convolution Algorithms Classic examples of convolution algorithms are the ubiquitous digital filters used in IP for blurring, edge detection, sharpening, denoising, etc (Gonzalez and Woods, 2001). Another exam-




Figure 4.5: Prefetching window for sequential access patterns.

ple is the chamfer distance transform, which also takes into account the neighbourhood of a pixel in the calculation of distances Borgefors (1986); A. Butt and Maragos (1998). Although these filters are implemented as a sequential traversal through the data, as in a raster scan, the pixel and its neighbourhood within a mask of a certain radius are considered in each step.

Despite the sequential pattern, the cache structure of Figure 4.5 and its sliding protocol are not suitable for these neighbour-based algorithms. The problem is that it triggers unnecessary prefetches each time the algorithm's mask crosses the border between C and N or P, as depicted in Figure 4.6. If, for example, the mask invades the next block (Figure 4.6.b), block N is prefetched (Figure 4.6.c). The mask, however, remains at the border and the block preceding the updated previous block is prefetched when the mask moves backward (Figure 4.6.d). These two prefetches are issued several times until the mask completely crosses the border. Although it is very likely that the blocks prefetched will already be in the cache buffer (thus avoiding I/O access), this possibly large number of prefetch calls is very inefficient.

To overcome these problems, let us define the neighbourhood mask of the IP algorithm as having the same dimensions as the original image and size $(2r + 1)^d$ along each dimension, where r is referred to as the mask's radius, in number of pixels. Let \mathbf{S} be a virtual expansion of the central block, by r pixels, along its borders with the neighbour blocks. \mathbf{S} is illustrated in Figure 4.7 as the dark gray area at each side of the central block, in a 2D image. Whenever the pixel requested crosses the border of the central block, as in Figure 4.7.b, the window does not slide and no prefetch is triggered, provided that the pixel is inside \mathbf{S} . When the pixel requested

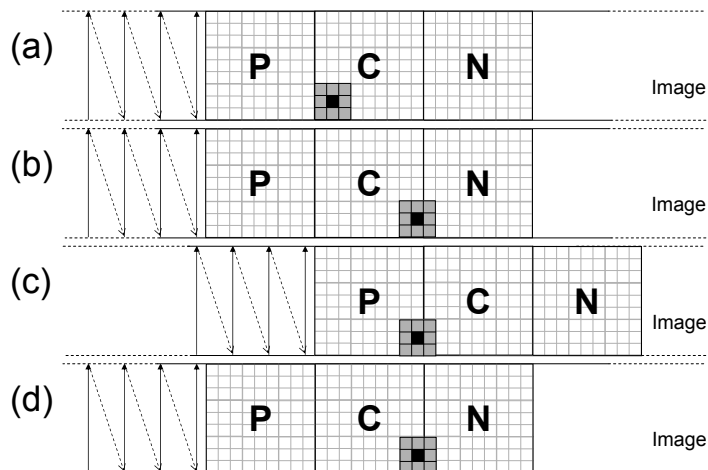


Figure 4.6: Situation when the mask reaches the border between two blocks: (a) the mask at the beginning of the central block; (b) the mask at the border; (c) the cache slides and the last block is prefetched; (d) the mask remains at the border. The arrows indicate the traversal pattern.

is finally beyond \mathbf{S} (Figure 4.7.c), the prefetch is triggered and the pointers to the cache blocks are updated as described in the previous section. However, the mask is now entirely located in the central block, and there is no risk that it will cross again the border left behind while the algorithm keeps moving forward (Figure 4.7.d).

Propagative Algorithms In general, the next pixel to be visited in propagative algorithms is decided upon the evaluation of a certain logical function. For this reason, such algorithms are much less predictable than those with sequential access patterns. Nevertheless, they still exhibit a good level of locality because the decision of which path to take is usually based on the connected neighbourhood of the current pixel (Cucchiara et al., 2004). Region growing and edge tracing algorithms fall into this category (Gonzalez and Woods, 2001).




Figure 4.7: *WindowCache* when the convolution mask reaches the border between two blocks: (a) the mask is in the central block; (b) the mask invades next block, but no prefetch is triggered; (c) the mask exceeds the extra space **S** and a prefetch is triggered; (d) the mask is completely inside the central block. The arrows indicate the traversal pattern.

Take as an example propagative access patterns following 6-connected neighbourhoods in a 3D image. Although such access pattern can produce paths with different connectivities on a global level, the cache configuration proposed in Figure 4.8 is an intuitive candidate to locally match the pattern. The blocks are organized in a 6-connected manner and are named after their cardinal direction in space - Central (C), Front (F), Back (B), North (N), South (S), East (E) and West (W). Again, the central block points to the data being processed. Whenever the window moves towards block N, block S references the block previously referenced by block C; block C references the block referenced by N, and blocks N, W, E, F, and B are prefetched. Similar reasoning is used to slide the window in all other directions.

This ends the description of *WindowCache* and how it can be used with different types of IP algorithms. The next section presents several experiments carried out on real data, using the algorithms described above,

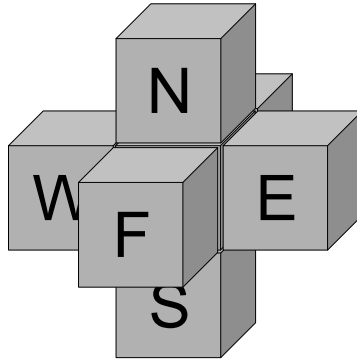


Figure 4.8: Prefetching window for 6-connected, propagative access patterns.

with their respective cache structures and window sliding protocols.

4.3 Experiments

The purpose of the experiments carried out was to compare the performance of the cached IP algorithms with their *in-core*¹ versions and with per-pixel I/O requests. The algorithms were also tested with a cache strategy that was not theoretically the most appropriate for its access pattern. In this way, this chapter aims to show that the best cache structures and prefetching strategies are those tailored for a certain type of algorithm and that the proposed method outperforms native operating system caches.

In order to evaluate the cache performance, the following were considered in each run of the IP algorithms:

- S_c : size of the cache;
- N_B : number of blocks in the cache;
- N_{ref} : number of pixel references;

¹When the entire image is loaded into main memory.

- N_{miss} : number of cache misses (requested pixel not in the cache);
- N_{rl} : number of synchronous block reloads (requested block not in the cache);
- N_{pf} : number of block prefetches;
- FPS : frame rate of the algorithm (where applicable);
- T : execution time of the IP algorithm (where applicable).

It is worth mentioning that read delays as illustrated in Figure 4.3 were not counted as misses, since the block containing the requested pixel is already scheduled for a prefetch.

For the in-core versions, only part of the images could be loaded to memory, and the results obtained were scaled to estimate those for the complete image. For the tests involving per-pixel I/O requests, a special cache abstraction was implemented, which issues one read (or write) operation for every pixel access. To reduce the effects of the operating system cache, the memory was cleaned with a large block of dummy data before each run of the IP algorithms.

The experiments were executed on a Intel[®] Core[™] 2 Quad CPU, at 2.4 GHz, with 2GB of addressable RAM, running under Windows Vista[™] Ultimate 64-bits, and a SATA hard disk drive of 750GB, 7200RPM, average seek time of 8.9ms, and average latency of 4.2ms. All image algorithms and the cache library were implemented in C++. The files used in our experiments corresponded to CT scans of the chest, organized in slice, row, column order. The file sizes varied from 1GB to 16GB.

4.4 Results

4.4.1 Raster Scan Algorithms

Let us begin with the results obtained with the sequential access pattern of Section 4.2.3. In these experiments, the IP algorithm implemented was a simple loop to traverse the entire image and perform a task with a pixel

Table 4.1: Cache sizes and number of blocks in the experiments with sequential traversal patterns.

slice traversal order												
S_c (slices)	20				40				100			
N_B	5	10	15	20	10	20	30	40	25	50	75	100
row and column traversal orders												
S_c (rows or columns)	10				20				50			
N_B	3	5	7	10	5	10	15	20	12	25	37	50

$p = (x, y, z)$, the result of which was stored in the output image. Two files were used in these experiments. The first contained 16GB (2048 columns \times 2048 rows \times 2048 slices, 2 bytes per pixel) and was traversed in slice order. The other contained approximately 1GB (768 \times 768 \times 997 pixels, 2 bytes per pixel) and was traversed in row and column orders. Although in the latter case the image fits in memory, input and output buffers do not fit at the same time. The cache size, S_c , was configured as approximately 1%, 2%, and 5% of the image size (given in number of slices, rows, or columns) and the number of blocks, N_B , corresponded to 25%, 50%, 75%, and 100% of S_c . This resulted in 12 runs of the same algorithm for each traversal direction. Table 4.1 shows all combinations.

The cache demonstrated very stable behaviour in terms of the computed statistics, irrespective of the traversal direction. For either $N_{ref} \approx 8.59 \times 10^9$ (16GB file) or $N_{ref} \approx 588 \times 10^6$ (1GB file), N_{miss} was equal to 1 in all runs, which is a remarkably low miss ratio. This was in fact expected, since, given that there are no jumps in the traversal, only the first pixel request requires a cache reload (compulsory miss Hennessy and Patterson (1990)). For the same reason, only the first two requested blocks needed to be synchronously loaded from disk, meaning that $N_{rl} = 2$. All subsequent pixel requests were satisfied by prefetches. In fact, due to the very predictable behaviour of the access pattern, even N_{pf} is known prior to the execution of the algorithms. It is simply a function of the number of blocks in which the image was subdivided along the direction of the traversal. This predictability does not mean, however, that the prefetches prevented all read delays. The

Table 4.2: Statistics computed for the sequential access pattern.

Traversal	S_c	N_B	N_{ref}	N_{miss}	N_{rl}	N_{pf}	FPS
slices	(in-core)	-	$\approx 8.59 \times 10^9$	-	-	-	6.833
	(per-pixel I/O)	-		-	-	-	0.027
	20	5		1	2	510	1.130
		15		1	2	2046	1.169
	100	50		1	2	1022	1.145
100		1	2	2046	1.1703		
rows	(in-core)	-	$\approx 588 \times 10^6$	-	-	-	36.079
	(per-pixel I/O)	-		-	-	-	0.182
	10	5		1	2	382	3.056
		10		1	2	766	1.439
	50	12		1	2	382	2.824
37		1	2	766	2.769		
columns	(in-core)	-	$\approx 588 \times 10^6$	-	-	-	7.189
	(per-pixel I/O)	-		-	-	-	0.005
	10	3		1	2	254	0.205
		7		1	2	766	0.079
	20	10		1	2	382	0.171
20		1	2	766	0.055		

FPS of the cached versions of the IP algorithms were considerably lower than in the respective in-core versions, which means that the higher per-pixel complexity of the cache and the lower speed of the I/O operations w.r.t main memory accesses took their toll. Table 4.2 shows some of the statistics captured for the sequential access patterns.

In terms of speed, special attention should be given to the in-core versions of the algorithms. In the slice traversal order, each slice had 2048×2048 ($\approx 4M$) pixels, and traversing one such slice takes long even in primary memory. For the row order, traversing the 768×997 ($\approx 766K$) pixels of each row was less computationally intensive, not to mention that the row organization in main memory is still cache friendly, since the memory organization of the image mimics the file layout. This explains the high FPS ratio. The column order, on the other hand, is not cache friendly because jumps in the image buffer are necessary to access each pixel. Even

though the number of pixels was the same as in the row order traversal, the *FPS* ratio was much lower (in this case, nearly 5 fold).

It should also be noted that for the row and column traversal orders the choice of N_B was an important factor. Since these access patterns do not match the file layout, block reloads and prefetches may take long to run, given the possible high number of I/O operations issued. The lower the number blocks in the cache, the faster the algorithms run, especially for small caches. This happens because the consequently larger blocks take longer to be consumed relative to the time the cache spends in I/O operations. Still, in favour of *WindowCache* is the fact that the versions of the algorithms that issued one I/O operation for each pixel request needed to be interrupted because they were taking too long to run (see the *FPS* entry of the “per-pixel I/O” lines of Table 4.2). This happened even in the slice traversal order, which in theory would benefit from the native operating system cache, given that the traversal matches the file layout.

Finally, the experiment of the slice traversal order was also executed with the 6-connected window of Section 4.2.3. The cache size was set to $S_c = 50$ slices and $N_B = 1000$ blocks. The algorithm also needed to be interrupted due to a very low frame rate of $FPS = 0.018$. The proposed sequential caches thus proved to be a valuable tool in this type of out-of-core IP operations.

4.4.2 Convolution Algorithms

The results obtained with the experiments with convolution patterns were analogous to the ones above. In this case, the algorithm implemented was a blur filter that averages the 26-connected neighbourhood of every pixel in the image and stores the result in the output image. The file used in these experiments was the same 1GB file used for the sequential patterns, and the algorithm traversed it in slice order.

The main difference to be noted w.r.t. the simple sequential patterns is that the frame rate was much lower with the blur filter. This was of course predictable due to the inherently higher complexity of the algorithm. Yet, for $N_{ref} \approx 15.8 \times 10^9$, N_{miss} was equal to 1. This miss ratio, although

Table 4.3: Statistics computed for the blur filter. FPS_s refers to the frame rate using the sequential caches.

S_c	N_B	N_{ref}	N_{miss}	N_{rl}	N_{pf}	FPS	FPS_s
(in-core)	-	$\approx 15.8 \times 10^9$	-	-	-	1.351	-
(per-pixel I/O)	-		-	-	-	0.011	-
10	5		1	2	497	0.529	0.301
20	5		1	2	248	0.540	0.322
	10		1	2	497	0.528	0.294
50	25		1	2	497	0.525	0.292

impressive, was once again expected, since the behaviour of the blur algorithm is identical to the sequential access pattern above. The sliding window protocol proposed in Section 4.2.3 makes a difference in the number of prefetches and processing speed of the algorithm, not in the number of misses. Indeed, the blur filter with the neighbour based cache was up to 80% faster than with the sequential cache of Section 4.2.3. Table 4.3 presents some statistics.

4.4.3 Propagative Algorithms

Finally, for the experiments with propagative access patterns, a region growing algorithm was used. This algorithm followed local 6-connected paths through the image, and pixels were processed on a LIFO basis. The aim of the application was to segment the airways and lungs from the chest CT scan whose file had 1GB in size (the same used for the blur filter above). S_c was equal to the previous experiments, while $N_B \in \{100, 200, 500, 1000, 2000, 4000, 8000\}$. In total, 21 runs of the algorithm were executed, and Table 4.4 shows the results.

Given the low predictability of the access pattern and the large N_{ref} , the performance of *WindowCache* was remarkably good, since N_{miss} varied from 1 to a maximum of 105. As one would expect, *WindowCache* was faster with larger cache sizes. What is worth observing as well is the fact that efficiency was affected by the choice of N_B . As N_B increases, for the

Table 4.4: Statistics computed for the region growing algorithm.

S_c	N_B	N_{ref}	N_{miss}	N_{rl}	N_{pf}	T (min)
10	100	$\approx 205.8 \times 10^6$	37	14082	399203	≈ 35
	200		55	19107	425826	≈ 24
	500		105	23969	448593	≈ 14
	1000		81	22096	371315	≈ 9.5
	2000		69	21697	307867	≈ 6.5
	4000		10	17196	230883	≈ 5
	8000		3	26569	344092	≈ 6.7
20	100	$\approx 205.8 \times 10^6$	21	6097	193014	≈ 29
	200		12	6542	184425	≈ 17
	500		23	6014	131884	≈ 8
	1000		3	3202	56779	≈ 4.5
	2000		2	3087	54839	≈ 4
	4000		1	6389	103471	≈ 4.4
	8000		1	13179	193739	≈ 5.5
50	100	$\approx 205.8 \times 10^6$	7	1561	62729	≈ 18
	200		1	944	30477	≈ 7
	500		1	449	8944	≈ 3.5
	1000		1	684	14913	≈ 3.6
	2000		1	1137	24735	≈ 3.8
	4000		1	2229	43464	≈ 4
	8000		1	4290	78956	≈ 4.5

same S_c , occasional block searches tend to take longer to run, which affects the execution speed. In addition, a larger N_B implies that the blocks of the cache encompass a smaller region of the image. On the one hand, these smaller blocks can be read from disk reasonably fast. On the other hand, if the traversal pattern revisits a certain region of the image, there is a higher probability that this region will not be covered by any block of the cache. Although a lower N_B results in bigger cache blocks, these blocks are updated more slowly. Since the access pattern is not globally predictable, only very few pixels of a big block may eventually be used before the block is discarded to give space to another one.

In another experiment, the same region growing algorithm was executed with the window structure and sliding protocol for a slice based sequential pattern (Section 4.2.3). Although the number of misses and reloads was still remarkably low ($N_{miss} = 1$ and $N_{rl} = 8$), for $S_c = 50$ slices and $N_B = 10$ blocks, the number of prefetches soared to $N_{pf} = 174614$. Since each prefetch loads 5 image slices at once (S_c/N_B) and the region growing algorithm may move very quickly from a pixel in one block to one in the neighbour block, read delays become very common. As a result, the execution time was $T \approx 63\text{min}$.

4.5 Discussion

Despite the very low miss ratios demonstrated with *WindowCache*, the negative impact on the efficiency of the algorithms relative to the in-core implementations should not be neglected. This impact, however, was predictable, since even CPU-level caches tend to slow down instruction processing (Hennessy and Patterson, 1990). In the case of *WindowCache*, the added complexity of the cache structure increased the number of instructions to be executed per pixel reference, which was responsible for the speed reduction. The proper comparison with respect to cache benefits, however, is always with a scenario in which the cache is not present and slow memory access are performed for every pixel requested. In out-of-core applications, disk access is the slow memory access in question, and it was also shown

that *WindowCache* outperforms approaches that issue one I/O operation for every pixel access.

The versions of the algorithms that issued one I/O operation for every pixel access also showed that cache and prefetching schemes offered by the operating system may not always help. This was especially true for access patterns that did not match the file layout, as the experiments demonstrated. The experiments further proved that the best cache and prefetching strategies are those tailored for a specific access pattern. The results obtained for the region growing application were the most prominent example, meaning that caches made for sequential access patterns (such as those in (Cucchiara et al., 2004; Rhodes et al., 2005a,b)) are not suitable.

4.6 Conclusion

This chapter presented a cache strategy to efficiently use primary and secondary memories in out-of-core processing of large image files. The strategy, named *WindowCache*, was developed upon the concept of a sliding window that represents the region of the image being processed in a certain moment. The window takes advantage of the temporal and spatial localities of the access pattern of the IP algorithm and issues block prefetches before they are requested by the application. These prefetches are executed asynchronously, preventing the main application from stalling while data is read (written) from (to) disk. Several experiments showed the benefits of using *WindowCache* when compared to strategies relying only upon operating system caches. The chapter also showed the importance of using a cache strategy that matches the access pattern of the IP algorithm by comparing executions of the same algorithm with different versions of *WindowCache*.

Chapter 5

Airway Tree Segmentation

“You are stupid! You are stupid!
Oh, and don’t forget, you are STUPID!”
Dexter, from Dexter’s Laboratory

“If your experiment needs statistics,
then you ought to have done a better experiment.”
Ernest Rutherford (1871 – 1937)

Parts of this chapter appear in:

- R. Pinho, J. Sijbers, and W. Vos. Efficient approaches to intrathoracic airway tree segmentations. In *Proc. of the Biomedical Engineering IEEE/EMBS Benelux Symposium*, volume 2, pages 151–154, Brussels, Belgium, 2006.
- R. Pinho, S. Luyckx, and J. Sijbers. Robust region growing based intrathoracic airway tree segmentation. In *2nd International Workshop on Pulmonary Image Analysis*, pages 261–271, London, England, 2009.

5.1 Introduction

As will be clear in the next chapter, the segmentation of the trachea from CT images is an integral step of the proposed system for assessment of stenosis. Although this segmentation can be achieved in different ways, as discussed in Chapter 3, a problem that needed to be solved in this thesis was the automatic segmentation of the trachea from the very beginning (below the cricoid cartilage) to the very end (at the level of the carina). This was in fact integrated with the segmentation of the whole airway tree, which is the subject covered in this chapter.

Airway tree segmentation is the process of identifying and extracting from volumetric medical images the structures of the respiratory system that lead the air into the lungs. With the result of the segmentation, physicians and researchers can make measurements, check for abnormalities and be generally assisted in diagnosing diseases in the respiratory system.

Due to the natural complexity of the airways, with several branching levels, and noise or other artefacts present in the image, the segmentation is a task far from trivial. A common method to solve the problem is region growing (Gonzalez and Woods, 2001), and semi- and fully-automated region growing algorithms have been used to segment the airways (Chiplunkar et al., 1997; Kiraly et al., 2002; Law and Heng, 2000; van Ginneken et al., 2008). In this process, the user provides one or more seed points inside

the airway structure. From these points, a region is grown by recursively aggregating voxels that pass a certain test of similarity. Common similarity tests check differences in intensity between neighbouring voxels.

One common problem of region growing algorithms is leakage. In the case of the airway tree segmentation, a thin wall separates the structure from neighbouring organs and air inside the lungs. Noise, partial volume effect, or other artefacts can create holes in this wall and, since the airway lumen and the lung interior have similar voxel intensities, the entire lung can be aggregated to the region. Another problem specific to airway tree segmentation is the early collapse of branches. In this case, the growing process stops too early, resulting in only partially segmented branches.

Tschirren et al. (2005) proposed an algorithm that takes advantage of the fact that the airway tree is a hierarchical combination of cylindrically shaped objects. In their algorithm, cylinders of adaptive radius and orientation bound the segmentation, facilitating the process of leak detection. The method proposed in this chapter extends the mentioned approach, adding the following contributions:

- propose a heuristic algorithm to automatically select a seed point inside the trachea, since chest CT scans often include the upper airways and other regions, complicating this task;
- use cylinders of adaptive height as well as adaptive radius in order to bound the segmentation;
- use anatomical, instead of pure image information to detect leaks;
- propose a more efficient way of computing an approximation of the centre line of the segmented airways;
- propose a new strategy to avoid leaks, by taking into account the fact that they grow through small holes on the edges of the structure being segmented.

The proposed method was evaluated with a dataset of 40 patients, subdivided into training and testing groups. Measures of number of branches,

airway tree length, and leakage were taken in order to evaluate the method, by comparing it to ground truth data. The next section describes the proposed methods for the segmentation of the complete trachea and of the airways. Section 5.2 presents the experiments with the referred dataset and the obtained results. Section 5.5 discusses the results obtained and some particularities of the proposed algorithm. The chapter is then concluded in Section 5.6.

5.2 Method

The proposed algorithm uses ideas first proposed by (Tschirren et al., 2005), in which the region growing segmentation is bound by cylinders (ROIs) of adaptive dimensions. These ROIs limit the expansion of leaks, making them more easily detectable. The basic algorithm starts from a seed point somewhere in the trachea, where the origin of the first ROI is set. The following ROIs are computed from the intersections between the region grown inside the current ROI and its borders. The skeleton of this region and the diameter of the intersection determine the location, orientation, and diameter of the next ROI. Leaks in this approach are detected with a morphological operator, based on the assumption that they have a "spongy" structure in the image. They are potentially eliminated by repeating the segmentation within the ROI, but following only paths along its direction, called "directional affinity". Figure 5.1 illustrates how the algorithm evolves.

The algorithm above is extended with following new ideas:

ROI dimensions In the previous work, the height of the ROIs only change when the region grown stops exactly at a bifurcation point of the airway. Here, the height is also a function of the branching level, given by

$$h_l = H_0 \rho^{(l-1)}, h_l \geq H_{min} \text{ ,} \quad (5.1)$$

where $l \geq 1$ is the current branching level ($l = 1$ is the trachea), h_l is the height of the ROI, in millimetres, H_0 is the default initial height, $\rho \in [0, 1]$ is the height change ratio, and H_{min} is the minimum height limit. By doing




Figure 5.1: Steps of the basic algorithm.

this, short branches at higher branching levels, which otherwise would not intersect an ROI, can be detected. The diameter of each ROI, in turn, is set to 50% longer than the diameter of the intersection between the region grown and the ROI of the previous step.

Leak detection The proposed detection of leaks is based on prior anatomical knowledge about the airway structure. Figure 5.2 shows that the diameter of the airways from one branching level to the next hardly ever increases. Moreover, there are usually no more than 3 or 4 branches per level, as illustrated in Figure 2.1. The proposed method therefore detects leaks by checking the variation in diameter from one branch to the next and the number of branches within each ROI. Whenever a branch intersects the current ROI, the area of the region of intersection, computed as the number of pixels, is taken as a reference for the diameter of the branch. A simple test determines if the increase in area from the beginning of the ROI to the point of intersection is above a factor of 2, in which case a leak is detected. The test for the number of branches, or intersections, is made with two thresholds. If $l < 4$, a leak is detected if the number of branches within the ROI is greater than 5. For lower levels, the threshold is reduced to 3, since the ROIs' radii and sizes will also be smaller.





Figure 5.2: Diameter distributions of the airway branches (Tschirren et al., 2005)

Skeleton efficiency In the previous work, the skeleton of the segmented surface is computed with distance transforms and path cost minimization. This can be time consuming, especially being computed for every ROI. Since the barycentres of the intersections between the branches and the ROI are originally used as anchor points of the skeleton, the proposed idea is to directly take them as an estimator of the spatial orientation of the ROI of the next step. To this end, the walls of the ROIs are classified as top, bottom and side borders. The top border is where the segmentation always starts (where the seed pixel for the ROI is placed). The segmentation process can naturally stop somewhere inside the current ROI or by intersecting any of its borders. The regions obtained from these intersections have their


barycentres calculated. When the intersection occurs at the bottom border, the barycentres of the start and end regions are directly connected. For intersections with side and top borders, a "global" barycentre (barycentre of the barycentres) is computed and connected to the respective side and top barycentres. This is a very straightforward approach which is also very fast. The accuracy of the estimated skeleton depends on the height of the cylinders, but tends to be good enough to compute the location and orientation of the ROIs. Figure 5.3 illustrates the procedure.

Avoiding leaks Finally, whenever the segmentation within an ROI is repeated due to a leak, each candidate voxel and its neighbourhood within a mask are analysed. Only if the voxel and all its non-visited neighbours pass the similarity test (voxel intensity lower than a threshold T in our case) is the voxel aggregated to the region. Each time the segmentation is repeated, a mask of higher radius is used, until no leaks are detected or a maximum number of attempts is reached. The reasoning behind this approach is that a leak always occurs due to the presence of holes on the boundaries of the region being segmented, so the algorithm basically tries to discover their sizes. In contrast to the direction affinity adopted in (Tschirren et al., 2005), the neighbour affinity technique allows more possibilities of continuing with the segmentation, while trying to avoid leaks. The 2D scheme of Figure 5.4 illustrates this idea.

Algorithm 3 details the segmentation loop. Step 1 comprises a sequence of steps to compute the first ROI, using the given image and seed point, and push it onto an ROI-queue. After growing a region within the ROI at the front of the queue, Step 5 detects leaks using anatomical information as described above. The regions of intersection between the ROI and the region grown that correspond to leaks are put in a list. If a leak is detected, Step 7 switches neighbour affinity on by providing the next neighbourhood mask to be used in the similarity tests. The ROI is reset and pushed back onto the queue in Steps 9 and 10, respectively. If leaks are still present after trying all masks, Step 13 removes from the list obtained in Step 5 the corresponding branches. Step 15 processes the remaining branches of the



(a) Segmented region, ROI, and intersections (centre). True skeleton (left). Proposed improvement using only the barycentres (the point in the middle is the “global” barycentre) (right).



(b) Regions and their underlying directions (left). The corresponding ROIs (right).

Figure 5.3: (a) Schematic view of the proposed skeleton estimation. (b) An example with real data.




Figure 5.4: Avoiding leaks. The segmentation is repeated with an increasing neighbourhood mask until no leaks are detected.

current ROI and returns a list of ROIs for the next iteration, which are pushed onto the queue. Finally, Step 20 switches neighbour affinity off and the process restarts.

5.2.1 Seed Point Selection

As mentioned previously, region growing algorithms need one or more seed points to mark the start of the segmentation. Algorithms that automate the seed point selection for the segmentation of the airways typically detect a circular region near the centre of a slice of the image volume. The region is supposed to correspond to the trachea, and the seed point is taken as its centre of gravity. This process may fail if the chosen slice contains misleading regions (e.g., if the CT scan contains parts of the upper airways) or does not contain the trachea at all. Figure 5.5 illustrates the former case with a slice containing the trachea, the oesophagus, and a possible tumour. This chapter presents a more robust method to automatically select a seed point inside the trachea.

For one axial slice $i = 1 \dots N$ of the image volume, the method works as in Algorithm 4. Let us use the *threshold below* operation to turn all voxels with intensity below a certain threshold to white and the rest to black (Shapiro and Stockman, 2002). Step 1 thus finds the best threshold to segment the air in the image, which includes the areas inside the lungs and airway lumen, using Otsu’s method (Otsu, 1979). Step 2 applies a masked,

Algorithm 3 region_grow(image, seed)

```

1: /* initialization and computation of 1st ROI */
2: while  $\neg$  empty(roi_queue) do
3:   r  $\leftarrow$  pop(roi_queue)
4:   roi_region_grow(image, r)
5:   has_leak  $\leftarrow$  detect_leak(r) {using anatomical information}
6:   if has_leak  $\wedge$  (count_leak < max_count_leak) then
7:     set_neighbour_affinity(mask[count_leak]) {set mask to avoid leak}
8:     count_leak  $\leftarrow$  count_leak + 1
9:     reset(image, r)
10:    push(roi_queue, r)
11:  else
12:    if has_leak then
13:      remove_leaking_branches(r) {may remove all branches}
14:    end if
15:    roi_list  $\leftarrow$  process_roi(r)
16:    for all  $r_i \in$  roi_list do
17:      push(roi_queue,  $r_i$ )
18:    end for
19:    count_leak  $\leftarrow$  0
20:    set_neighbour_affinity(null) {switch neighbour affinity off}
21:  end if
22: end while

```

morphological closing operation to the slice in order to fill all holes. Step 3 identifies 8-connected regions in the resulting image and labels them. Step 6 removes noise, i.e., all regions with size $s \leq S_{min}$ pixels. Step 10 takes care of eliminating narrowed regions, i.e., with eccentricity $e > e_{max}$, and steps 12 through 15 identify the region of the slice with highest eccentricity, R_{e_i} .

After these steps, a number of regions may be left in each slice. These regions comprise the trachea and areas corresponding to air outside the lungs, the upper airways, the lungs, etc. The challenge is then to choose




Figure 5.5: A situation where other structures may mislead the detection of the trachea in an axial slice of the CT scan.

the slice containing only the trachea or at least to correctly identify it when other structures are present. For this, a function of several parameters is minimized so as to favour:

- slices with fewer regions, since, in general, the upper part of the trachea tends to appear alone in the image;
- slices that maximize e , of R_{e_i} , given that the upper part of the trachea, just below the cricoid cartilage, tends to be elliptical;
- slices in which the major axis of R_{e_i} is aligned with the sagittal plane;
- slices with lower indexes, since the search is for the top of the trachea (assuming slice 0 coincides with the top position of the CT scan);
- slices in which R_{e_i} is small, which avoids confusion with the lungs;
- slices in which R_{e_i} maximizes the area of the ellipse, so that only “regularly” shaped ellipses are chosen.

Algorithm 4 find_trachea(slice_{*i*})

```

1: threshold(slicei) {segment air}
2: close(slicei)
3: regionsi ← label(slicei)
4: for all Rj ∈ regionsi do
5:   if size(Rj) < Smin then
6:     remove(regionsi, Rj) {remove noise}
7:   else
8:     e ← eccentricity(Rj)
9:     if e > emax then
10:      remove(regionsi, Rj) {remove narrow regions}
11:    else
12:      if e > maxe then
13:        maxe ← e
14:        Rei ← Rj
15:      end if
16:    end if
17:  end if
18: end for

```

The minimization is then defined as

$$\arg \min_{i=1..N} f(i, e_i, a_i, s_i, r_i) = n_{r_i} \left[\frac{i}{N} + \frac{s_i}{S} + (1 - e_i) + a_i + (1 - r_i) \right], \quad (5.2)$$

where $\{e_i, a_i, r_i\} \in [0, 1]$. In this equation, i is the slice index, $n_{r_i} \geq 1$ is the number of regions of the slice, s_i is the size of R_{e_i} in pixels, with S being a maximum size threshold, e_i is the eccentricity of R_{e_i} , a_i is the angle between R_{e_i} 's major axis and the sagittal direction, and r_i is a measure of area maximization. The latter is computed by taking the ratio between the number of pixels of R_{e_i} and the area of its corresponding ellipse. Lastly, the selected seed point is the centre of gravity of the R_{e_i} that minimizes $f(\cdot)$.

5.3 Experiments

As stated in Section 5.1, the proposed method was tested with a dataset of 40 patients. This dataset was provided as part of the workshop and airway segmentation challenge *EXACT09: Extraction of Airways from CT 2009* (Lo et al., 2009). The data was subdivided into one training and one testing group, each with 20 patients, numbered CASE01...CASE20 and CASE21...CASE40, respectively. The segmentation was evaluated by a team of trained observers. The aim of the workshop was to compare the performance of different algorithms. For this purpose, a ground truth was constructed from all submitted segmentations and all submissions were evaluated with respect to this ground truth.

The objective of the experiments was to check, for the testing group, how many branches were detected, the segmented tree length and the amount of leakage. The following measures were used to compare the submitted results:

- Branch count: number of branches detected.
- Branch detected: the fraction of branches that were detected with respect to the branches present in the ground truth.
- Tree length: the sum of the length of the centre lines of all correctly detected branches.
- Tree length detected: the fraction of tree length that was detected correctly, relative to the tree length of the ground truth.
- Leakage count: the number of unconnected groups of “correct” regions that are neighbours of a “wrong” region.
- Leakage volume: the volume of regions that are wrongly detected.
- False positive rate: the fraction of the volume of regions that are detected wrongly relative to the volume of all detected regions.

The trachea was excluded from the branch length and branch count related measurements. For the voxel based measures of leakage, both trachea and main bronchi were excluded. Furthermore, the exact airway shape and dimensions were not taken into account.

Algorithms 1 and 2 were implemented in C++, and the programs were executed on an Intel[®] Core[™] 2 Quad CPU, at 2.4 GHz, with 8GB of RAM, running under Windows Vista[™] Ultimate 64-bits. The region growing algorithm used a single threshold value $T = -800HU$ for the whole airway tree and did not employ multiseeded connectivity, as opposed to Tschirren et al. (2005), since only the airway lumen is segmented, not the walls. In addition, intensities of candidate voxels were averaged within a 6-connected neighbourhood to reduce noise artefacts. The parameters ρ and H_{min} were primarily chosen empirically for the training group, but adjustments were necessary during the experiments with the test set. Eventually, $\rho = 0.85$ provided the best results except for CASE32, for which it was set to $\rho = 0.75$, with $H_{min} = 2\text{mm}$ in all cases. For the neighbour affinity, spherical and cubic masks with radii from 1 to 7 voxels were used, defining, in this order, 6, 18, 26, 92, 124, 342, 728, 1330, and 2197-neighbourhoods. These masks remained unchanged during the experiments with the test set, but new masks were added until the results for the training group were, at least visually, acceptable. With respect to the seed point selection, in Algorithm 4, $S_{min} = 250$, $e_{max} = 0.75$, S was equal to the number of pixels of the slice, and $N = \min(300, N_a)$ slices, where N_a is the number of axial slices of the image volume. Again, these values were empirically chosen for the training group, but remained unchanged with the test set.

5.4 Results

The results obtained with the region growing algorithm applied to the testing group¹ can be seen in Table 5.1. The main difficulty in the segmentation of the airways is to find the balance between the number of segments detected and leakage. In general, it is very difficult to increase the former

¹Provided by the organizers of the workshop.

without allowing the latter to increase as well. The proposed approach thus remained on the conservative side in terms of branch count and reach, but mostly with low leakage count.

With respect to the automatic seed point selection, the proposed algorithm performed very well in all cases. The selected point was always located inside the trachea, at the top. In very few situations, however, the point was set at a lower location. This happened when the trachea had an almost circular shape along all or nearly all of its length. Since the algorithm favours elliptical regions, such a shape happened to appear at slices with higher indexes. Given that the trachea was not considered in these experiments, it was not a problem, but adjustments to the function of Eq. (5.2) may still be necessary. Figure 5.6 shows an example of the trachea detection algorithm for one patient of the image database.

Table 5.2 presents the execution times of the proposed algorithms applied to the testing group. For the region growing, execution time is naturally an increasing function of the number of detected branches, but all executions ran in less than 1 minute, with half of them below 3 seconds and 75% below 8 seconds. The seed point selection, in turn, showed less variation, mostly because $N = 300$ slices for all cases. The differences between cases lay mainly in the complexity of each slice processed by the algorithm. Finally, Figure 5.7 presents the segmentation results for 2 patients.

5.5 Discussion

In the proposed algorithm, the choice of the radii of the ROIs has a clear impact on the behaviour of the segmentation. If the ROIs have small radii, they will be placed too close to the airways, and more intersections will occur. Not only will the complexity increase, but invalid intersections might appear, possibly leading to more leaks. On the other hand, cylinders far from the structure will not restrain leaks to small sizes. This makes the leak detection process less efficient and also prone to error. Finding a good balance can be difficult, but the chosen 50% increase in diameter relative to the diameter of the intersection between region grown and the previous

Table 5.1: Evaluation measures for the twenty cases in the test set.

	Branch count	Branch detected (%)	Tree length (cm)	Tree length detected (%)	Leakage count	Leakage volume (mm ³)	False positive rate (%)
CASE21	69	34.7	39.4	35.7	0	0.0	0.00
CASE22	132	34.1	86.4	26.1	7	160.0	1.14
CASE23	89	31.3	56.6	21.7	6	56.1	0.52
CASE24	69	37.1	56.0	34.4	14	277.2	1.66
CASE25	76	32.5	58.5	23.2	8	557.8	3.27
CASE26	35	43.8	24.3	37.0	0	0.0	0.00
CASE27	36	35.6	25.9	31.9	0	0.0	0.00
CASE28	53	43.1	35.2	32.1	1	473.7	7.60
CASE29	73	39.7	46.9	34.0	4	27.6	0.40
CASE30	47	24.1	33.2	21.7	0	0.0	0.00
CASE31	61	28.5	39.0	22.2	7	578.0	6.98
CASE32	64	27.5	46.6	21.4	2	1740.7	14.34
CASE33	70	41.7	50.2	34.2	5	670.3	11.25
CASE34	140	30.6	85.4	23.9	10	2407.9	12.70
CASE35	95	27.6	61.1	19.8	3	39.7	0.32
CASE36	83	22.8	69.7	16.9	0	0.0	0.00
CASE37	67	36.2	52.3	29.4	2	105.9	1.11
CASE38	28	28.6	23.5	35.3	0	0.0	0.00
CASE39	109	21.0	84.8	20.7	2	93.5	1.04
CASE40	88	22.6	63.6	16.4	13	1420.6	10.24
Mean	74.2	32.1	51.9	26.9	4.2	430.4	3.63
Std. dev.	29.5	6.9	19.6	6.9	4.4	672.3	4.92
Min	28	21.0	23.5	16.4	0	0.0	0.00
1st quartile	53	27.5	35.2	21.4	0	0.0	0.00
Median	70	31.9	51.3	25.0	3	99.7	1.07
3rd quartile	95	39.7	69.7	34.4	8	670.3	10.24
Max	140	43.8	86.4	37.0	14	2407.9	14.34




Figure 5.6: Example of the trachea detection algorithm for CASE24 in the image database. (a) The slice of the image chosen as the beginning of the trachea (top) and a schematic view of the parameters of the tracheal contour used by the detection algorithm (bottom). (b) 3D views of the selected tracheal contour (in red).

Table 5.2: Execution times of the algorithms applied to the test set.

	Region growing (secs)	Seed selection (secs)	Total time (secs)
CASE21	1.98	82.9	84.88
CASE22	6.85	79.6	86.45
CASE23	15.88	64.2	80.08
CASE24	15.02	63.2	78.22
CASE25	31.79	60	91.79
CASE26	1.59	68.8	70.39
CASE27	1.22	57.2	58.42
CASE28	1.83	61	62.83
CASE29	1.67	61.4	63.07
CASE30	0.97	61.2	62.17
CASE31	3.51	64.2	67.71
CASE32	25.74	66.8	92.54
CASE33	5.46	65.6	71.06
CASE34	44.62	57.1	101.72
CASE35	3.6	59.4	63
CASE36	1.62	69	70.62
CASE37	5.3	59.2	64.5
CASE38	2.79	58	60.79
CASE39	1.31	76.5	77.81
CASE40	2.22	70.2	72.42
Mean	8.75	65.27	74.02
Std. dev.	12.08	7.38	12.27
Min	0.97	57.1	58.42
1st quartile	1.66	59.85	63.05
Median	3.15	63.7	70.84
3rd quartile	8.89	68.85	81.28
Max	44.62	82.9	101.72




Figure 5.7: Segmentations with ROIs coloured per level (top) and respective 3D reconstructions (bottom) for CASE34 (left) and CASE39 (right), respectively.

ROI proved very good for the image database.

An issue when positioning ROIs is the fact that its orientation may differ considerably from the previous ROI. Simply placing the ROI on the barycentre of the intersection between the region and the previous ROI may not entirely cover the region to be segmented. This situation is shown in Figure 5.8. Unnecessary intersections with the top border of the cylinder may occur in this case, increasing the complexity of the segmentation. To solve the problem, the longest perpendicular distance from the region of intersection to the top border of the ROI, initially placed at the barycentre of the region, is calculated. The ROI is then shifted this distance in the opposite direction of its orientation and is guaranteed to cover the entire region to be segmented.

One characteristic of the proposed neighbour affinity algorithm is that the resulting segmentation becomes thinner as the neighbourhood mask increases. This, however, can easily be corrected with local dilation operations. Another aspect to be taken into account is that, at very high levels, the diameter of the airway segment may be very narrow, maybe only 2 or 3 pixels wide. At this point, the proposed neighbour affinity has no mechanism to differentiate between a very narrow segment and a small hole in the airway wall. Therefore, it is possible that these segments appear collapsed in the resulting segmentation.

5.6 Conclusions

This chapter presented a semi-automatic region growing method for the segmentation of the intrathoracic airways from tomographic scans. The method uses cylinders (or ROIs) of adaptive orientation and dimensions to bound the segmentation. The role of these ROIs is to set a limit to leaks, a common problem with region growing algorithms, and to allow them to be more easily detected. The proposed approach uses anatomical information about the airways in order to detect the leaks and includes a novel algorithm to avoid new leaks once they are detected. This chapter also presented a heuristic algorithm to automatically select a seed point at the




Figure 5.8: Positioning ROIs from level L to $L+1$. (a) Uncovered regions (in red) if the direction of ROI_{L+1} differs from that of ROI_L . (b) Zoomed area showing the proposed solution, i.e., move ROI_{L+1} "backwards" along the corresponding centre line of ROI_L so as to cover the longest perpendicular distance from the top of ROI_{L+1} to the indicated intersection.

top of the trachea, which is later provided to the region growing algorithm. The method was tested on a dataset of 40 patients, and remained on the conservative side in terms of branch detection, but with a low number of leaks in most cases.

Chapter 6

Estimation of Healthy Tracheas

“I’ve said it before, and I’ll say it again...
aye carumba!”
Bart Simpson

“The most exciting phrase to hear in science,
the one that heralds new discoveries,
is not ‘Eureka!’ (I found it!)
but ‘That’s funny ...’ ”
Isaac Asimov (1920 – 1992)

Parts of this chapter appear in:

- R. Pinho, T. Huysmans, W. Vos, and J. Sijbers. Tracheal stent prediction using statistical deformable models of tubular shapes. In *Proc. of SPIE Medical Imaging*, page 69144O, San Diego, CA, USA, 2008.
- R. Pinho, K. G. Tournoy, R. Gosselin, and J. Sijbers. Assessment of tracheal stenosis using active shape models of healthy tracheas: A surface registration study. In *2nd International Workshop on Pulmonary Image Analysis*, pages 125–136, London, England, 2009.
- R. Pinho, K. G. Tournoy, R. Gosselin, and J. Sijbers. Assessment and stenting of tracheal stenosis using deformable shape models. *Medical Image Analysis*, (Under Review), 2010a.
- R. Pinho, K. G. Tournoy, and J. Sijbers. A decision support system for the treatment of tracheal stenosis. In *Proc. of Workshop on Discrete Geometry and Mathematical Morphology (WADGMM)*, pages 72–76, Istanbul, Turkey, 2010b.
- R. Pinho, K. G. Tournoy, and J. Sijbers. Computer-aided assessment and stenting of tracheal stenosis. In Ayman El-Baz, editor, *Computer Aided Diagnosis of Lung Imaging*, volume (in press). American Scientific Press, 2011.

6.1 Introduction

This chapter presents the most important concepts behind the proposed system for the assessment and stenting of tracheal stenosis. When physicians assess the stenosis, it is intuitive to imagine that they try to guess from the available data how the patient’s trachea would appear if the narrowing were not present. This image of the healthy trachea is the basis to determine the location, length, and degree of the stricture. The image of

the calibre and curvature of the estimated healthy shape also aid in choosing the type and dimensions of the stent to be used, or even the shape the trachea should have after a reconstruction surgery.

Inspired by this procedure, the proposed method aims at providing to the physician the desired healthy shape. This shape not only represents the necessary visual aid in the diagnostic and pre-operative processes, but also enables the automatic assessment of the stenosis and suggestion of patient-specific stent dimensions. The estimation of the healthy shape of the trachea is obtained thanks to the use of an Active Shape Model of healthy tracheas and to a new algorithm to register such model to CT data.

An Active Shape Model (ASM), as defined by Cootes et al. (1995), is a statistical model of shapes whose objective is to capture the geometric variation present in a set of training shapes. This model is then capable of generating new shape instances which are constrained by the statistical variation of the training set, referred to as the allowable shape domain. As a result, a defining characteristic of ASMs, which is an advantage with respect to other deformable models, is that they can only generate shapes that resemble those in the training set.

This characteristic is the motivation to use an ASM of healthy tracheas to estimate the healthy tracheal shape of a patient with stenosis. Since the model contains only instances of healthy tracheas, local geometric variations typical of stenosis cannot be generated. The details of the construction and use of the proposed ASM are presented in the following.

6.2 Training Set

Before building a statistical shape model, a set of N shapes whose statistical variation will be captured must be selected. Here, a set of surfaces of healthy tracheas segmented from chest CT scans is used. The region growing method presented in Chapter 5 segments the trachea from the very beginning, below the cricoid, down to the region of the carina, where the trachea splits into the two main bronchi. It also generates a good approx-

imation of the centre line of the tracheas, which are used as the basis to achieve correspondence between shapes of the training set.

Note that when selecting images for the training set, the method relies on the assumption that the CT scans include the complete trachea. If the scans start above the trachea, the region growing method mentioned above takes care of identifying its starting point. CT scans starting slightly below the point where the trachea starts, on the other hand, should not be used, since they will introduce errors in the construction of the statistical model.

6.3 Establishment of Correspondences

An ASM needs a one-to-one correspondence between the points (or landmarks) describing the shapes in the training set in order to capture correct shape variations. A direct result from the segmentations above is that the region below the cricoid and the carina already correspond within the training set. For each trachea, a set $\mathcal{C}_i, i = 1 \dots N$, defines the sequence of points along its centre line, and, possibly, $|\mathcal{C}_i| \neq |\mathcal{C}_j| : j \neq i$, meaning that the correspondence along the centre line has not been established for all points in the set. This correspondence is achieved with arc-length parametrization, by which a new set $\mathcal{C}'_i = \{\mathbf{c}'_{i1}, \dots, \mathbf{c}'_{in'_c}\}$ is computed for each trachea, in such a way that $|\mathcal{C}'_i| = n'_c, \forall i$. Once the centre lines of all tracheas correspond, another set

$$\mathcal{C}''_i = \{\mathbf{c}''_{i1}, \dots, \mathbf{c}''_{in''_c}\} = \beta(\mathcal{C}'_i, n''_c), |\mathcal{C}''_i| = n''_c, \forall i,$$

is obtained by smoothing \mathcal{C}'_i with the cubic spline interpolation function $\beta(\cdot)$, using the elements of set \mathcal{C}'_i as control points. The cubic spline is the preferred interpolation rather than a more refined arc-length parametrization so as to achieve C^2 continuity along the curve.

The shapes of the training set, however, are described by points on their boundaries. A set \mathcal{B}'_{ik} , with n'_b points, is thus first computed for each centre line point $\mathbf{c}''_{ik}, k = 1 \dots n''_c$, by sampling the segmented trachea around the point at constant angular intervals. The set \mathcal{B}'_{ik} lies on a plane Π''_{ik} centred at point \mathbf{c}''_{ik} and perpendicular to the curve defining the centre line. The

normal vector of Π''_{ik} , \mathbf{n}''_{ik} , is obtained with an approximation of the first derivative of the curve at point \mathbf{c}''_{ik} .

Finally, a cubic spline interpolation of \mathcal{B}'_{ik} generates a set $\mathcal{B}''_{ik} = \beta(\mathcal{B}'_{ik}, n''_b)$, approximating the boundary of the surface intersecting Π''_{ik} with a smooth contour containing n''_b points.

The correspondence attained with the described procedure may still be incorrect due to one remaining degree of freedom, which is the rotation of the tracheas around their centre lines. Supposing that all patients are scanned in the same position (prone or supine)¹, this degree can be removed by computing the orientation of the contour defining the region of the carina, which is assumed to determine the orientation of the trachea with respect to the axial direction. This contour can be approximated by an ellipse, whose major axis forms an angle α_c with the coronal direction once the contour is aligned with the axial plane. The boundary sampling procedure is then modified in such way that the first point of each set \mathcal{B}'_{ik} is computed at an angle α_c with respect to the coronal direction.

Figure 6.1 illustrates the whole concept. Each trachea is finally described by a smooth, triangular surface $\mathcal{X}_i = (\mathcal{L}_i, \mathcal{T}_i)$, where \mathcal{L}_i is the set of points, or landmarks, with $n = |\mathcal{L}_i| = n''_c \times n''_b$, and \mathcal{T}_i is the set of triangles connecting the points of \mathcal{L}_i .

6.4 Active Shape Model

After establishing the correspondences, each shape is represented by n 3D landmarks concatenated into a $3n$ -dimensional vector $\mathbf{x}_i = (x_{i11}, x_{i12}, x_{i13}, \dots, x_{in1}, x_{in2}, x_{in3})^T$. The shapes are then aligned with respect to location and pose and normalized with respect to scale such that $|\mathbf{x}_i| = 1$. Principal Component Analysis (PCA) extracts the eigenvectors, \mathbf{p}_i , and non-negative eigenvalues, λ_i , of the set's covariance matrix. The eigenvectors, also referred to as the main modes of variation of the training set, are grouped in an orthonormal matrix $\mathbf{P}_{3n \times N}$. New shape instances can be obtained with this model, called a point distribution model (PDM),

¹The image can always be flipped if this supposition fails.

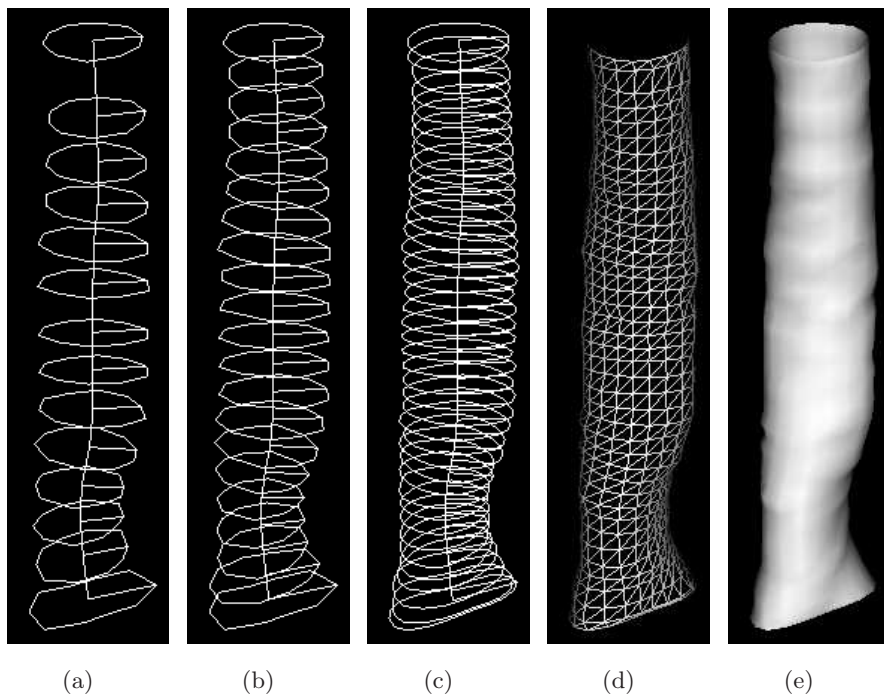


Figure 6.1: Example of the process to transform the shapes of the training set from a centre line to a boundary representation and to establish correspondences between them. (a) Centre line \mathcal{C}_i , obtained directly from the image segmentation, displayed with contours \mathcal{B}'_{ik} . (b) Centre line \mathcal{C}'_i , after arc-length parametrization of \mathcal{C}_i , also shown with contours \mathcal{B}'_{ik} . (c) Centre line \mathcal{C}''_i and contours \mathcal{B}''_{ik} , after b-spline interpolation of \mathcal{C}'_i and \mathcal{B}'_{ik} , respectively. (d) Smooth surface \mathcal{X}_i , in wireframe mode. (e) Smooth surface \mathcal{X}_i , in shaded mode. The lines crossing the contours of (a), (b), and (c) indicate the position at which the sampling of the boundary starts (the region of the carina corresponds to the bottom contour).

by linearly combining the mean shape with a weighted version of the eigenvector matrix \mathbf{P} , that is

$$\mathbf{x} = \bar{\mathbf{x}} + \mathbf{P}\mathbf{b}. \quad (6.1)$$

The vector $\mathbf{b} = (b_1, b_2, \dots, b_N)^T$ is the set of parameters of the model and represents the contribution of each eigenvector. Since λ_i represents the variance of the training set along the eigenvector \mathbf{p}_i , the range $[-3\sqrt{\lambda_i}, +3\sqrt{\lambda_i}]$ is a suitable limit for b_i (Cootes et al., 1995).

In general, only the first t modes of variation are used, such that

$$\lambda_t = \frac{1}{\lambda_T} \sum_{i=1}^t \lambda_i \leq T_{max} \quad , \quad (6.2)$$

where λ_T is the total variance of the model. In fact, care should be taken in the case of the trachea, because if T_{max} is high the shape generated by the ASM may be too flexible and more susceptible to the attraction caused by the regions with stenosis during the registration of the ASM. Empirical experience showed that setting $T_{max} = 90\%$ proved to be good enough.

Figure 6.2 shows an example of the first 6 modes of variation captured by a model built with $N = 38$ healthy tracheas. For $T_{max} = 90\%$, t corresponded to the first 10 modes of variation. In the figure, the average shape appears in the middle, with shapes corresponding to $\pm 3\sqrt{\lambda_i}$ next to it. The mode with largest variation indicates changes in the calibre. The second mode represents global bending, and from then onwards different levels of local bending or twisting.

6.5 ASM Registration

ASMs can be registered to an object of the class they represent by adjusting the parameter vector \mathbf{b} . When the model is applied to an image, the registration is usually an iterative, edge-based search. It encompasses an initialization step, followed by rigid and non-rigid registration stages. In the latter two, the landmarks of the shape generated by the model, \mathbf{x} , are




Figure 6.2: From (a) to (f), the first 6 modes of variation. Each figure shows the average shape (middle), $-3\sqrt{\lambda_i}$ (left), and $+3\sqrt{\lambda_i}$ (right) along the mode.

moved along their normals towards high gradients corresponding to edges of the target, generating a candidate shape \mathbf{y} . Depending on the stage, either just pose or pose and shape are changed so as to let the model fit to \mathbf{y} . The registration is described in detail in the following.

6.5.1 Initialization

Since the registration is an iterative, edge based search, neighbouring organs and structures as well as noise may mislead the search. Therefore, the registration is dependent on the initial search location. This can be done automatically, but automatic initialization is a problem on its own and as such is out of the scope of the present work. At this point, the initialization of the registration is done manually, and the operator is given controls to interactively change the position, scale, and orientation of the average shape

of the model. Another option would be to let the user define the bounding box of the trachea, or mark the location of the vocal fold and of the carina (defining scale) and indicate their orientation. Whichever the method, the initialization should take in the order of seconds even for a non-experienced user, so that the average shape can be quickly and conveniently placed in the image, namely, near the target trachea.

6.5.2 Rigid Alignment

In this step, the shape generated by the model is rigidly aligned with the candidate shape of each iteration, where the latter is supposed to match high gradients, or edges, of the target trachea in the image. The search for high gradients is done at discrete steps within long profiles along the normals of the shape generated by the model. The location $\mathbf{p} = (x, y, z)$ that maximizes the function

$$f_j = |\nabla\mathbf{I}(\mathbf{p})| \left\langle \frac{\nabla\mathbf{I}(\mathbf{p})}{|\nabla\mathbf{I}(\mathbf{p})|}, \mathbf{n}_j \right\rangle \quad (6.3)$$

is chosen as the candidate location of each landmark j , where $j = 1 \dots n$, \mathbf{n}_j is the landmark's normal, $\nabla\mathbf{I}(\mathbf{p})$ is the gradient of the image² at location \mathbf{p} along \mathbf{n}_j , and $\langle \cdot \rangle$ denotes inner product. The rigid alignment stops when no significant change is made to the shape generated by the model at consecutive iterations or when a maximum number of iterations is reached. The purpose of this step is to find the best position and orientation of the average shape of the model with respect to the target shape. It is also expected that average shape will already be closer to the healthy regions of the trachea when the alignment converges.

6.5.3 Pose and Shape Registration

In the last step of the ASM registration, a non-rigid registration, which optimizes pose and shape, is started, and the landmarks of the shape generated by the model at the current iteration are again displaced along their

²Computed with sub-voxel accuracy using tri-linear interpolation.

normals. Once the candidate shape \mathbf{y} is projected into model space, which aligns \mathbf{y} with the average shape, a new set $\hat{\mathbf{b}}$ that defines the best fit of the model to \mathbf{y} is obtained by minimization of the squared error between \mathbf{y} and \mathbf{x} , represented by the following error function:

$$\xi(\mathbf{b}) = (\mathbf{y} - \mathbf{x})^T(\mathbf{y} - \mathbf{x}) . \quad (6.4)$$

Expanding Eq. (6.4) with Eq. (6.1) and minimising ξ with respect to \mathbf{b} results in:

$$\hat{\mathbf{b}} = \mathbf{P}^T(\mathbf{y} - \bar{\mathbf{x}}) . \quad (6.5)$$

This minimization is herein referred to as *StandardLS*. A new shape $\hat{\mathbf{x}}$ is generated from $\hat{\mathbf{b}}$, using Eq. (6.1), and is transformed back into the image space. The displacement of landmarks and the update of $\hat{\mathbf{b}}$ and $\hat{\mathbf{x}}$ are repeated until no significant changes are made to $\hat{\mathbf{x}}$ (Cootes et al., 1995).

Robust ASM Fitting

Although the registration of the model of healthy tracheas to a trachea with stenosis tends not to be affected by the local geometric variations of the stricture, the registration may still produce globally narrowed tracheas. As a result, fitting a healthy trachea shape model to a stenotic trachea is a challenge in itself. In other application domains, *Robust ASM Fitting* algorithms have been proposed to avoid the influence of non-target regions in the registration. In the following, some of these algorithms are reviewed, and a new Robust ASM Fitting algorithm, called *Fixed Landmarks*, is proposed for the case of tracheal stenosis.

Weighted Least Squares When the distribution of the residuals, $\{r_j | j = 1, \dots, n\}$, between \mathbf{x} and \mathbf{y} in Eq. (6.4) is not Gaussian, due to the presence of outliers, *StandardLS* may produce suboptimal results. The influence of outliers can be reduced by assigning weights to the contribution of each residual, modifying Eq. (6.4) to

$$\xi_w(\mathbf{b}) = (\mathbf{y} - \mathbf{x})^T \mathbf{W}(\mathbf{y} - \mathbf{x}) , \quad (6.6)$$

where \mathbf{W} is a diagonal matrix of weights. Minimising ξ_w with respect to \mathbf{b} yields

$$\hat{\mathbf{b}} = (\mathbf{P}^T \mathbf{W} \mathbf{P})^{-1} \mathbf{P}^T \mathbf{W} (\mathbf{y} - \bar{\mathbf{x}}) , \quad (6.7)$$

which is the basic formulation of weighted least squares (*WLS*) minimization.

From the above definition, it is clear that a good choice of weights is key for an effective use of *WLS*. In the field of Robust Statistics, estimators that are less affected by deviations from Gaussian or other model assumptions can be devised to further improve the effects of *WLS* (Huber, 1981). Rogers and Graham (2002) used robust statistics with ASMs in different medical applications and filled matrix \mathbf{W} with the Huber weighting function. Theobald et al. (2006) compared several weighting functions for landmark occlusion detection, among which the Talwar, the Cauchy, and the Gaussian weighting functions performed best. Figure 6.3 shows the four functions described. In all of them, σ is the standard deviation of the residuals, which can be estimated at each iteration from the median of their absolute values (Rogers and Graham, 2002).

$$w_{\text{huber}_j} = \begin{cases} 1, & r_j < \sigma \\ \sigma/|r_j|, & \sigma \leq r_j < 3\sigma \\ 0, & r_j \geq 3\sigma \end{cases} \quad w_{\text{talwar}_j} = \begin{cases} 1, & r_j < \sigma \\ 0, & r_j \geq \sigma \end{cases}$$

$$w_{\text{cauchy}_j} = \frac{1}{1 + \left(\frac{r_j}{\sigma}\right)^2} \quad w_{\text{gauss}_j} = \frac{1}{\sigma\sqrt{2\pi}} e^{-\frac{(r_j - \mu)^2}{2\sigma^2}}$$

Figure 6.3: Huber, Talwar, Cauchy, and Gaussian weighting functions.

Assuming that the rigid alignment above already leaves the average shape of the model near the healthy regions of the target, the outliers in the residuals immediately correspond to the regions with stenosis. Therefore,

it would only be necessary to find the appropriate weighting function to remove the influence of the narrowed regions on the registration.

Surface Extrapolation

In this approach, the purpose is to use the model to predict missing parts of the target shape. At each iteration k of the registration, $\mathbf{y}^{(k)} \approx (\bar{\mathbf{x}} + \mathbf{P}\hat{\mathbf{b}}^{(k)})|_{\mathcal{L}}$, where \mathcal{L} , of size m , denotes the set of landmarks of the model actually used. It is possible that $m \ll n$, where n is the total number of landmarks of the model. The parameter set $\hat{\mathbf{b}}^{(k)}$ is computed as in Eq. (6.5), but using only the components of \mathbf{P} and \mathbf{x} corresponding to the m target landmarks.

Rajamani et al. (2007) used extrapolation to predict the shape of the femur from manually sampled points during hip surgery. In their method, the m sampled points are matched to the nearest landmarks of the shape generated by the model at each iteration. Furthermore, a weighting term added to the error function restricts the deformation freedom of the ASM as m decreases, forcing the model to produce shapes similar to the average shape. In (Zachow et al., 2005), extrapolation was used to plan reconstructions of mandibular dysplasia. The ASM is registered to parts of the mandible that are considered as being regularly shaped.

In the case of the trachea, the m points could be those corresponding to the healthy regions. These regions could be manually determined as in (Rajamani et al., 2007; Zachow et al., 2005), but this could be impractical to be done on a 3D surface. Assuming the rigid alignment places the average shape near the healthy regions of the trachea, the correspondence could instead be computed dynamically, with a closest point criterion between the shape generated by the model and the candidate shape of each iteration. The $(n - m)$ candidate landmarks that are too far from the model shape, could then be removed from the computation of $\hat{\mathbf{b}}$.

Fixed Landmarks

The proposed *Fixed Landmarks* somewhat borrows from the methods above by setting a restriction on the displacement of landmarks, which defines the candidate shape \mathbf{y} . Namely, if the candidate locations are not within a short threshold distance $d > 0$ from \mathbf{x} , the corresponding landmarks of \mathbf{y} remain fixed, while the other landmarks are allowed to move as usual.

Let $\mathbf{x}_{v_j} = (x_{j1}, x_{j2}, x_{j3})^T$ be a landmark in a 3n-dimensional shape vector $\mathbf{x} = (x_{11}, x_{12}, x_{13}, \dots, x_{n1}, x_{n2}, x_{n3})^T$. Let $\hat{\mathbf{x}}^{(k)} = \bar{\mathbf{x}} + \mathbf{P}\hat{\mathbf{b}}^{(k)}$ be the shape generated by the model at iteration k of the non-rigid registration. Let $\mathbf{y}^{(k+1)}$ be the candidate shape generated by displacing the landmarks of $\hat{\mathbf{x}}^{(k)}$, assuming that $\mathbf{y}^{(k+1)}$ has already been projected into model space and aligned with $\bar{\mathbf{x}}$.

Let now

$$d\mathbf{y}_{v_j}^{(k+1)} = \begin{cases} \mathbf{y}_{v_j}^{(k+1)} - \hat{\mathbf{x}}_{v_j}^{(k)}, & |\mathbf{y}_{v_j}^{(k+1)} - \hat{\mathbf{x}}_{v_j}^{(k)}| \leq d, \\ \mathbf{0}, & \text{otherwise} \end{cases}$$

with $d\mathbf{y}^{(k+1)}$ the concatenation of all $d\mathbf{y}_{v_j}^{(k+1)}$.

Using the definitions above, Eq. (6.5) can be rewritten as:

$$\hat{\mathbf{b}}^{(k+1)} = \mathbf{P}^T(\hat{\mathbf{x}}^{(k)} + d\mathbf{y}^{(k+1)} - \bar{\mathbf{x}}). \quad (6.8)$$

Let us now group the $d\mathbf{y}_{v_j}^{(k+1)} = \mathbf{0}$ and the corresponding rows of $\hat{\mathbf{x}}^{(k+1)}$ and $\bar{\mathbf{x}}$, and columns of \mathbf{P}^T , using row and column permutations, respectively. This results in two subsets of landmarks, \mathcal{L}' and \mathcal{L}'' , of sizes n' and n'' , respectively, such that $d\mathbf{y}^{(k+1)}|_{\mathcal{L}''} = \mathbf{0}$. Eq. (6.8) can then be split into

$$\hat{\mathbf{b}}^{(k+1)} = \left[\mathbf{P}^T(\hat{\mathbf{x}}^{(k)} + d\mathbf{y}^{(k+1)} - \bar{\mathbf{x}}) \right]_{\mathcal{L}'} + \left[\mathbf{P}^T(\hat{\mathbf{x}}^{(k)} + d\mathbf{y}^{(k+1)} - \bar{\mathbf{x}}) \right]_{\mathcal{L}''}, \quad (6.9)$$

which does not affect the result. Since $d\mathbf{y}^{(k+1)}|_{\mathcal{L}''} = \mathbf{0}$, the parameters are finally given by

$$\hat{\mathbf{b}}^{(k+1)} = \left[\mathbf{P}^T(\mathbf{y}^{(k+1)} - \bar{\mathbf{x}}) \right]_{\mathcal{L}'} + \left[\mathbf{P}^T(\hat{\mathbf{x}}^{(k)} - \bar{\mathbf{x}}) \right]_{\mathcal{L}''}. \quad (6.10)$$

Eq. (6.10) shows that $\hat{\mathbf{b}}^{(k+1)}$ is determined by both the displaced landmarks $\mathbf{y}^{(k+1)}|_{\mathcal{L}'}$ and the landmarks $\hat{\mathbf{x}}^{(k)}|_{\mathcal{L}''}$, which remained fixed³. Consequently, when computing

$$\hat{\mathbf{x}}^{(k+1)} = \bar{\mathbf{x}} + \mathbf{P}\hat{\mathbf{b}}^{(k+1)} \quad , \quad (6.11)$$

$\hat{\mathbf{x}}^{(k+1)}$ will be the best fit, in a least squares minimization sense, to $\mathbf{y}^{(k+1)}|_{\mathcal{L}'}$ and $\hat{\mathbf{x}}^{(k)}|_{\mathcal{L}''}$. Figure 6.4 illustrates the whole concept. Provided that there are enough healthy areas around regions with stenosis, the fixed landmarks force the shape generated by the model to remain far from those regions, while enabling correct matches at the healthy areas. Compared to the previous methods, the *Fixed Landmarks* method on the one hand immediately removes the influence of the outliers on the registration. On the other hand, not removing the fixed landmarks from the computation of $\hat{\mathbf{x}}^{(k+1)}$ (as opposed to the *Surface Extrapolation* method) in fact aids the registration in yielding the desired healthy tracheal shape.

Figure 6.5 shows how the *Fixed Landmarks* method acts on two examples, one with simulated data and the other with clinical data. From left to right, the figure shows the image with the stenotic trachea, two non-consecutive steps of the rigid alignment and two non-consecutive steps of the *Fixed Landmarks* method. The green shape indicates the shape $\hat{\mathbf{x}}^{(k)}$ and the red dots represent the landmarks of the candidate shape $\mathbf{y}^{(k+1)}$, as described above. Note that during the rigid alignment all the edges of the target trachea attract the landmarks, since the search profiles are long. This is necessary in order to find the correct position and orientation of the target. When the rigid alignment converges, the *Fixed Landmarks* immediately starts. Now, the landmarks for which $|d\mathbf{y}_{v_j}^{(k+1)}| > d$ remain fixed with respect to $\hat{\mathbf{x}}^{(k)}$. Note that they remain far from the regions with stenosis. As the registration proceeds, the green shape takes the form of the target shape, guided by the healthy areas of the target and by the landmarks $\hat{\mathbf{x}}^{(k)}|_{\mathcal{L}''}$, which are those that remained fixed. Figures 6.5(e) and 6.5(j) represent the final result, i.e., when the *Fixed Landmarks* converges.

³Note that \mathcal{L}' and \mathcal{L}'' can be different at each iteration.




Figure 6.4: The *Fixed Landmarks* concept: (a) $\hat{\mathbf{x}}^{(k)}$ (dashed) is attracted by high gradients of the stenotic trachea in the image (thick lines) – the arrows indicate the landmark displacements $d\mathbf{y}^{(k+1)}$. (b) the set of landmarks that remain fixed, $\hat{\mathbf{x}}^{(k)}|_{\mathcal{L}''}$, because the displacement was longer than d . (c) $\hat{\mathbf{x}}^{(k+1)}$ – the fixed landmarks keep the shape far from the stenotic area.

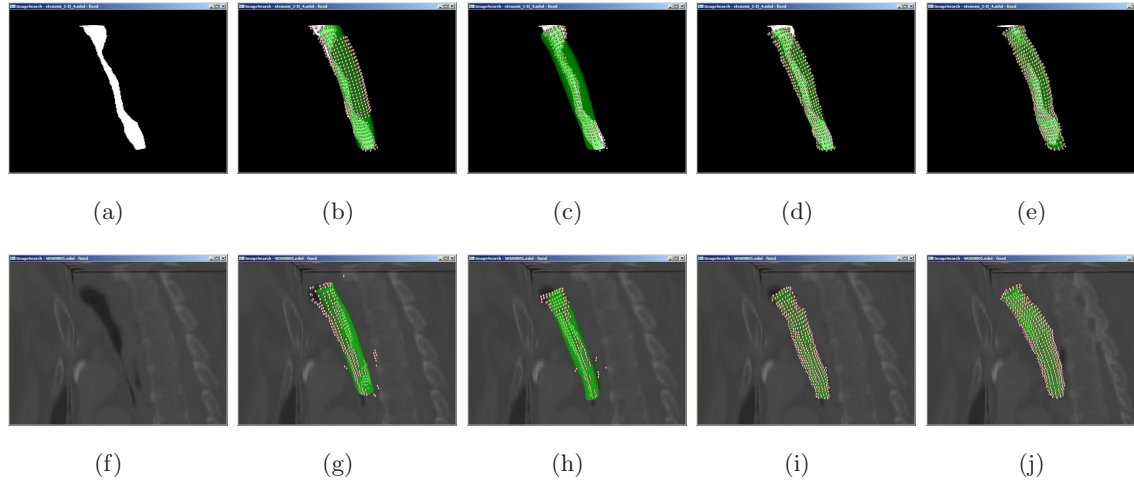


Figure 6.5: Examples of the *Fixed Landmarks* method in action with simulation data (top) and clinical data (bottom). The green shape is the shape generated by the model at one iteration and the red dots indicate the landmarks of the candidate shape, which match high gradients in the image. From left to right, the figures show the trachea with stenosis, two non-consecutive steps of the rigid alignment and two non-consecutive steps of the *Fixed Landmarks*.

6.6 Experiments

Experiments on simulation as well as on clinical data were carried out in order to evaluate the proposed *Fixed Landmarks* method and to compare it with the different registration techniques. With the simulation data, ground truths of healthy tracheas were formally established in order to quantitatively evaluate the method. The experiments with clinical data were used for qualitative evaluations.

To build the ASM, $N = 9$ healthy tracheas were used, each with $n = 1024$ landmarks. The low-dose, chest CT scans of their respective patients were obtained from pulmonary medication studies carried out at the University Hospital of Antwerp, Belgium, and patient data were anonymized before the images were used. For this set of experiments, the tracheas were segmented from the images and converted to 3D surfaces with the marching cubes algorithm (Lorensen and Cline, 1987). The 3D surfaces were later supplied to a correspondence optimization algorithm (Huysmans et al., 2010), which eventually produced the shapes used in the model.

6.6.1 Quantitative Comparison

In order to quantitatively compare the registration methods, a large set of leave-one-out tests was performed using simulation data. First, for each of the $N = 9$ healthy tracheas, 72 phantoms of stenosis – 24 anteriorly located (A), 24 posteriorly (P), and 24 roughly symmetrically narrowed (S) – were created with a software tool developed specifically for this purpose. In this tool, a healthy trachea previously segmented from the original CT image (using the method presented in Chapter 5) is visualized in the axial, sagittal, and coronal planes, and also in three dimensions. The user then places a configurable erosion mask on top of the region where the stenosis is to be inserted. The level of stenosis is specified by the overlap between the mask and the image. When overlapping is satisfactory, the user triggers the erosion operator, which in turn removes from the original image the areas below the mask. Figure 6.6 shows a snapshot of the user interface and of a phantom obtained with the application.




Figure 6.6: Application to build phantoms of tracheal stenosis.

The created phantoms followed the categories of Table 6.1 and Table 6.2, based on (Freitag et al., 2007), and were validated by an expert in the pulmonology field. For example, phantom 5.II-2S represents a symmetrical stenosis between 25% and 50% along the middle third of trachea number 5, as seen in Figure 6.6.

For each run of the $T = N$ leave-one-out tests, the model was built with $N - 1$ tracheas and was then registered to the phantoms created from the trachea not present in the training set. In all tests, the average shape of the ASM was initially roughly placed near the target trachea in the image.

The objective of this set of experiments was to register the model to

Table 6.1: Categories of stenosis based on location and length.

CATEGORY	LOCATION AND LENGTH
I	Upper third of the trachea
II	Middle third of the trachea
III	Lower third of the trachea
I-II	Upper third extending to middle third
II-III	Middle third extending to lower third
I-III	Upper third extending to lower third

Table 6.2: Categories of stenosis based on its degree of narrowing.

CATEGORY	DEGREE
1	<25%
2	26–50%
3	51–75%
4	>75%

the phantoms using all the registration methods of Section 6.5.3 and to measure the distance between the estimated tracheas and their originally healthy counterparts. The algorithm proposed in (Aspert et al., 2002) was employed to compute errors between surfaces using the Hausdorff distance.

The iteration limit for the registration was set to 200. The minimum squared error between shapes generated at subsequent iterations, i.e., $\xi_r^{(k)} = (\mathbf{x}^{(\hat{k})} - \mathbf{x}^{(\hat{k}-1)})^T (\mathbf{x}^{(\hat{k})} - \mathbf{x}^{(\hat{k}-1)})$, was set to 10^{-7}mm^2 . In the *Fixed Landmarks*, the landmarks were displaced within a distance $d = 1 \text{mm}$ along their normals. For the *Surface Extrapolation*, the m landmarks of the candidate shape that guided the deformations were those that remained near the target surface ($d \leq 1 \text{mm}$) after each iteration of the registration. In addition, the weighting term defined in (Rajamani et al., 2007) was dropped, since shapes similar to the average shape of the model are very unlikely to produce a good estimation of a specific healthy trachea. In this way,

the *Surface Extrapolation* becomes equivalent to using *WLS* with a step function yielding binary weights. Besides the approaches of Section 6.5.3, the experiments included the *StandardLS*, i.e., not distinguishing between healthy and stenotic areas on the target surfaces.

6.6.2 Qualitative Comparison

In addition to the simulation experiments, a retrospective study was made with chest CT scans from 3 patients. The use of the CT scans was approved by the ethics committee of the Ghent University Hospital (doc. ECUZG2009/140), Belgium, and patient data were anonymized before the images were used in the experiments. The 3 patients had stenosis with the following characteristics:

Patient 1 – Severe posterior stenosis along the lower two thirds of the trachea,

Patient 2 – Side-to-side stenosis along the lower two thirds of the trachea,

Patient 3 – Severe symmetrical stenosis along the lower two thirds of the trachea,

No preprocessing was applied to any of the 3D CT images. In addition, they were very anisotropic in the axial direction, with pixel resolution, in mm, (0.62, 0.62, 3.00), (0.98, 0.98, 5.00), and (0.44, 0.44, 3.00) respectively.

For each patient, the initial position of the average shape of the model was manually set for one registration method, recorded, and then replicated for all other methods. The results were reviewed by an expert in the pulmonology field in order to qualitatively compare all methods with respect to the estimated healthy trachea.

The ASM built for these experiments contained all the $N = 9$ healthy tracheas. As before, the maximum number of iterations was set to 200, the minimum $\xi_r^{(k)}$ was set to 10^{-7}mm^2 , and $d = 1\text{mm}$. *WLS* and *StandardLS* were used in the same way as in the experiments with simulation data.

6.7 Results

6.7.1 Quantitative Comparison

For the quantitative comparison between the registration methods using the simulation data, the phantoms of each healthy trachea were subdivided into $G = 10$ groups, according to the categories defined in tables 6.2 and 6.2, each with a different size S_g . The reasoning behind this subdivision is to show how the methods behaved relative to variations in location, length, and degree of stenosis across the whole set of T leave-one-out tests. Let us then define, for a test instance t , $\delta_{\max_{gtp}}$ and $\delta_{\text{mean}_{gtp}}$ as the maximum and mean distances, respectively, between the estimated trachea for phantom p , of group g , and its original, healthy equivalent. As stated in Section Section 6.6.1, these distances are obtained using the algorithm proposed in (Aspert et al., 2002). Afterwards,

$$\bar{\delta}_{\max_{gt}} = \frac{1}{S_g} \sum_{p=1}^{S_g} \delta_{\max_{gtp}} \quad \text{and} \quad \bar{\delta}_{\text{mean}_{gt}} = \frac{1}{S_g} \sum_{p=1}^{S_g} \delta_{\text{mean}_{gtp}} \quad (6.12)$$

can be calculated as the average, per-group maximum and mean distances, respectively, for one test instance.

The final comparison used each group's average maximum error, μ_{\max_g} , and average mean error, μ_{mean_g} , for each method, across the whole set of T leave-one-out tests. That is,

$$\mu_{\max_g} = \frac{1}{T} \sum_{t=1}^T \bar{\delta}_{\max_{gt}} \quad \text{and} \quad \mu_{\text{mean}_g} = \frac{1}{T} \sum_{t=1}^T \bar{\delta}_{\text{mean}_{gt}} \quad (6.13)$$

The results and their respective standard deviation bars are shown in Figure 6.7.

As expected, the *StandardLS* did not perform well. The influence of the stenotic regions on the registration indeed made the resulting shape much narrower or deformed than desired. The use of *WLS* brought some improvements, but not enough to completely remove the influence of stenotic

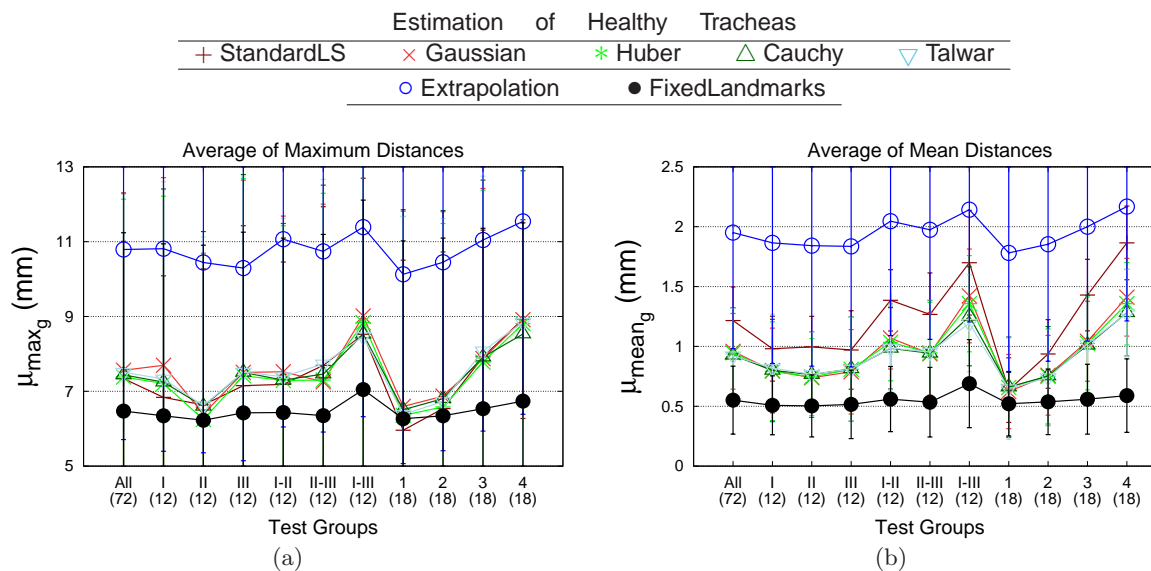


Figure 6.7: Per-group μ_{\max_g} (a), μ_{mean_g} (b), and respective standard deviation bars for each method across the whole set of leave-one-out tests. Along the horizontal axis, the number of phantoms in the test group, S_g , is shown in parentheses.

regions. The problem with the weighting approaches is the difficulty in finding the proper weight assigning function to act only on the regions with stenosis. If the weighting scheme is too tight, the shape may not deform enough, remaining similar to the mean shape. If it is too loose, the shape may be strongly attracted by the areas with stenosis. Regarding the *Surface Extrapolation*, using only the points near the target surface to guide the deformations eventually resulted in few, very localized points, especially in the most severe cases. Without a stronger clue to indicate the shape to be obtained in a global level, the method could not converge to the desired result, which explains its poor performance. As a conclusion, the *Fixed Landmarks* was the best registration method. It is especially worth noting how other methods performed worse as the length and degree of stenosis increased, while the *Fixed Landmarks* was hardly affected. Its μ_{mean_g} 's remained near 0.5mm in all but one test group, I-III, which represented the longest and most severe types of stenosis in the simulation data.

Figure 6.8 presents an example of the estimation of the healthy trachea for phantom II-4A generated from one of the healthy tracheas used in our experiments, using the *StandardLS*, *Gauss WLS*, *Huber WLS*, *Surface Extrapolation*, and *Fixed Landmarks* methods. The dashed, outermost silhouette in each case represents the original healthy trachea and the *Fixed Landmarks* yielded the best approximation to it.

6.7.2 Qualitative Comparison

As mentioned before, the results from the experiments with clinical data were reviewed by an expert in the pulmonology field. It can be seen in Figure 6.9 that the *Fixed Landmarks* produced very plausible healthy tracheas. They have an acceptable calibre and follow the curvature of the patient's trachea. Figure 6.10 shows the results of other methods applied to the CT scan of **Patient 1**. It can be seen how the severely narrowed trachea influenced the registration and either made the estimated tracheas too narrow and deformed or led them astray.




Figure 6.8: Shape estimation for phantom II-4A of one healthy trachea from the simulation experiments, using, from left to right, the *StandardLS*, *Gauss WLS*, *Huber WLS*, *Surface Extrapolation*, and *Fixed Landmarks* methods. The dashed, outermost silhouettes represent the healthy trachea used to build the phantom.

6.8 Discussion

As with any other open tubular surface, the registration of the ASM to tracheas may be hindered by two degrees of freedom. In the first, given that there may be no cues indicating the top and bottom limits of the trachea, the shape generated by the model may freely drift and eventually converge to unwanted locations. This can easily be circumvented by forcing limits to these displacements. Here, the limits are derived from the bounding box of the average shape of the ASM after the operator chooses its correct location and scale in order to initialize the registration. Setting the limits to 1 voxel




Figure 6.9: Results of the estimation of the healthy trachea with the *Fixed Landmarks* for patients 1, 2, and 3, from left to right. The estimated trachea, in green, is shown in the CT scan of the patient, overlaid on their segmented stenotic trachea.




Figure 6.10: *Gauss WLS* (a), *Surface Extrapolation* (b), and *StandardLS* (c) registration methods applied to the CT scan of Patient 1. The estimated trachea, in green, is placed in the image, overlaid on the segmented stenotic trachea. With *Surface Extrapolation*, the estimated surface failed to match the lower part of the trachea. With other methods, the surface was too narrow or deformed (compare with Figure 6.9(a)).

above and below the bounding box is usually enough.

The second degree of freedom is related to rotation. A perfect tubular surface has no cues indicating correspondence with respect to the ASM registered to it. Consequently, the shape generated by the model may freely rotate around the target's centre line until convergence, possibly when the iteration limit is reached. In principle, some sort of restriction could be forced on rotations around the centre line of the trachea. However, a few cues are expected to be present in the case of the trachea, such as the carina. In addition, the operator can chose the initial orientation of the shape before the registration, as discussed earlier. Thus, no restrictions on rotations are set during the registration.

Chapter 7

Segmentation of Tracheal Stenosis

Brian: “Please, please, please listen!

I’ve got one or two things to say.”

The Crowd: “Tell us! Tell us both of them!”

Monty Python’s Life of Brian

“Somewhere, something incredible is waiting to be known.”

Carl Sagan (1934 – 1996)

Parts of this chapter appear in:

- R. Pinho, K. G. Tournoy, R. Gosselin, and J. Sijbers. Assessment and stenting of tracheal stenosis using deformable shape models. *Medical Image Analysis*, (Under Review), 2010a.
- R. Pinho, K. G. Tournoy, and J. Sijbers. A decision support system for the treatment of tracheal stenosis. In *Proc. of Workshop on Discrete Geometry and Mathematical Morphology (WADGMM)*, pages 72–76, Istanbul, Turkey, 2010b.
- R. Pinho, K. G. Tournoy, and J. Sijbers. Computer-aided assessment and stenting of tracheal stenosis. In Ayman El-Baz, editor, *Computer Aided Diagnosis of Lung Imaging*, volume (in press). American Scientific Press, 2011.
- L. Vanacken, R. Pinho, J. Sijbers, C. Raymaekers, and K. Coninx. Force feedback to assist active contour modelling for tracheal stenosis. (In Preparation), 2010.

7.1 Introduction

The healthy trachea estimated above is the basis for the automatic assessment of stenosis. This assessment is performed via a comparison between the cross-sectional area profiles of the estimated healthy trachea and the narrowed trachea in the image. It is thus necessary to segment the narrowed trachea in question from the given CT image, which is done with an active contour model tailored for this purpose. The steps involved in this segmentation are described in the following.

7.2 Active Contour Models

ACMs, commonly known as Snakes, are curves defined within an image domain that are able to move under the influence of internal forces derived

from the curve itself and of external forces derived from the image data. The internal and external forces are defined in such a way that the curve will register to an object boundary or other desired features within an image. As defined by Kass et al. (1988), a snake can be represented in 2D by a curve $\mathbf{v}(s) = (x(s), y(s))$, $s \in [0, 1]$, responding to an energy functional of the form:

$$E = \int_0^1 [\kappa E_{int}(\mathbf{v}(s)) + (1 - \kappa) E_{ext}(\mathbf{v}(s))] ds \quad , \quad (7.1)$$

where $\kappa \in [0, 1]$ is a weighting factor.

The internal energy E_{int} commonly restricts the deformations taking into account curvature and elasticity. Considering elasticity, the functional prevents the curve from breaking apart, maintaining its continuity. As for elasticity, the internal energy avoids the appearance of sharp corners. The external energy represents the gradient of an image I convolved with a Gaussian function G at scale σ , which causes the curve to be attracted by contours with high image gradients.

The objective is then to minimize Eq. (7.1), making the system a force balance equation of the form:

$$\mathbf{F}_{int} + \mathbf{F}_{ext} = \mathbf{0} \quad , \quad (7.2)$$

where

$$\mathbf{F}_{int} = -\nabla E_{int} \quad (7.3)$$

and

$$\mathbf{F}_{ext} = -\nabla E_{ext}. \quad (7.4)$$

The minimization is solved iteratively, often in the discrete domain after discretization of $\mathbf{v}(s)$. The expected result is a curve that matches the high gradients of the image while being restricted by the internal constraints, according to the assigned weighting factors.

A classical problem with ACMs, however, is initialization. The starting curve should be placed near the features to be detected in the image, otherwise the capture range of the external energy may not extend enough to attract the curve. In addition, in the absence of external energies, the internal energies cause the curve to shrink into itself.

Several methods have been proposed in order to overcome the limitations of the original ACM formulation. Cohen (1991), for instance, added a balloon force to the energy formulation that pushes the curve outwards in the direction of its normals, avoiding it to shrink. Later, Cohen and Cohen (1993) proposed to increase the capture range of the external energy by using a force based on a map of distances from any point in the image to the nearest edge. The Gradient Vector Flow (GVF) of Xu and Prince (1998) and the recent Magnetostatic Active Contour Model (MAC) of Xie and Mirmehdi (2008) are other good examples of methods that affect the external force.

The following section describes the method proposed for assessment of tracheal stenosis and prediction of patient-specific stents using ASMs of healthy tracheas and ACMs.

7.3 Segmentation of the Stenotic Trachea

The ideas presented in Section 7.2 naturally extrapolate to 3D and can easily be adapted to a discrete domain. Within the context of the proposed application, a discrete surface can be defined as $\mathcal{X} = (\mathcal{L}, \mathcal{T})$, where \mathcal{L} is the set of landmarks and \mathcal{T} is the set of triangles connecting the points of \mathcal{L} . The internal and external forces are defined locally, and the minimization iterates through all the points of the surface. The process stops when no significant deformation has been applied to the surface. In the following, the integration of the ACM into the proposed method is explained.

7.3.1 Initialization

The ACM is initialized with the estimated healthy shape obtained with the method of Chapter 6. As a result, the shape generally shrinks under the influence of the external forces, which guide the boundaries of the surface towards the edges of the stenotic trachea in the image. In this setting, it is not necessary to employ a balloon force. The internal forces control the deformations and do not allow the shape geometry to degenerate. If the stricture is extremely severe, the model is eventually capable of providing




Figure 7.1: Distance map I_D used in the establishment of the external force of the ACM.

a good approximation to the tracheal wall at that area, while keeping the surface continuous.

7.3.2 External Force

In this chapter, the ideas first presented in Cohen and Cohen (1993) are adapted in order to create an external force to guide the deformation of the surface. At first, the original image, I , is thresholded to segment air (at -200 Hounsfield Units) and the result is then inverted. This will generate two binary images, I_{B1} and I_{B2} . Next, a 3-4-5 chamfer distance transform is applied to each binary image, yielding two distance maps, I_{D1} and I_{D2} , respectively. In the last step, both distance transforms are combined into a new distance map $I_D = I_{D1} + I_{D2}$. Figure 7.1 illustrates the result of such procedure.

In general, a Gaussian smoothing filter at some scale σ can be convolved with I prior to the above operations. This, however, may cause significant loss of detail, especially at areas of severe stenosis, which might remove important intensity transitions between the tracheal lumen and neighbouring structures. Therefore, the method does not use the smoothing filter. Nevertheless, such filter can still be applied if the noise levels of the image

are considered too high.

Eventually, the external force term of Eq. (7.2) applied to the landmark \mathbf{x}_{v_j} of \mathcal{X} is defined as

$$\mathbf{F}_{\text{ext}_j} = -\frac{|\nabla I_D(\mathbf{x}_{v_j})|}{M} \nabla I_D(\mathbf{x}_{v_j}) , \quad (7.5)$$

where M is the maximum gradient magnitude in I_D .

7.3.3 Internal Forces

The internal forces, which control stretching and bending, have similar directional components. Their scalar components are obtained according to the local deviation from the minimization conditions. Local directional and scalar components are altogether derived from the positions of the points w.r.t. their neighbours. Since there is no balloon force in this work, differentiating between border and non-border points of the surface (the former being located at the surface's open ends) is important in order to reduce the shrinking effects of the internal forces.

The internal elastic force tries to keep the points of \mathcal{X} equally spaced by moving them towards a central point relative to its neighbours, as illustrated in Figure 7.2. This force is defined as

$$\mathbf{F}_{\text{elast}_j} = D_j \frac{\mathbf{d}_{\text{elast}_j}}{|\mathbf{d}_{\text{elast}_j}|} . \quad (7.6)$$

The directional component $\mathbf{d}_{\text{elast}_j}$ is given by

$$\mathbf{d}_{\text{elast}_j} = \mathbf{c}_{\text{elast}_j} - \mathbf{x}_{v_j} , \quad (7.7)$$

where $\mathbf{c}_{\text{elast}_j}$ is a point at a central location relative to the neighbours of \mathbf{x}_{v_j} , such that:

$$\mathbf{c}_{\text{elast}_j} = \begin{cases} \delta_j(\mathbf{x}_{v_j} + \mathbf{n}_j), & \text{if } \mathbf{x}_{v_j} \text{ is a border point} \\ \frac{1}{|\mathcal{O}_r(\mathbf{x}_{v_j})|} \sum_{\mathbf{p} \in \mathcal{O}_r(\mathbf{x}_{v_j})} \mathbf{p}, & \text{otherwise} \end{cases} , \quad (7.8)$$

In the above, δ_j is the shortest distance from \mathbf{x}_{v_j} to the bisector plane of the line connecting its two immediate border neighbours. This plane is specified by the normal vector \mathbf{n}_j . \mathcal{O}_r , in turn, is the ordered set of points directly connected to \mathbf{x}_{v_j} , namely, its One-ring.

The scalar component, D_j , of $\mathbf{F}_{\text{elast}_j}$ is a normalized measure of how much \mathbf{x}_{v_j} deviates from $\mathbf{c}_{\text{elast}_j}$:

$$D_j = \left(1 - \frac{\sum_{\mathbf{p} \in \mathcal{N}(\mathbf{x}_{v_j})} |\mathbf{p} - \mathbf{c}_{\text{elast}_j}|}{\sum_{\mathbf{p} \in \mathcal{N}(\mathbf{x}_{v_j})} |\mathbf{p} - \mathbf{x}_{v_j}|} \right). \quad (7.9)$$

In Eq. (7.9), \mathcal{N} contains the two immediate border neighbours of \mathbf{x}_{v_j} if \mathbf{x}_{v_j} is a border point, or $\mathcal{N} = \mathcal{O}_r$ otherwise.

The internal bending force tries to move the points of the surface in such a way that they are all coplanar. This force is defined as

$$\mathbf{F}_{\text{bend}_j} = K_{G_j} \frac{\mathbf{d}_{\text{bend}_j}}{|\mathbf{d}_{\text{bend}_j}|}. \quad (7.10)$$


Its directional component, $\mathbf{d}_{\text{bend}_j}$, also varies according to the location of point \mathbf{x}_{v_j} , that is:

$$\mathbf{d}_{\text{bend}_j} = \begin{cases} \mathbf{d}_{\text{elast}_j}, & \text{if } \mathbf{x}_{v_j} \text{ is a border point or } \mathbf{d}_{\text{elast}_j} \perp \mathbf{n}_{v_j} \\ \text{sign}(\langle \mathbf{d}_{\text{elast}_j}, \mathbf{n}_{v_j} \rangle) \mathbf{n}_{v_j}, & \text{otherwise} \end{cases}, \quad (7.11)$$

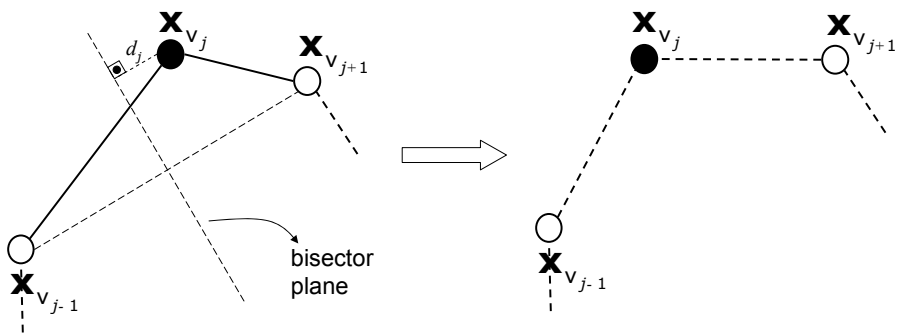
where \mathbf{n}_{v_j} is the normal of \mathcal{X} at point \mathbf{x}_{v_j} . This means that $\mathbf{F}_{\text{bend}_j}$ will either move \mathbf{x}_{v_j} as $\mathbf{F}_{\text{elast}_j}$ does or it will move the point along its normal. The latter case is illustrated in Figure 7.3

The scalar component, K_{G_j} , of $\mathbf{F}_{\text{bend}_j}$ is taken from the normalized, discrete Gaussian curvature of \mathcal{X} at \mathbf{x}_{v_j} and is given by

$$K_{G_j} = \left(1 - \frac{1}{2\pi} \sum_{\mathbf{p}_1, \mathbf{p}_2 \in \mathcal{N}(\mathbf{x}_{v_j})} \angle_{\mathbf{x}_{v_j}} \mathbf{p}_1 \mathbf{p}_2 \right). \quad (7.12)$$



(a)



(b)

Figure 7.2: The elastic force acting on a point \mathbf{x}_{v_j} of the surface. (a) The general case, in which the point is moved to a central location relative to its One-ring. (b) When \mathbf{x}_{v_j} is at one of the open ends of the cylindrical surface.

In Eq. (7.12), \mathcal{N} is defined as in Eq. (7.9) and $\angle \mathbf{x}_{v_j} \mathbf{p}_1 \mathbf{p}_2$ determines the angle between edges connecting two consecutive points of \mathcal{N} with \mathbf{x}_{v_j} .




Figure 7.3: The general case of the application of the bending force, so that sharp angles are removed from the surface.

Finally, the internal force term of Eq. (7.2) at landmark \mathbf{x}_{v_j} of \mathcal{X} is defined as

$$\mathbf{F}_{\text{int}_j} = \gamma \mathbf{F}_{\text{elast}_j} + (1 - \gamma) \mathbf{F}_{\text{bend}_j} , \quad (7.13)$$

and for each iteration k of the deformation algorithm of the ACM,

$$\mathbf{x}_{v_j}^{(k)} = \mathbf{x}_{v_j}^{(k-1)} + \kappa \mathbf{F}_{\text{int}_j} + (1 - \kappa) \mathbf{F}_{\text{ext}_j} . \quad (7.14)$$

7.3.4 Surface Manipulation

The deformations caused by the forces applied to the surface may severely modify its geometry. As a consequence, surface manipulation strategies involving edge collapses and splits become necessary. In order to approximately maintain the original surface geometry, in terms of number of points

and triangular connections, the average distance between points of the surface is computed prior to the execution of the ACM. After each iteration of the deformation algorithm through all the points, edges whose length is more than twice the pre-computed average distance are split. Likewise, edges with less than half the pre-computed length are collapsed. Care must be taken with these operations, especially at the borders of the surface, so that no invalid triangles are generated. The algorithm used was adapted from Botsch et al. (2002).

7.4 Discussion

One may argue that the segmentation of the stenotic trachea, as opposed to the approach adopted in this work, could be easily accomplished with, e.g., a 3D region growing method, as the one in Chapter 5, or a 2D slice based ACM to segment the tracheal lumen, as described in (Triglia et al., 2002). However, such segmentation may be problematic if the stenosis is too severe, given little or no presence of tracheal lumen in the image, as also reported in (Triglia et al., 2002). The proposed solution in that work was to manually reconstruct the missing parts of the trachea.

Patient 4 was an example of such a severely narrowed trachea among our experiments with clinical data. The lumen around the narrowed region was barely visible in the CT image and there was only a hint of the tracheal wall. Figure 7.4(a) shows the result of a simple 3D region growing applied to this CT scan. Note that a considerable part of the trachea is missing in the segmentation. The advantage of using the ACM to segment the trachea as proposed in this work is that it is able to reconstruct the tracheal wall for absent regions such as the one in this example, without manual intervention, thanks to the geometric constraints imposed by the internal forces of the model. The result of using the ACM is shown in Figure 7.4(b).

Nonetheless, it is worth mentioning that **Patient 4** was an especially difficult case. The problem here was that the oesophagus of the patient was also visible in the image, exactly behind the narrowed region of the trachea. Since the initial shape for the ACM is the healthy estimation




Figure 7.4: Illustration of a difficult case (**Patient 4**) due to severe stenosis (a) and how the ACM managed to reconstruct the missing part of the tracheal wall (b).

given by the ASM, the edges of the oesophagus in the image incurred a strong external force to the ACM. In fact, so strong that the ACM could only be successfully applied to this image after a pre-processing step to manually segment the trachea. The result of this pre-processing step was similar to that in Figure 7.4(a). After the pre-processing, the ACM was successfully applied to obtain the complete shape of the narrowed trachea, generating the result of Figure 7.4(b).

In conclusion, there might be difficult situations that require some level of operator intervention even when using the proposed ACM. Still, the use of the proposed ACM not only naturally fits into the proposed application framework, but is also able to overcome such difficult situations.

Chapter 8

Assessment and Stenting of Tracheal Stenosis

“Wabbits wove cawwots.”

Elmer Fudd

“If I have seen further than others,
it is by standing upon the shoulders of giants.”

Isaac Newton (1643 – 1727)

Parts of this chapter appear in:

- R. Pinho, K. G. Tournoy, R. Gosselin, and J. Sijbers. Assessment and stenting of tracheal stenosis using deformable shape models. *Medical Image Analysis*, (Under Review), 2010a.
- R. Pinho, K. G. Tournoy, and J. Sijbers. A decision support system for the treatment of tracheal stenosis. In *Proc. of Workshop on Discrete Geometry and Mathematical Morphology (WADGMM)*, pages 72–76, Istanbul, Turkey, 2010b.
- R. Pinho, K. G. Tournoy, and J. Sijbers. Computer-aided assessment and stenting of tracheal stenosis. In Ayman El-Baz, editor, *Computer Aided Diagnosis of Lung Imaging*, volume (in press). American Scientific Press, 2011.

8.1 Introduction

This chapter shows that once healthy and narrowed versions of the trachea are obtained, the assessment of the stenosis is straightforward. It is achieved by a simple comparison of the area profiles of the two surfaces along their centre lines. This comparison determines the parameters of the stenosis (start, end, degree of narrowing). The parameters of the stent suggested for a specific patient (deployment location, length, external diameter) are then immediately derived from the parameters of the stenosis. The proposed algorithms for the automatic assessment and stenting of stenosis are detailed in the following sections. An extensive set of experiments to validate the method is presented in Section 8.4

8.2 Assessment of Stenosis

Let $\mathcal{S}_h = (\mathcal{X}_h, \mathcal{T}_h)$ be the triangular surface corresponding to the estimated healthy trachea and $\mathcal{S}_s = (\mathcal{X}_s, \mathcal{T}_s)$ the surface corresponding to the trachea

Algorithm 5 predict_stent(\mathcal{B}_h'' , \mathcal{A}_h , \mathcal{A}_s)

```

1:  $\mathcal{S} \leftarrow \emptyset$ ,  $\mathcal{S}_T \leftarrow \emptyset$ 
2: for  $k = 1$  to  $|\mathcal{B}_h''|$  do
3:   if ( $A_{sk}/A_{hk} < R_0$ ) then
4:      $\mathcal{S}_T \leftarrow \mathcal{S}_T + \mathcal{B}_{hk}''$ 
5:   else if ( $\text{length}(\mathcal{S}_T) \leq L_0$ ) then
6:      $\mathcal{S}_T \leftarrow \emptyset$ 
7:   else
8:      $\mathcal{S} \leftarrow \mathcal{S} + \mathcal{S}_T$ 
9:      $\mathcal{S}_T \leftarrow \emptyset$ 
10:  end if
11: end for
12: return  $\mathcal{S}$ 

```

with stenosis. Since \mathcal{S}_h preserves the $|\mathcal{X}_h| = n_c \times n_b = n$ structure defined in Section 6.2, its centre line can easily be obtained as the n_c centres of gravity of the contours \mathcal{B}_{hk} with n_b points. Let \mathcal{A}_h be the area profile of \mathcal{S}_h , computed as the areas of the contours \mathcal{B}_{hk} .

The surface \mathcal{S}_s is intersected with the planes containing the contours \mathcal{B}_{hk} , generating a set of contours \mathcal{B}_{sk} . In this way, both the healthy and narrowed versions of the trachea are now represented by a set of contours perpendicular to their center lines. Let now \mathcal{A}_s be the area profile of \mathcal{S}_s , analogously to \mathcal{S}_h . \mathcal{A}_s will show a decrease in area at some points, and a comparison between \mathcal{A}_h and \mathcal{A}_s determines the point where the stenosis starts, its length, and the degree of narrowing.

With the area profiles along the centre lines of both \mathcal{X}_h and \mathcal{X}_s , a very simple algorithm can be used to quantify the stenosis of a patient, as follows:

Algorithm 5 receives as input the set of contours of the estimated healthy trachea, \mathcal{B}_h , along with its area profile, \mathcal{A}_h , and the area profile of the contours of the stenotic trachea, \mathcal{A}_s . Let \mathcal{S}_T be the set that stores all the contours of the estimated healthy shape covering regions of stenosis along the extension of the narrowed trachea. The process starts with the initialization of \mathcal{S}_T and of a temporary variable, \mathcal{S}_0 , which, as \mathcal{S}_T , is a set

of contours. Next, for all contours of the estimated healthy trachea, if the ratio between the areas of healthy and stenotic surfaces is below a given threshold, R_0 , \mathcal{S}_0 is incremented with the corresponding contour from \mathcal{B}_h . If not, and if the length of \mathcal{S}_0 is below another given threshold, L_0 , the temporary variable is emptied, as shown in Step 6. This test avoids the inclusion of very short segments in the set. If \mathcal{S}_0 is long enough, it is finally added to \mathcal{S}_T , in Step 8. In the end, the algorithm returns as a result the contours contained in \mathcal{S}_T .

From the resulting set \mathcal{S} , the parameters of the stenosis are calculated as follows:

1. start point: taken as the centre point of the contour \mathcal{B}_{sk}'' that precedes the first contour of \mathcal{S}_T ;
2. end point: taken as the centre point of the contour \mathcal{B}_{sk}'' that succeeds the last contour of \mathcal{S}_T ;
3. degree of narrowing: taken as the smallest ratio computed between the area profiles of \mathcal{B}_h'' and \mathcal{B}_s'' .

8.3 Prediction of Stents

The stent parameters can also be obtained from set \mathcal{S}_T returned by Algorithm 5. The calibre of the stent is the average diameter of all contours in \mathcal{S}_T . The deployment location is equal to the start point of the stenosis. Finally, the length of the stent is the path length between the start and end points.

By adjusting the thresholds R_0 for the severity and L_0 for the extent of the stenosis, the number and dimensions of the predicted stents can be further adjusted. The applied changes can immediately be evaluated through a visual interaction between the operator and the application, which aids in deciding if the stent is the ideal treatment and in the surgery planning in general.

At this point, all the steps necessary for the assessment and stenting of tracheal stenosis using ASMs and ACMs have been described. Figure 8.1

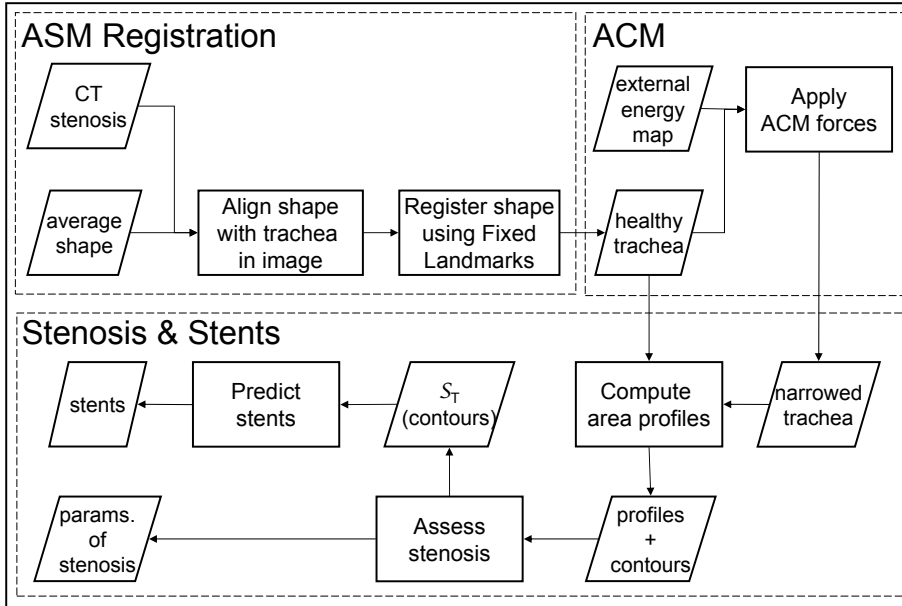


Figure 8.1: Flowchart of the proposed method.

shows a flowchart which illustrates the whole method. The next section proceeds with the experiments with simulation and clinical data in order to validate the proposed method.

8.4 Experiments

Experiments on simulation as well as on clinical data were carried out in order to evaluate the proposed method. With the simulation data, ground truths of measurements of stenoses were formally established in order to quantitatively evaluate the method. The experiments with clinical data were used for qualitative and quantitative evaluations.

To build the ASM, $N = 38$ healthy tracheas were used, each with $n = 820$ landmarks, following the steps described in sections 6.2, 6.3,

and 6.4. The chest CT scans of their respective patients were obtained from pulmonary medication studies carried out at the University Hospital of Antwerp, Belgium, and from the Lung Image Database Consortium (LIDC). Patient data were anonymized before the images were used. The tracheas were segmented from the images and converted to 3D surfaces with the procedure described in Section 6.2 before the model was built.

8.4.1 Simulation Data

The simulation experiments followed the same principle of those in Section 6.6. That is, the same number of phantoms of stenosis, 72, was generated for each trachea, with the narrowing located anteriorly, posteriorly, or symmetrically, following the same classification tables, which are reproduced here in Table 8.1 and Table 8.2. The phantoms, however, were created from $N' = 10$ tracheas randomly taken from the training set of $N = 38$ tracheas above.

Analogously to Section 6.6, a set of $T = N'$ leave-one-out tests was executed. The model was built with $N - 1$ tracheas and was then registered to the phantoms created from the trachea not present in the training set. The quality of the assessment of stenosis and prediction of stents using the ASM strongly depends on the estimation of the healthy trachea of the patient. Therefore, the objective of this set of experiments was to register the model to the phantoms using the proposed *Fixed Landmarks* method (Section 6.5), segment the stenotic phantoms (Section 7.3), quantify the stenosis (Section 8.2), and to measure the difference between the detected parameters of the stenosis and the ground truth parameters. The ground truth parameters were also obtained with the methods of sections 7.3 and 8.2, but using the original healthy trachea as the reference.

The same set of experiments was run with the *Talwar WLS*, one of the best performing Robust ASM Fitting methods discussed in Section 6.5.3 and Section 6.6.1. The purpose of running these experiments was to confirm that in order to achieve good results in the assessment of the stenoses, an ASM registration method that avoids the narrowed regions of target trachea must be employed.

Table 8.1: Categories of stenosis based on location and length.

CATEGORY	LOCATION AND LENGTH
I	Upper third of the trachea
II	Middle third of the trachea
III	Lower third of the trachea
I-II	Upper third extending to middle third
II-III	Middle third extending to lower third
I-III	Upper third extending to lower third

Table 8.2: Categories of stenosis based on its degree of narrowing.

CATEGORY	DEGREE
1	<25%
2	26–50%
3	51–75%
4	>75%

The iteration limit for the registration of the ASM was set to 200, high enough to enable the registrations to find a resting point, in which few surface changes would be observed between iterations. The minimum squared error between shapes generated at consecutive iterations, i.e., $\xi_r^{(k)} = (\mathbf{x}^{\hat{k}} - \mathbf{x}^{\hat{k-1}})^T (\mathbf{x}^{\hat{k}} - \mathbf{x}^{\hat{k-1}})$, was set to 10^{-5}mm^2 . This parameter is somewhat coupled with the number of iterations and was small enough to interrupt the registration when there was no significant change in the shapes of consecutive iterations. Regarding the *Fixed Landmarks*, the landmarks were only displaced if the corresponding high gradients were found within a distance $d = 1 \text{mm}$ along their normals.

For the ACM, the values used as weights for the forces where $\kappa = 0.7$ and $\gamma = 0.5$. These are traditionally difficult parameters to define and usually depend on the quality of the image and on the application itself. Here, the purpose was to give some freedom to the external force, since the

local deformations of the stenoses needed to be correctly modelled. The iteration limit was set to 200 and $\xi_r^{(k)} = 10^{-3}\text{mm}^2$.

Finally, for the quantification of the stenosis, Algorithm 5 used $R_0 = 0.8$, which balances the false positive and detection rates of the start, end, and narrowest points of the stenoses. The threshold L_0 was equivalent to the distance covered by two consecutive centre line points.

8.4.2 Clinical Data

In addition to the simulation experiments, a retrospective study with 11 CT scans, from 9 patients, was executed. The use of the CT scans was approved by the ethics committee of the Ghent University Hospital (doc. ECUZG2009/140), Belgium, and patient data were anonymized before the images were used in the experiments. The patients had stenosis with the following characteristics:

Patient 1 – Severe symmetrical stenosis along the upper third of the trachea,

Patient 2 – Side-to-side stenosis along the lower two thirds of the trachea,

Patient 3 – Severe symmetrical stenosis along the lower two thirds of the trachea,

Patient 4 – Severe posterior stenosis along the lower two thirds of the trachea (2 scans),

Patient 5 – Suspected, short anterior stenosis at the lower half of the trachea,

Patient 6 – Severe long segment posterior stenosis,

Patient 7 – Severe posterior stenosis along the upper two thirds of the trachea,

Patient 8 – Short anterior stenosis above the carina (2 scans),

Patient 9 – Symmetrical stenosis along the upper third of the trachea.

For the registration of the ASM, no preprocessing was applied to the CT images, even though some of them were very anisotropic in the axial direction. The original images were converted to a distance map through the steps described in Section 7.3 and the stenotic tracheas were segmented with the ACM. Finally, the quantification of the stenosis was performed using the procedure described in Section 8.2. The results were reviewed by experts in the pulmonology field in order to evaluate the proposed method with respect to the estimated healthy trachea, the assessment of the stenosis, and the prediction of stents.

The ASM built for these experiments contained all the $N = 38$ healthy tracheas. All other parameters were set as in the simulation experiments.

8.5 Results

8.5.1 Simulation Data

Figure 8.2 shows the results obtained from the experiments with simulation data for the assessment of stenosis. The graphs show area profiles of healthy and stenotic tracheas for both the estimation obtained with the ASM registration and the ground truth obtained with the originally healthy tracheas used to build the phantoms. Each row of Figure 8.2 shows three categories of stenosis for a given trachea, located posteriorly, anteriorly, and symmetrically, respectively. The reader should note that in most cases the parameters of the stenosis were very close to the ground truth parameters. Yet, there were a few difficult situations.

In Figure 8.2(c), for instance, it was difficult to detect the start of the stenosis and the point of maximum degree. In the case of the start point, the reason for the error was that the calibre of the estimated trachea was slightly overestimated at the top, while in the point of maximum degree there was an underestimation of the calibre along the rest of the centre line. It is worth mentioning, however, that this is the most difficult simulation case, i.e., a long segment stenosis, in which few healthy areas are available

in the target trachea to guide the registration of the ASM. Still, the method performed very well in other similar situations.

A different situation occurred in Figure 8.2(a). The absence of healthy posterior tracheal wall along the upper two thirds of the phantom misled the registration of the ASM, which eventually yielded an overestimation of the calibre of the trachea along its lower part. The result was that the estimated extension of the stenosis was much longer than the ground truth measurement, due to the faulty detection of the end point.

Figure 8.3 presents statistics of the error of the estimated parameters of the stenoses relative to the ground truth for each of the $N' = 10$ tracheas used in the simulations. Note that the average error in degree between estimation and ground truth was around only 0.05 percentage points, while the error in the other parameters was around 5mm. At first, these results can be considered less than ideal. However, physicians can only choose between stents differing in length per centimetre, thus an average error of 5mm is in reality acceptable. Moreover, Figure 8.3 reveals that the average and the standard deviation were severely affected by the presence of large outliers. In fact, most measurements within the 75 percentile were below the average, and the median of the errors in nearly all cases was less than 1.5mm.

The presence of outliers is explained by the large errors that can occur if the estimated healthy calibre is not adequate. The most difficult cases were those in which the stenosis had very low degree (a situation where stents are normally not employed) or extended through long segments. This can be evidenced in Figure 8.4, which shows how the error varied per simulated category of stenosis. Indeed, categories 2 and I-III had the largest errors.

Figure 8.5 shows the results obtained from the simulation experiments that used the *Talwar WLS* method instead of the *Fixed Landmarks*. The aim here is to confirm that the correctness in the assessment of the stenosis strongly depends on the quality of the estimation of the healthy trachea of the patient. Although the results for the start and end points of the stenosis did not differ too much from those obtained with the *Fixed Landmarks*, the degree of narrowing, as expected, was considerably underestimated. The net result was that the errors in the assessment of the degree with respect




Figure 8.2: Evaluation of the proposed method with simulation data. The graphs show the cross-sectional area profiles and the parameters of stenosis for the estimated and the ground truth versions of the phantom tracheas.

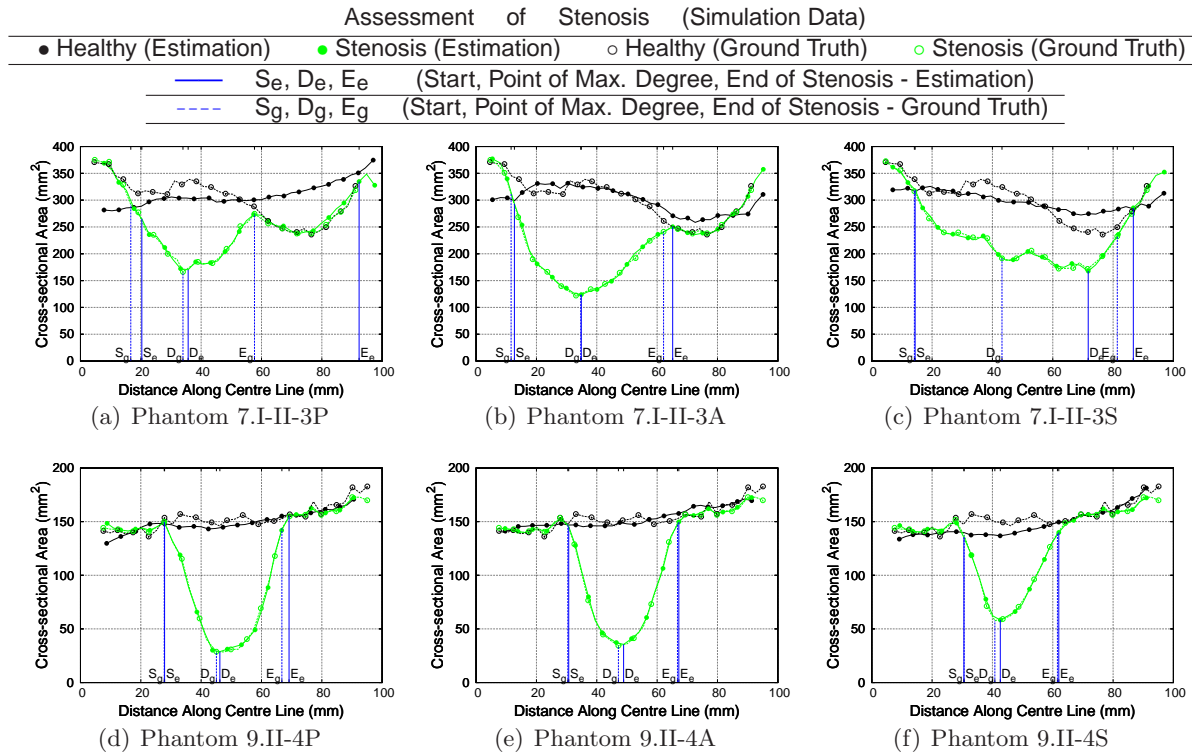


Figure 8.2: (Cont.) Evaluation of the proposed method with simulation data. The graphs show the cross-sectional area profiles and the parameters of stenosis for the estimated and the ground truth versions of the phantom tracheas.

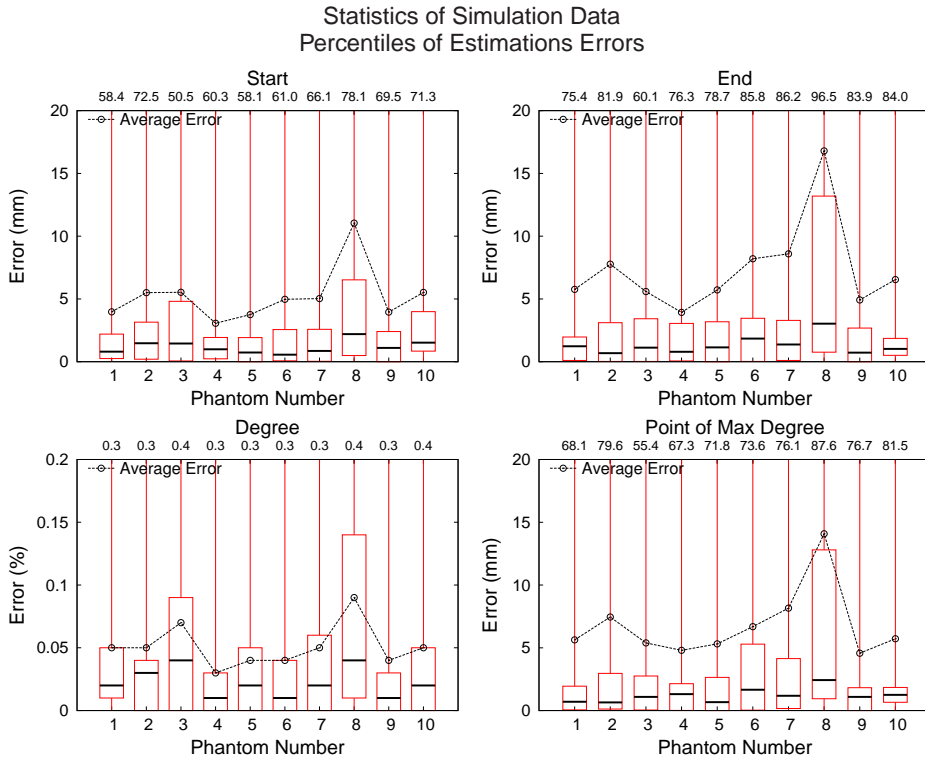


Figure 8.3: The graphs show, for each phantom, the minimum, 1st quartile, median, 3rd quartile, maximum, and average errors of the estimated parameters of the 72 simulated categories of stenosis relative to the ground truth (the maximum values are shown at the top of the graph).

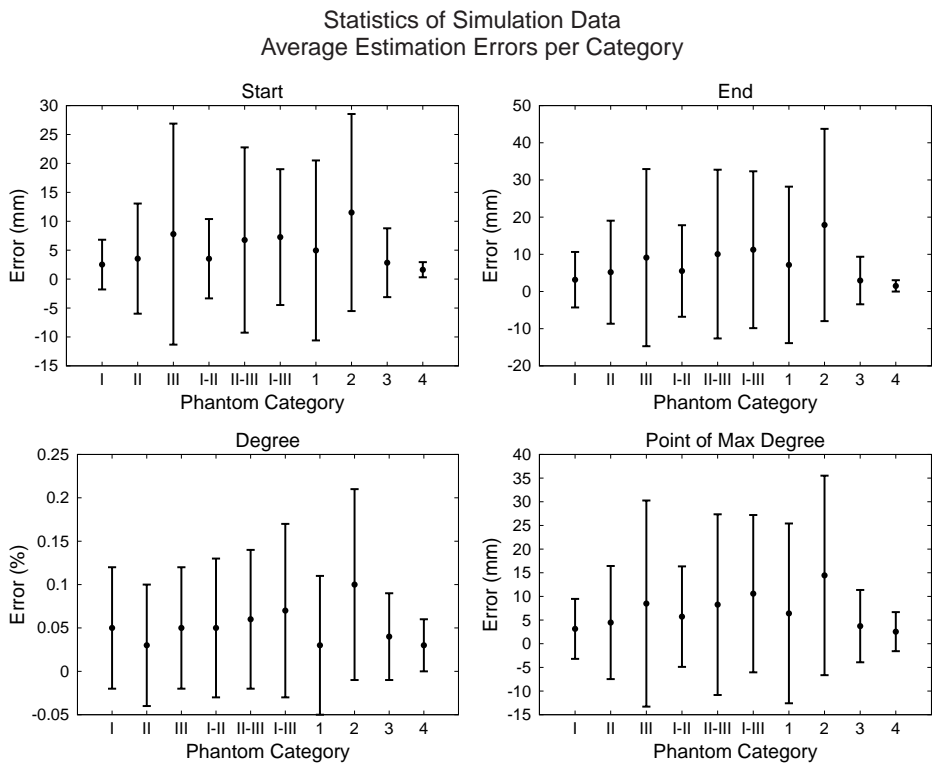


Figure 8.4: Average estimation error per simulated category of stenosis, along with their respective standard deviations.

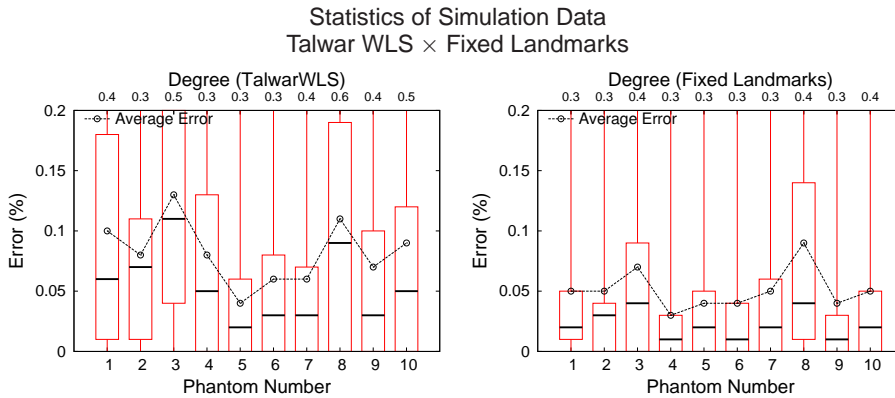


Figure 8.5: The graphs show, for each phantom, the percentiles and average errors of the degree of narrowing of the 72 simulated categories of stenosis relative to the ground truth when using the *Talwar WLS* and the *Fixed Landmarks* methods.

to the ground truth were larger than when using the *Fixed Landmarks* method, as the figure demonstrates.

8.5.2 Clinical Data

For the evaluation with clinical data, both objective and subjective evaluations were planned. For the objective evaluations, a special application was developed to allow the physicians to determine the parameters of stenosis based on the patient's CT image, in a way similar to what they do in their clinical practice. In this application, the physician visualizes the multiplanar reconstructions of the patients' CT scans and can navigate through the slices along each plane (see Figure 8.6). Zooming and panning was provided in order to ease the visualization of the images. The objective was to manually select in the images, with a simple mouse click, the points where the stenoses started, where they ended, where they reached their maximum degree, and to select in the application's GUI the level of severity according to the values presented in Table 8.2.

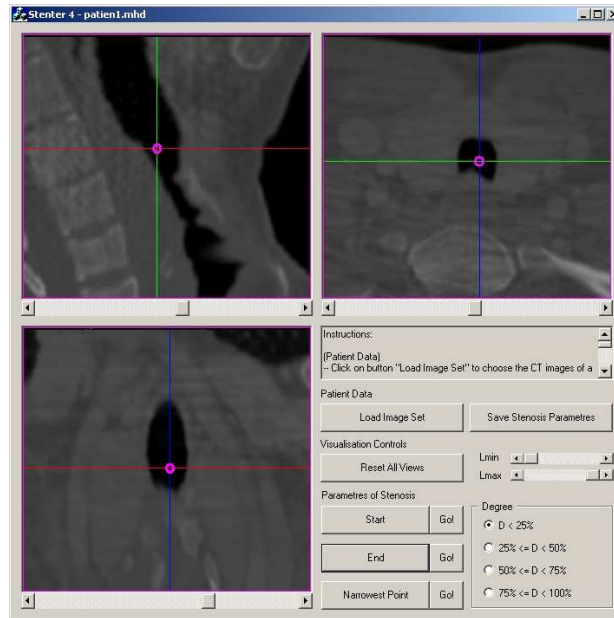


Figure 8.6: GUI of the application developed for manual/visual assessment of stenosis. The expert marks on the patient's CT image the parameters of the stenosis (start, end, point of maximum degree) and chooses the degree of narrowing from ranges of values. Zooming and panning enhances the visualization.

The purpose of creating the application above was to cross-check the results obtained with the proposed automatic assessment of stenosis and the results obtained from experts' visual/manual assessments (the gold standard). The application was thus run by 3 experts in the pulmonology field and the results of their visual/manual assessments were recorded. Table 8.3 and Table 8.4 show a significant variability present in the recorded results for the stenoses of the 9 patients. With such variability, it is difficult to make an objective validation of the proposed automatic method against the manual evaluations. As a consequence, these data could not be used.

Table 8.3: Samples of the manual/visual measurements performed by 3 experts in the pulmonology field. All values are in millimetres and indicate the position along the axial axis where the corresponding parameters of the stenosis were marked by the experts, relative to the start of the patients' CT scan.

	Patient 1			Patient 3		
	Exp.1	Exp.2	Exp.3	Exp.1	Exp.2	Exp.3
Start	93	93	94	540	576	495
Pt. Max Degree	107	95	95	594	657	576
End	110	113	112	684	675	576
	Patient 6			Patient 7		
	Exp.1	Exp.2	Exp.3	Exp.1	Exp.2	Exp.3
Start	151	170	145	55	64	73
Pt. Max Degree	189	185	174	87	81	86
End	198	194	187	99	96	91

Instead, they proved that stenosis assessment is indeed very operator dependent.

Concerning the results obtained with the proposed method, healthy tracheas of plausible calibre and shape were yielded in nearly all cases, as can be seen in Figure 8.7. In addition, the automatically detected parameters of the stenoses conformed to the diagnosed parameters, as described in Section 8.4.2. In the subjective evaluation, the consulted experts agreed that the proposed automatic assessment of the stenosis and the suggested stents were accurate enough to support their diagnostic decisions. The proposed method therefore serves its purpose and sheds light in the diagnostic pro-

Table 8.4: Variability in experts' manual/visual assessment of stenosis, for the 9 patients used in the experiments. The table shows the standard deviation (in millimetres) observed when the experts manually determined the parameters of the stenoses.

	P1	P2	P3	P4.1	P4.2	P5
Start	0.58	5.00	13.53	19.97	18.25	2.89
Pt. Max Degree	6.93	5.77	14.18	1.73	6.93	2.89
End	1.53	21.79	19.97	4.58	16.52	11.55
	P6	P7	P8.1	P8.2	P9	
Start	13.05	9.00	0.00	0.00	9.81	
Pt. Max Degree	7.77	3.21	0.00	0.00	3.46	
End	5.57	4.04	2.89	2.89	5.20	

cess, by providing a common ground from which specialists can do their own assessment.

The good quality of the results does not mean, however, that difficult cases did not exist. The first difficult case was **Patient 3**, in Figure 8.7(c). The very unusual geometry of the patient's trachea hindered the registration of the ASM, which could not produce a plausible healthy shape along the narrowed areas. Eventually, the stenosis, notably its degree, was largely underestimated. The second difficult case was **Patient 5**, shown in Figure 8.7(f). Since the stenosis had very low degree (even the official diagnostic mentioned only a suspicion that stenosis was present), the *Fixed Landmarks* did not succeed in differentiating between healthy and narrowed regions of the trachea. As a consequence, the narrowed regions had stronger influence on the registration of the ASM, and the method could not detect any narrowing at all.

Figure 8.8 shows the predicted stents for some of the patients used in our experiments with clinical data. The stents are directly derived from the computed parameters of the stenosis, as described in Section 8.3. Note that they cover the stenosis according to the parameter R_0 defined at 0.8 (at least 20% of narrowing) and provide a plausible healthy calibre for the

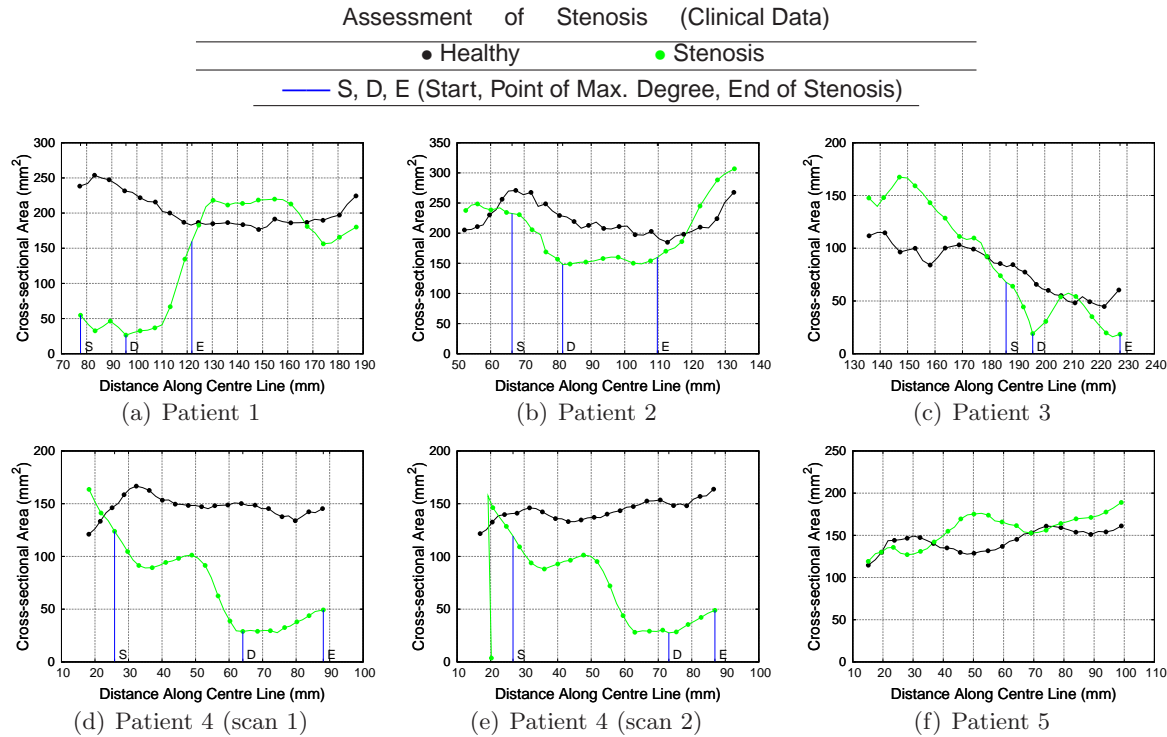


Figure 8.7: Stenosis quantification for the experiments with clinical data. The graphs show the cross-sectional area profiles and the parameters of stenosis for the narrowed and estimated healthy versions of the patients' tracheas.

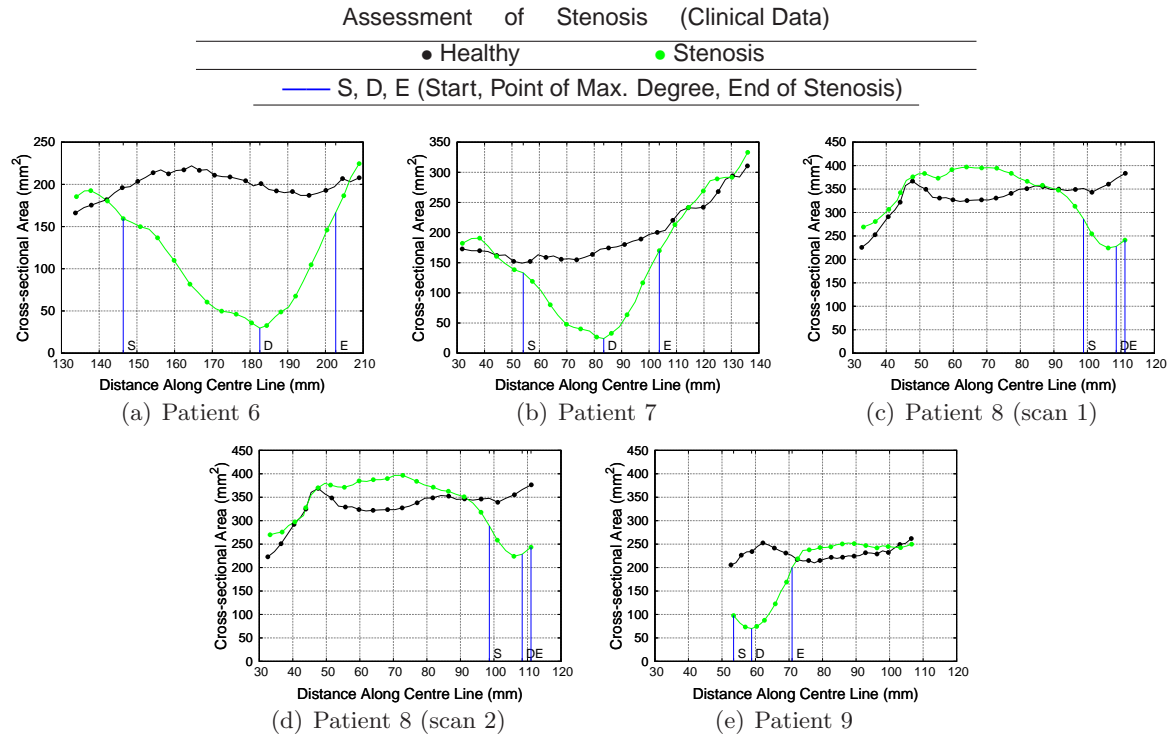


Figure 8.7: (Cont.) Stenosis quantification for the experiments with clinical data. The graphs show the cross-sectional area profiles and the parameters of stenosis for the narrowed and the estimated healthy versions of the patients' tracheas.

narrowed trachea. Nonetheless, an exception can be observed in **Patient 3**. As shown in Figure 8.7(c), the assessment of the stenosis in this case was faulty due to the highly deformed nature of the patient's trachea, which is evidenced in Figure 8.8(c).

8.6 Discussion

8.6.1 Fixed Landmarks

Despite the good results obtained with the *Fixed Landmarks*, it is important to mention that the quality of the proposed method depends on the choice of the sets \mathcal{L}' and \mathcal{L}'' , which correspond to landmarks that are allowed to move and those that remain fixed at each iteration, respectively. These sets, in turn, depend on the choice of the threshold distance d . Intuitively, as d increases, the *Fixed Landmarks* tends to perform like the *StandardLS*, since \mathcal{L}'' will tend to be empty and no landmarks will remain fixed. The registration will thus not be guarded against the attraction of stenotic regions. Conversely, if d is too short, \mathcal{L}' will tend to be empty and all landmarks will remain fixed. In the worst case, the latter situation will cause the resulting shape to be equal to the average shape of the ASM. Figure 8.9 demonstrates this behaviour after the whole set of simulation experiments were performed with different values of d . The graphs show how the error in the estimation of the degree of stenosis, the most sensitive parameter, varied with d .

8.6.2 Execution Times

Although a detailed time analysis was out of the scope of this thesis, a quick evaluation showed that each step of the overall method takes only a few seconds to run with the provided parameters. The machine used in the tests was an Intel[®] Core[™] 2 Quad CPU, at 2.4 GHz, with 8GB of RAM, running under Windows Vista[™] Ultimate 64-bits, and all algorithms were implemented in C++. For one case, the measured times are given in Table




Figure 8.8: Visualization of the stents computed for the patients in the experiments with clinical data.




Figure 8.8: (cont.) Visualization of the stents computed for the patients in the experiments with clinical data.

8.5. These results mean that the method is a viable tool to be used in the clinical setting.

Table 8.5: Approximate execution times of each step of the proposed method, measured for one example case.

EXECUTION TIMES		
ASM	ACM	Stenosis + Stents
~30s	~60s	~1s

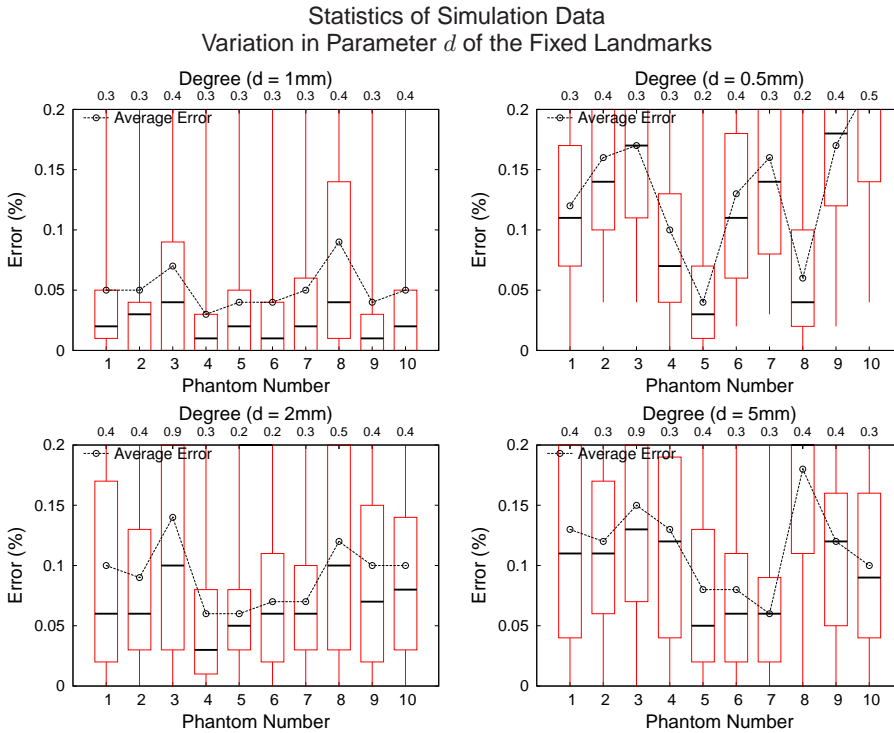


Figure 8.9: Variation in the results obtained from the simulation experiments when varying the parameter d of the *Fixed Landmarks* method. The graphs show how the error in the computation of the degree of narrowing of the stenosis increased when $d \neq 1\text{mm}$.

Part III

Conclusions

Chapter 9

Conclusion

“Kids, you tried your best and you failed miserably.
The lesson is, never try.”
Homer Simpson

“I have not failed. I’ve just found
10,000 ways that won’t work.”
Thomas A. Edison (1847 – 1931)

This thesis presented a decision support system for the assessment and stenting of tracheal stenosis. Although bronchoscopy remains the gold standard method for the assessment of the strictures, it is invasive, render subjective results, and may even fail if the stenosis is too narrow. Current methods that benefit from advances in the imaging field and integrate sophisticated algorithms to aid the physician in the assessment are still very operator dependent.

In clinical practice, physicians visually guess the shape of the healthy trachea of a patient, i.e., without stenosis, in order to determine the parameters of the stricture (start, end, degree of narrowing) and make the choice of treatment. The idea behind the proposed system is to provide to the physician this healthy shape. An Active Shape Model (ASM) of healthy tracheas was proposed in order to achieve such results. A defining characteristic of ASMs is that they only generate shapes that resemble those in the training set from which it was built. Since the proposed model contains only instances of healthy tracheas, geometrical variations typical of stenosis can not be generated. Still, when such model is registered to a CT chest scan of a patient with stenosis, the narrowed areas of the trachea can influence the results, making the registration a challenge in itself. A novel ASM registration method called *Fixed Landmarks* was proposed to avoid the influence of these narrowed regions. Experiments showed that the *Fixed Landmarks* outperforms other approaches used previously in different application domains.

The obtained estimation of the healthy trachea is the basis for the automatic assessment of the stenosis and the suggestion of a stent. In addition to the healthy shape, however, it is also necessary to segment the narrowed trachea from the CT images. An adapted Active Contour Model (ACM) was thus proposed for this task. It uses the healthy shape as the initialization and deforms this shape towards the boundary of the stenotic trachea in the image. Special external and internal forces were develop for the ACM so as to guarantee that the shape is attracted to the correct boundaries while maintaining its continuity and smoothness.

The healthy and narrowed versions of the trachea are then employed in the automatic calculation of the parameters of the stenosis and of the

corresponding stent. A comparison between the cross sectional area profile of both surfaces makes the process straightforward. An extensive set of experiments with simulation as well as with clinical data validated the proposed method. It was shown that the obtained results, which were validated by specialists in the pulmonology field, could be used as a common ground to aid physicians in the diagnostic process and to reduce the usually observed operator dependency.

In the development of the methods above, a couple of extra challenges needed to be solved. The first was the manipulation of large medical image files. With current imaging technology, chest CT scans can yield images of a couple of gigabytes. Manipulating such images often arose memory restrictions, because they could not be kept in memory all the time. To this end, a cache and prefetching strategy for out-of-core image processing algorithms was also proposed. This strategy uses a sliding window protocol to prefetch image data from disk, taking advantage of a priori knowledge about the access pattern of the algorithms through the images. Experiments demonstrated that the proposed strategy generates very low miss ratios, while keeping execution times to acceptable levels.

To build the aforementioned ASM of healthy tracheas, these organs must first be segmented from CT images of healthy patients. A segmentation algorithm that extracts the trachea from beginning (right below the cricoid cartilage) to end (at the level of the carina) was thus the second extra challenge to be solved in this thesis. Such an algorithm was in fact part of one proposed to segment the entire intrathoracic airway tree. The proposed algorithm uses a heuristic to identify the beginning of the trachea and compute a seed point to be provided to an adaptive region growing algorithm. In the latter, cylindrical regions of interest (ROIs) bound the segmentation and control the appearance and removal of leaks, a common problem in region growing techniques. Experiments with 40 images validated the tracheal detection algorithm and showed that the region growing algorithm indeed reduces the number of leaks in the segmentation.

Chapter 10

Future Work

“I’m a laboratory mouse in the first stages of an elaborate plan to take over the world.”

Brain, from Pink and Brain

“Science may set limits to knowledge, but should not set limits to imagination.”

Bertrand Russell (1872 – 1970)

It is very rare that any piece of research work can no longer be improved right on the next day after it is considered to be “finished”. This is by no means different with this thesis, and in fact this is one of the marvels of doing research. In the following sections, possible improvements to each part of the proposed method are described and discussed.

10.1 Cache and Prefetching

A natural future development for the cache and prefetching is the validation of the proposed window structures and sliding protocols with other image processing algorithms, such as registration, compression, stitching, editing, etc.

Although the focus of the proposed method was on out-of-core image processing algorithms for medical applications, the sliding window concept of *WindowCache* can be applied to a wide range of applications where memory and latency management are important. For instance, transport of large amounts of data through networks, processor level cache optimizations, visualization algorithms, data transfers between main memory and GPUs and the like.

10.2 Airway Tree Segmentation

Improvements to the proposed region growing algorithm include the use of adaptive intensity thresholds, employed in, e.g., (van Ginneken et al., 2003; Kiraly et al., 2002). In fact, this approach has already been tried and produced some preliminary results. Although the number of branches may increase considerably, so may the number of leaks. As a consequence, this technique must be coupled with more efficient leak detection and removal.

10.3 Fixed Landmarks

One option to reduce the dependency on the threshold distance d discussed in Section 8.6.1 is to let this parameter be set by the operator. Different

values of d should then be tried until acceptable results are yielded. Another possibility is to devise an adaptive algorithm to change d as needed during the registration. For instance, d could start with a larger value and be progressively reduced as a function of the iteration number until it reached a lower threshold. Nevertheless, the value $d = 1\text{mm}$ proved to be a good empirical choice in our experiments with simulation and clinical data.

Most of the errors observed in the estimation of the healthy trachea were related to geometric variations that could not be completely modelled by the ASM. This can be overcome with a more careful selection of shapes for the training set, one which covers the largest set of geometric variations for healthy subjects. In addition, further improvements may also be achieved with more sophisticated methods to establish correspondences between the shapes of the training set.

Although the focus of the present work was on tracheal stenosis, it is reasonable to expect that the *Fixed Landmarks* method could also be applied on vascular stenoses. The difficulties in this case would mostly lie in the fact that arteries have a much larger global shape variability than tracheas and that longitudinal versus radial scale is also much larger in arteries. As a consequence, it could be difficult to build an ASM of healthy arteries that would be completely resistant to the regions where the stenoses occur. Still, arteries can usually be efficiently segmented with the use of contrast agents, which could at least remove the interference of noise and the presence of neighbouring structures.

There are also reasons to believe that the *Fixed Landmarks* method could be used as a *Robust ASM Fitting* approach in general. Since ASM registrations yield only objects that resemble those in the training set, keeping landmarks fixed with respect to their position in the previous iteration of the registration immediately removes the influence of outliers among the residuals. The fixed landmarks, together with those that were allowed to move, eventually help the registration generating the desired shape.

10.4 Functional Stent Analysis

A considerable step forward in the proposed framework would be to extend the stent prediction with a functional analysis of airways through airflow simulations. By doing this, other parameters such as the type of the stent could be added to the model. A numerical model of the tracheal deformation could take into account the responsiveness of the trachea to the predicted stent. All these results could be fed back into the stent prediction so that the choice of stent could be optimized. This type of analysis could also be useful in arterial stenosis, where blood flow simulations could influence the stent choice.

Contributions

Journal papers

- R. Pinho, K. J. Batenburg, and J. Sijbers. Windowcache: Pre-fetching strategies for out-of-core image processing algorithms. *Image Processing, IEEE Transactions on*, (Under Review), 2010a.
- R. Pinho, K. G. Tournoy, R. Gosselin, and J. Sijbers. Assessment and stenting of tracheal stenosis using deformable shape models. *Medical Image Analysis*, (Under Review), 2010b.
- L. Vanacken, R. Pinho, J. Sijbers, C. Raymaekers, and K. Coninx. Force feedback to assist active contour modelling for tracheal stenosis. (In Preparation), 2010.

Book chapters

- R. Pinho, K. G. Tournoy, and J. Sijbers. Computer-aided assessment and stenting of tracheal stenosis. In Ayman El-Baz, editor, *Computer Aided Diagnosis of Lung Imaging*, volume (in press). American Scientific Press, 2011.

Conference proceedings (full paper)

- R. Pinho, J. Sijbers, and W. Vos. Efficient approaches to intrathoracic airway tree segmentations. In *Proc. of the Biomedical Engineering IEEE/EMBS Benelux Symposium*, volume 2, pages 151–154, Brussels, Belgium, 2006.
- R. Pinho, J. Sijbers, and T. Huysmans. Segmentation of the human trachea using deformable statistical models of tubular shapes. In *Proc. of Advanced Concepts for Intelligent Vision Systems*, volume 4678 of *Lecture Notes in Computer Science*, pages 531–542, 2007.
- R. Pinho, K. J. Batenburg, and J. Sijbers. Seeing through the window: Pre-fetching strategies for out-of-core image processing algorithms. In *Proceedings of SPIE Medical Imaging*, volume 6919, page 69190D, San Diego, CA, USA, 2008a. SPIE.
- R. Pinho, T. Huysmans, W. Vos, and J. Sijbers. Tracheal stent prediction using statistical deformable models of tubular shapes. In *Proc. of SPIE Medical Imaging*, page 69144O, San Diego, CA, USA, 2008b.
- R. Pinho, S. Luyckx, and J. Sijbers. Robust region growing based intrathoracic airway tree segmentation. In *2nd International Workshop on Pulmonary Image Analysis*, pages 261–271, London, England, 2009a.
- R. Pinho, K. G. Tournoy, R. Gosselin, and J. Sijbers. Assessment of tracheal stenosis using active shape models of healthy tracheas: A surface registration study. In *2nd International Workshop on Pulmonary Image Analysis*, pages 125–136, London, England, 2009b.
- R. Pinho, K. G. Tournoy, and J. Sijbers. A decision support system for the treatment of tracheal stenosis. In *Proc. of Workshop on Discrete Geometry and Mathematical Morphology (WADGMM)*, pages 72–76, Istanbul, Turkey, 2010.

Conference proceedings (abstract)

- A. Bernat, T. Huysmans, F. van Glabbeek, J. Sijbers, R. Pinho, and J. L. Gielen. Exploring the clavicle: Morphometric differences using a 3d model. In *AAOS Annual Meeting*, Las Vegas, Nevada, 2009.
- T. Huysmans, A. Bernat, R. Pinho, J. Sijbers, F. van Glabbeek, P. M. Parizel, and H. Bortier. A framework for morphometric analysis of long bones: Application to the human clavicle. In *Liege Image Days 2008: Medical Imaging*, 2008.
- R. Pinho, T. Huysmans, W. Vos, and J. Sijbers. Tracheal stent prediction using statistical deformable models of healthy tracheas. In *Liege Image Days 2008: Medical Imaging*, 2008.

Bibliography

- M. A. Butt and P. Maragos. Optimum design of chamfer distance transforms. *Image Processing, IEEE Transactions on*, 7(10):1477–1484, 1998.
- A. Agrawal, J. Kohout, G. J. Clapworthy, N. J. Mcfarlane, F. Dong, M. Viceconti, F. Taddei, and D. Testi. Enabling the interactive display of large medical volume datasets by multiresolution bricking. *Journal of Supercomputing*, 51(1):3–19, 2010.
- N. Aspert, D. Santa-Cruz, and T. Ebrahimi. Mesh: Measuring errors between surfaces using the hausdorff distance. In *Proc. of the IEEE International Conference on Multimedia and Expo*, volume I, pages 705–708, 2002.
- M. A. G. Ballester, A. P. del Palomar, J. L. L. Villalobos, L. L. Rodríguez, O. Trabelsi, F. Pérez, Á. G. Ca namaque, E. B. Cortés, F. R. Panadero, M. D. Castellano, and J. H. Jover. Surgical planning and patient-specific biomechanical simulation for tracheal endoprotheses interventions. In *Medical Image Computing and Computer-Assisted Intervention – MICCAI 2009*, volume 5762 of *Lecture Notes in Computer Science*, pages 275–282. Springer, 2009.
- H. D. Becker and B. R. Marsh. History of the rigid bronchoscope. In

- C. T. Bolliger and P. N. Mathur, editors, *Interventional Bronchoscopy*, volume 30, pages 2–15. Karger, 2000.
- P. M. Boiselle, J. Catena, A. Ernst, and D. A. Lynch. Tracheobronchial stenoses. In Phillip M. Boiselle and David Lynch, editors, *CT of the Airways*, pages 121–149. Humana Press - Springer, 2008.
- G. Borgefors. Distance transformations in digital images. *Computer Vision Graphics and Image Processing*, 34(3):344–371, 1986.
- M. Botsch, S. Steinberg, S. Bischoff, and L. Kobbelt. Openmesh – a generic and efficient polygon mesh data structure. In *Proc. of the OpenSG Symposium*, Bonn, Germany, 2002.
- S. Bruckner. Efficient volume visualization of large medical datasets. Master’s thesis, Institute of Computer Graphics and Algorithms, Vienna University of Technology, Favoritenstrasse 9-11/186, A-1040 Vienna, Austria, 2004.
- R. Bruckschen, F. Kuester, B. Hamann, and K. I. Joy. Real-time out-of-core visualization of particle traces. In *PVG '01: Proceedings of the IEEE 2001 symposium on parallel and large-data visualization and graphics*, pages 45–50, 2001. ISBN 0-7803-7223-9.
- V. Callanan, K. Gillmore, S. Field, and A. Beaumont. The use of magnetic resonance imaging to assess tracheal stenosis following percutaneous dilatational tracheostomy. *The Journal of Laryngology and Otology*, 111(10):953–957, 1997.
- E. Caron, O. Cozette, D. Lazure, and G. Utard. Virtual memory management in data parallel applications. In *LNCS: Proceedings of HPCN Europe*, volume 1593, pages 1107–1116, 1999.
- A. Carretta, G. Melloni, P. Ciriaco, L. Libretti, M. Casiraghi, A. Bandiera, and P. Zannini. Preoperative assessment in patients with postintubation tracheal stenosis. *Surgical Endoscopy*, 20(6):905–908, 2006.

-
- C. S. Chin, V. Litle, J. Yun, T. Weiser, and S. J. Swanson. Airway Stents. *Annals of Thoracic Surgery*, 85(2):S792–796, 2008.
- Wu Ching-Yang, L. Yun-Hen, H. Ming-Ju, W. Yi-Chen, L. Ming-Shian, K. Po-Jen, and L. Hui-Ping. Airway stents in management of tracheal stenosis: have we improved? *ANZ Journal of Surgery*, 77:27–32(6), 2007.
- R. Chiplunkar, J. M. Reinhardt, and E. A. Hoffman. Segmentation and quantitation of the primary human airway tree. In *Proc. SPIE Medical Imaging*, pages 403–414, 1997.
- K. Chmura, S. Hines, and E. D. Chan. Airway anatomy and physiology. In Phillip M. Boiselle and David Lynch, editors, *CT of the Airways*, pages 3–24. Humana Press - Springer, 2008.
- L. D. Cohen. On active contour models and balloons. *Computer Vision, Graphics, and Image Processing. Image Understanding*, 53(2):211–218, 1991.
- L.D. Cohen and I. Cohen. Finite-element methods for active contour models and balloons for 2-d and 3-d images. *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, 15(11):1131–1147, 1993.
- T. F. Cootes, C. J. Taylor, D. H. Cooper, and J. Graham. Active shape models: their training and application. *Computer Vision and Image Understanding*, 61(1):38–59, 1995.
- R. Cucchiara and M. Piccardi. Exploiting image processing locality in cache pre-fetching. In *Proceedings of 5th Int. Conf. on High Performance Computing HiPC'98*, pages pp. 466–472, 1998.
- R. Cucchiara, M. Piccardi, and A. Prati. Neighbor cache prefetching for multimedia image and video processing. *Multimedia, IEEE Transactions on*, 6(4):539–552, 2004.
- P. Czaja, J. Soja, P. Grzanka, A. miel, A. Szczeklik, and K. Sladek. Assessment of airway caliber in quantitative videobronchoscopy. *Respiration*, (74):432–438, 2007.

- C. Doooms, T. De Keukeleire, A. Janssens, and K. Carron. Performance of fully covered self-expanding metallic stents in benign airway strictures. *Respiration*, 77:420–426, 2009.
- N. D. D’Souza, J. M. Reinhardt, and E. A. Hoffman. Asap: interactive quantification of 2d airway geometry. In *Medical Imaging 1996: Physiology and Function from Multidimensional Images*, volume 2709, pages 180–196. SPIE, 1996.
- M. Elliott, D. Roebuck, C. Noctor, C. McLaren, B. Hartley, Q. Mok, C. Dunne, N. Pigott, C. Patel, A. Patel, and C. Wallis. The management of congenital tracheal stenosis. *International Congress Series*, 1254:321 – 334, 2003. Advances in Pediatric ORL. Proc. of the 8th International Congress of Pediatric Otorhinolaryngology.
- L. F. Valencia, J. Montagnat, and M. Orkisz. 3D graphical models for vascular-stent pose simulation. In *Proc. of the International Conference on Computer Vision and Graphics*, pages 259–266, Zakopane, Poland, 2002.
- R. Farias and C. T. Silva. Out-of-core rendering of large, unstructured grids. *IEEE Computer Graphics and Applications*, 21(4):42–50, 2001.
- J. S. Ferguson and G. McLennan. Virtual bronchoscopy. *Proceedings of the American Thoracic Society*, 2(6):488–491, 2005.
- L. Freitag. Tracheobronchial stents. In Bolliger C. T. and Mathur P. N., editors, *Interventional Bronchoscopy*, volume 30, pages 171–186. Karger, 2000.
- L. Freitag. Airway stents. In *Interventional Pulmonology*, volume 48, pages 190–217. ERS, 2010.
- L. Freitag, A. Ernst, M. Unger, K. Kovitz, and C. H. Marquette. A proposed classification system of central airway stenosis. *European Respiratory Journal*, (30):7–12, 2007.

-
- R. C. Gonzalez and R. E. Woods. *Digital Image Processing*. Addison-Wesley Longman Publishing Co., Inc., 2001.
- S. M. Graham, G. McLennan, G. F. Funk, H. T. Hoffman, T. M. McCulloch, J. Cook-Granroth, and E. A. Hoffman. Preoperative assessment of obstruction with computed tomography image analysis. *American Journal of Otolaryngology*, 21(4):263–270, 2000.
- H. C. Grillo. *Surgery of the Trachea and Bronchi*. BC Decker, 2004.
- H. C. Grillo, D. M. Donahue, D. J. Mathisen, J. C. Wain, and C. D. Wright. Postintubation tracheal stenosis: Treatment and results. *The Journal of Thoracic and Cardiovascular Surgery*, 109(3):486 – 493, 1995.
- C. J. Hegedus. Detailed mathematical description of the geometry of airway bifurcations. *Respiratory Physiology & Neurobiology*, 141:99–114.
- J. L. Hennessy and D. A. Patterson. *Computer Architecture: A Quantitative Approach*. Morgan Kaufman Publishers, Inc., 1990.
- I. Hertel, G. Strauss, T. Schulz, J. Dornheim, J. Cordes, A. Krger, C. Tietjen, O. Burgert, A. Dietz, and B. Preim. Three-dimensional visualisation of ct-data for surgical planning in trachea resection: proof of concept. In *International Journal of Computer Assisted Radiology and Surgery*, volume 1, pages 312–313. Springer, 2006.
- E. A. Hoffman, D. Gnanaprakasam, K. B. Gupta, J. D. Hoford, S. D. Kugelmass, and Richard S. Kulawiec. Vida: an environment for multi-dimensional image display and analysis. In *Biomedical Image Processing and Three-Dimensional Microscopy*, volume 1660, pages 694–711. SPIE, 1992.
- H. Hoppe, H. Dinkel, B. Walder, G. von Allmen, M. Gugger, and P. Vock. Grading Airway Stenosis Down to the Segmental Level Using Virtual Bronchoscopy. *Chest*, 125(2):704–711, 2004.
- P. J. Huber. *Robust Statistics*. Wiley, 1981.

- T. Huysmans, J. Sijbers, and B. Verdonk. Automatic construction of correspondences for tubular surfaces. *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, 32(4):636–651, 2010.
- S Ikeda. *Atlas of flexible bronchofiberoscopy*. University Park Press, 1974.
- M. Kass, A. Witkin, and D. Terzopoulos. Snakes: Active contour models. *International Journal of Computer Vision*, V1(4):321–331, 1988.
- K. Kiesler, M. Gugatschka, E. Sorantin, and G. Friedrich. Laryngo-tracheal profile: a new method for assessing laryngo-tracheal stenoses. *European Archives of Oto-Rhino-Laryngology*, 264(3):251–256, 2007.
- A. P. Kiraly, W. E. Higgins, G. McLennan, E. A. Hoffman, and J. M. Reinhardt. Three-dimensional human airway segmentation methods for clinical virtual bronchoscopy. *Academic Radiology*, 9(10):1153 – 1168, 2002.
- P. Kohlmann, S. Bruckner, A. Kanitsar, and M. E. Gröller. Evaluation of a bricked volume layout for a medical workstation based on java. *Journal of the Winter School of Computer Graphics*, 15(1-3):83–90, 2007.
- T. Law and P. Heng. Automated extraction of bronchus from 3D CT images of lung based on genetic algorithm and 3D region growing. In *Proc. SPIE Medical Imaging*, pages 906–916, 2000.
- K. S. Lee and P. M. Boiselle. Tracheal and bronchial neoplasms. In Phillip M. Boiselle and David Lynch, editors, *CT of the Airways*, pages 151–190. Humana Press - Springer, 2008.
- K. S Lee, W. Lunn, D. Feller-Kopman, A. Ernst, H. Hatabu, and P. M. Boiselle. Multislice CT evaluation of airway stents. *Journal of Thoracic Imaging*, 20(2):81–88, 2005.
- P. Lee, E. Kupeli, and A. C. Mehta. Airway stents. *Clinics in Chest Medicine*, 31(1):141 – 150, 2010.

-
- D. R. Lipsa, P. J. Rhodes, R. D. Bergeron, and T. M. Sparr. Spatial prefetching for out-of-core visualization of multidimensional data. In *Proc. SPIE, Visualization and Data Analysis*, volume 6495, 2007.
- P. Lo, B. van Ginneken, J. M. Reinhardt, and M. de Bruijne. Extraction of airways from CT (exact'09). In *2nd International Workshop on Pulmonary Image Analysis*, pages 175–189, London, England, 2009.
- W. E. Lorensen and H. E. Cline. Marching cubes: A high resolution 3D surface construction algorithm. *SIGGRAPH Computer Graphics*, 21(4): 163–169, 1987.
- M. Mandour, M. Remacle, P. van de Heyning, S. Elwany, A. Tantawy, and A. Gaafar. Chronic subglottic and tracheal stenosis: endoscopic management vs. surgical reconstruction. *European Archives of Oto-Rhino-Laryngology*, 260(7):374–380, 2003.
- R. A. McLaughlin, J. P. Williamson, M. J. Phillips, J. J. Armstrong, S. Becker, D. R. Hillman, P. R. Eastwood, and D. D. Sampson. Applying anatomical optical coherence tomography to quantitative 3D imaging of the lower airway. *Optics Express*, 16(22):17521–17529, 2008.
- T. Michikawa, K. Tsuji, T. Fujimori, and H. Suzuki. Out-of-core distance transforms. In *SPM '07: Proceedings of the 2007 ACM symposium on Solid and physical modeling*, pages 151–158, 2007. ISBN 978-1-59593-666-0.
- T. Miyazawa. History of the flexible bronchoscope. In Bolliger C. T. and Mathur P. N., editors, *Interventional Bronchoscopy*, volume 30, pages 16–21. Karger, 2000.
- T. Miyazawa, Y. Miyazu, Y. Iwamoto, A. Ishida, K. Kanoh, H. Sumiyoshi, M. Doi, and N. Kurimoto. Stenting at the flow-limiting segment in tracheobronchial stenosis due to lung cancer. *American Journal of Respiratory and Critical Care Medicine*, 169(10):1096–1102, 2004.

- B. Mostafa. Endoluminal stenting for tracheal stenosis. *European Archives of Oto-Rhino-Laryngology*, 260(9):465–468, 2003.
- M. B. Nielsen, O. Nilsson, A. Söderström, and K. Museth. Out-of-core and compressed level set methods. *ACM Trans. Graph.*, 26(4):16, 2007.
- S. A. R. Nouraei, D. W. McPartlin, S. M. Nouraei, A. Patel, C. Ferguson, D. J. Howard, and G. S. Sandhu. Objective sizing of upper airway stenosis: A quantitative endoscopic approach. *The Laryngoscope*, 116(1):12–17, 2006.
- N. Otsu. A threshold selection method from gray-level histograms. *Systems, Man and Cybernetics, IEEE Transactions on*, 9(1):62–66, 1979.
- K. Palágyi and A. Kuba. A 3D 6-subiteration thinning algorithm for extracting medial lines. *Pattern Recognition Letters*, 19:613–627, 1998.
- P. K. Parida and A. K. Gupta. Role of spiral computed tomography with 3-dimensional reconstruction in cases with laryngeal stenosis—a radioclinical correlation. *American Journal of Otolaryngology*, 29(5):305 – 311, 2008.
- U. B. S. Prakash. Advances in Bronchoscopic Procedures. *Chest*, 116(5):1403–1408, 1999.
- M. Prasad, J. P. Bent, R. F. Ward, and M. M. April. Endoscopically placed nitinol stents for pediatric tracheal obstruction. *International Journal of Pediatric Otorhinolaryngology*, 66(2):155–160, 2002.
- S. Prohaska, A. Hutanu, R. Kahler, and H. Hege. Interactive exploration of large remote micro-ct scans. In *VIS '04: Proceedings of the conference on Visualization '04*, pages 345–352, 2004. ISBN 0-7803-8788-0.
- B. R., V. S. Kumar, T. Pan, T. Kurc, U. Catalyurek, J. Saltz, and Y. Wang. Distributed out-of-core preprocessing of very large microscopy images for efficient querying. *Cluster Computing, IEEE International Conference on*, 0:1–10, 2005.

-
- K. T. Rajamani, M. A. Styner, H. Talib, G. Zheng, L. P. Nolte, and M. A. Gonzalez Ballester. Statistical deformable bone models for robust 3D surface extrapolation from sparse data. *Medical Image Analysis*, 11(2):99–109, 2007.
- P. J. Rhodes, X. Tang, R. Bergeron, and T. SM. parr. Out of core visualization using iterator aware multidimensional prefetching. In *Proc. SPIE, Visualization and Data Analysis*, volume 5669, pages 295–306, 2005a.
- P. J. Rhodes, X. Tang, R. D. Bergeron, and T. M. Sparr. Iteration aware prefetching for large multidimensional datasets. In *SSDBM'2005: Proceedings of the 17th international conference on Scientific and statistical database management*, pages 45–54, 2005b.
- M. Rogers and J. Graham. Robust active shape model search. In *Proc. of the 7th European Conference on Computer Vision-Part IV*, pages 517–530. Springer-Verlag, 2002.
- Y. Saito and H. Imamura. Airway stenting. *Surgery Today*, 35(4):265–270, 2005.
- V. Sauret, P.M. Halson, I.W. Brown, J.S. Fleming, and A.G. Bailey. Study of the three-dimensional geometry of the central conducting airways in man using computed tomographic (CT) images. *Journal of Anatomy*, 200(2):123–134, 2002.
- L. G. Shapiro and G. C. Stockman. *Computer Vision*. Prentice Hall, 2002.
- D. Shitrit, P. Valdislav, A. Grubstein, D. Bendayan, M. Cohen, and M. R. Kramer. Accuracy of virtual bronchoscopy for grading tracheobronchial stenosis: correlation with pulmonary function test and fiberoptic bronchoscopy. *Chest*, 128(5):3545–3550, 2005.
- C. T. Silva, Y. J. Chiang, J. El-Sana, and P. Lindstrom. Out-of-core algorithms for scientific visualization and computer graphics. Course Notes for IEEE Visualization Conference 2002, Boston, MA, 2002.

- E. Sorantin, C. Halmai, B. Erdôhelyi, K. Palágyi, L.G. Nyúl, L. K. Ollé, B. Geiger, F. Lindbichler, G. Friedrich, and K. Kiesler. Spiral-CT-based assessment of tracheal stenoses using 3-D-skeletonization. *Medical Imaging, IEEE Transactions on*, 21(3):263–273, 2002.
- E. Sorantin, D. Mohadjer, L. G. Nyúl, K. Palágyi, F. Lindbichler, and B. Geiger. New advances for imaging of laryngotracheal stenosis by post processing of spiral-ct data. In W. Hrubby, editor, *Digital (R)Evolution in Radiology - Bridging the Future of Health Care*, pages 297–308. Springer, 2006.
- N. Spittle and A. McCluskey. Lesson of the week: Tracheal stenosis after intubation. *British Medical Journal*, 321(7267):1000–1002, 2000.
- S. Stamenkovic, R. Hierner, P. De Leyn, and P. Delaere. Long-segment tracheal stenosis treated with vascularized mucosa and short-term stenting. *Annals of Thoracic Surgery*, 83(3):1213–1215, 2007.
- Jr Stephens, K. E. and D. E. Wood. Bronchoscopic management of central airway obstruction. *Journal of Thorac and Cardiovascular Surgery*, 119(2):289–296, 2000.
- B. Theobald, I. Matthews, and S. Baker. Evaluating error functions for robust active appearance models. In *Proc. of the International Conference on Automatic Face and Gesture Recognition*, pages 149 – 154, 2006.
- J. Tian, J. Xue, Y. Dai, J. Chen, and J. Zheng. A novel software platform for medical image processing and analyzing. *IEEE Transactions on Information Technology in Biomedicine*, 12(6):800–812, 2008.
- J. Triglia, S. Marciano, B. Nazarian, G. Moulin, I. Sudre-Levillain, and A. Giovanni. Virtual laryngotracheal endoscopy based on geometric surface modeling using spiral computed tomography data. *The Annals of otology, rhinology & laryngology*, 111(1):36–43, 2002.

-
- J. Tschirren, E. A. Hoffman, G. McLennan, and M. Sonka. Intrathoracic airway trees: segmentation and airway morphology analysis from low-dose CT scans. *Medical Imaging, IEEE Transactions on*, 24(12):1529–1539, 2005.
- J. K. Udupa and S. Samarasekera. Fuzzy connectedness and object definition: Theory, algorithms, and applications in image segmentation. *Graphical Models and Image Processing*, 58(3):246 – 261, 1996.
- R. Valdés-Cristerna and O. Yáñez-Suárez. Active contours and surfaces with cubic splines for semiautomatic tracheal segmentation. *Journal of Electronic Imaging*, 12(1):81–96, 2003.
- R. Valdés-Cristerna, O. Yáñez-Suárez, and V. Medina. Trachea segmentation in CT images using active contours. In *Engineering in Medicine and Biology Society, 2000. Proc. of the 22nd Annual International Conference of the IEEE*, volume 4, pages 3184–3187, 2000.
- C. M. van Bommel, M. A. Viergever, and Wiro J. Niessen. Semiautomatic segmentation and stenosis quantification of 3D contrast-enhanced MR angiograms of the internal carotid artery. *Magnetic Resonance in Medicine*, 51(4):753–760, 2004.
- B. van Ginneken, M. de Bruijne, M. Loog, and M. A. Viergever. Interactive Shape Models. In *Proc. of SPIE Medical Imaging*, volume 5032, pages 1206–1216, San Diego, CA, USA, 2003.
- B. van Ginneken, W. Baggerman, and E. M. van Rikxoort. Robust segmentation and anatomical labeling of the airway tree from thoracic CT scans. In *Medical Image Computing and Computer-Assisted Intervention MICCAI*, pages 219–226, 2008.
- F. Venuta, E. A. Rendina, and T. de Giacomo. Airway stenting. http://www.ctsnet.org/sections/clinicalresources/thoracic/expert_tech-1.html, 2004.

-
- J. S. Vitter. External memory algorithms and data structures: dealing with massive data. *ACM Computing Surveys*, 33(2):209–271, 2001.
- Q. Wang, J. JaJa, and A. Varshney. An efficient and scalable parallel algorithm for out-of-core isosurface extraction and rendering. *Journal of Parallel and Distributed Computing*, 67(5):592–603, 2007.
- E. M. Webb, B. M. Elicker, and W. R. Webb. Using CT to diagnose nonneoplastic tracheal abnormalities: Appearance of the tracheal wall. *American Journal of Roentgenology*, 174(5):1315–1321, 2000.
- J. P. Williamson, J. J. Armstrong, R. A. McLaughlin, P. B. Noble, A. R. West, S. Becker, A. Curatolo, W. J. Noffsinger, H. W. Mitchell, M. J. Phillips, D. D. Sampson, D. R. Hillman, and P. R. Eastwood. Measuring airway dimensions during bronchoscopy using anatomical optical coherence tomography. *European Respiratory Journal*, 35(1):34–41, 2010.
- D. E. Wood, Y. Liu, E. Vallieres, R. Karmy-Jones, and M. S. Mulligan. Airway stenting for malignant and benign tracheobronchial stenosis. *Annals of Thoracic Surgery*, 76(1):167–174, 2003.
- X. Xie and M. Mirmehdi. MAC: Magnetostatic active contour model. *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, 30(4):632–646, 2008.
- C. Xu and J. L. Prince. Snakes, shapes, and gradient vector flow. *Image Processing, IEEE Transactions on*, 7(3):359–369, 1998.
- S. Zachow, H. Lamecker, B. Elsholtz, and M. Stiller. Reconstruction of mandibular dysplasia using a statistical 3D shape model. In *Proc. of Computer Assisted Radiology and Surgery*, volume 1281, pages 1238–1243. Elsevier, 2005.

Nederlandse Samenvatting

Tracheale stenose is een onnatuurlijke vernauwing van de trachea die, hoewel ze zeldzaam is, levensbedreigend kan zijn. Mogelijke opties voor behandeling van tracheale stenose zijn een chirurgische ingreep of het gebruik van een tracheale stent. Ongeacht de behandeling is kwantificatie van de stenose een belangrijke stap in de diagnose, met name het bepalen van de exacte positie en de graad van de vernauwing. Momenteel hangt de specificatie van dergelijke parameters nog sterk af van de expertise van de radioloog en beroept men zich, ondanks dat de beeldvorming en beeldverwerking vaak geavanceerd zijn, nog steeds op manuele of semi-automatische procedures.

Deze thesis beschrijft een beslissingsondersteunend systeem voor de automatische bepaling van stenoseparameters uit CT opnamen van de thorax. Bovendien worden voor de behandeling van de tracheale stenose op een automatische wijze de diameter en lengte van de stent berekend. Het belangrijkste concept van de methode is het accuraat schatten van de vorm van de trachea indien de stenose niet aanwezig zou zijn. De berekende gezonde tracheavorm kan door de arts worden gebruikt om operaties te plannen en is tevens de basis voor een automatische evaluatie van de stenose en de voorspelling van patiëntspecifieke stents.

Om tot een dergelijk stentbepalings- en predictiesysteem te komen moesten enkele moeilijkheden overwonnen worden. De eerste uitdaging was het efficiënt visualiseren en verwerken van grote medische bestanden. Met de huidige technologie kunnen CT scans van de thorax gemakkelijk meerdere gigabytes in beslag nemen. Als deze bestanden met de huidige

desktop computersystemen worden verwerkt, botst men vrij snel op computergeheugenproblemen. In deze thesis wordt een *out-of-core* beeldverwerkingstechniek voorgesteld die enkel een klein deel van dergelijke grote bestanden in het geheugen houdt. Deze techniek, *WindowCache*, gebruikt een cache en prefetching strategie die het doorkruisen van bestanden optimaliseert gebruik makend van voorkennis over het datatoegangspatroon van het beeldverwerkingsalgoritme.

Een tweede uitdaging was de automatische segmentatie van gezonde trachea's in CT beelden, een belangrijke voorverwerkingsstap van de voorgestelde methode. Naast segmentatie van de trachea wordt in deze thesis eveneens een methode beschreven voor de segmentatie van een meer volledige luchtwegenstructuur, bestaande uit de trachea en de bronchi. De voorgestelde methode identificeert automatisch het begin- en eindpunt van de trachea en introduceert een nieuwe techniek om lekken in de segmentatie te voorkomen.

De schatting van de gezonde tracheavorm van een patient met tracheale stenose wordt vervolgens bereikt mbv een statistisch vormmodel van gezonde trachea's. Een dergelijk model, zoals beschreven door Cootes et al. (1995) beschrijft de geometrische variatie die aanwezig is in een (trainings)set van vormen. Het vormmodel kan tevens nieuwe vormen genereren binnen, evenwel qua variatie beperkt door de vormen binnen de trainingsset; m.a.w: het hier voorgestelde model kan enkel gezonde tracheas genereren. Dit 'gezond' model kan vervolgens geregistreerd worden in een CT beeld van een patient met tracheale stenose. Hoewel de resulterende vorm van het model na registratie geen lokale geometrische vervormingen bevat die eigen zijn aan stenosis, beïnvloeden de vernauwde gebieden in de trachea van de patiënt toch globaal de registratieprocedure. In deze thesis wordt een nieuwe registratiemethode voorgesteld, *Fixed Landmarks* die op een robuustere wijze de invloed van deze vernauwde gebieden vermijdt en die een plausibele schatting genereert van de gezonde trachea van de patiënt.

In een volgende stap segmenteert het voorgestelde systeem, op basis van de gezonde trachea, automatisch de vernauwde trachea uit het beeld. De segmentatie is gebaseerd op het klassieke Actieve Contour Model van

Kass et al. (1988), met gepaste interne en externe krachten. De zo bekomen gezonde en vernauwde versies van de trachea worden uiteindelijk gebruikt in een automatische procedure die de parameters van de stenose berekent (beginpunt, eindpunt, graad van vernauwing), waarbij ahv deze parameters de lengte en de diameter van de corrigerende stent wordt voorspeld.

Uitgebreide experimenten werden uitgevoerd ter validatie van de voorgestelde methoden. Voor de *out-of-core* techniek werden verschillende algoritmen en caching structuren getest op zeer grote beelden. De luchtwe-gensegmentatiemethode werd getest op een databank van 40 beeldvolumes, onderverdeeld in training en test sets. Voor het statistisch vormmodel werden 38 gezonde tracheas gebruikt. Tenslotte werd de bepaling van stenose en de voorspelling van stentparameters gevalideerd via uitgebreide simulatie-experimenten en met een retrospectieve studie met 11 CT scans van 9 patienten. De resultaten tonen dat het voorgestelde systeem effectief de operator-afhankelijkheid reduceert bij de bepaling van tracheale stenose. Het voorgestelde systeem bleek bovendien snel en accuraat genoeg voor gebruik in de kliniek.

Agradecimentos

Apesar de contrário aos pré-requisitos, diz-se que um doutorado nunca é feito sozinho. De fato, a lista de pessoas que fizeram parte do meu projeto Doutorado não é curta. Vou começar com aqueles que estiveram mais próximos de mim durante todo o tempo.

Eu não poderia ser mais feliz por ter pais, Cleber e Madalena, tão maravilhosos e amáveis, que independente da angústia que sentiram ao terem o único filho morando fora por quase 5 anos¹, me deram muito apoio. Todo o seu sacrifício para me dar boa educação está sendo recompensado agora, e minha conquista nada menos é que deles também. Sobre esta conquista, há um suposto provérbio sueco que diz: “sofrimento dividido é meio sofrimento e alegria dividida é dupla alegria”. Meu Doutorado não teria feito qualquer sentido sem o amor da minha vida, Adriana, para dividi-lo comigo. Ela é parte de cada palavra deste texto, de cada linha de código que tive que escrever, de cada hora extra de trabalho. Ela estava lá para engrandecer os bons momentos e para animar aqueles não tão bons. Tenho certeza de que ela está tão feliz por mim agora quanto eu fiquei por ela quando ela defendeu o seu Doutorado.

Ainda sobre provérbios, minha avó Dimas disse repetidas vezes durante meu Doutorado duas coisas que nunca vou esquecer: “é muito importante trabalhar em algo de que se gosta.” e “Eu estou sempre rezando para que você vença”. Vencendo ou não, e apesar da distância física, eu sei que ela sempre esteve do meu lado, assim como sempre estive do seu. Minha

¹Bem-aventurados os que inventaram a webcam, não é mãe?

querida sogra, Regina, é alguém que também sempre esteve do meu lado, e eu tenho certeza de que ela me apoiou tanto quanto se eu fosse seu próprio filho.

Mas quando se é um filho, neto ou genro morando num país diferente, levando tempo para se adaptar, é o trabalho que nos preenche na maior parte do tempo. Um bom assunto no qual trabalhar e um bom orientador são portanto partes importantes do processo. Meu orientador foi o Prof. Dr. Jan Sijbers. Acredito que não vou exagerar se disser que Jan é unanimamente tido com a maior consideração por seus alunos de Doutorado. Ele sempre esteve disponível para discussões, mesmo via chat nos domingos à noite, e frequentemente de bom humor. Ele confiou cegamente na minha capacidade de trabalhar à distância, permitindo que eu terminasse minha tese fora da Bélgica, pelo que eu fiquei muito grato. Por estas e outras razões, ele provou ser não só um grande orientador, mas uma excelente pessoa também.

Ainda assim, um bom orientador é apenas parte da equação. Um bom ambiente onde se trabalhar é obviamente complementar. Também encontrei isto no VisionLab, por exemplo, na alegria do meu bom amigo Sander van der Maar (que foi sempre generoso em me ter como hóspede nas vezes que precisei viajar para a Bélgica), na simpatia de Ruben Capiou e nossos papos sobre fotografia, na prestatividade de Toon Huysmans e Gert van Gompel, nas partidas de badminton com Mai Zhenhua, no sorriso simpático de Juntu Jaber, nas discussões de projeto enriquecedoras com Joost Batenburg, etc. Todos eles, mesmo os que vieram e partiram, contribuíram do seu jeito para fazer o VisionLab um lugar muito bom onde se trabalhar.

Eu também tive o prazer de conhecer o Dr. Kurt Tournoy durante o desenvolvimento da minha tese. Em um momento no qual eu parecia bloqueado devido à falta de imagens para testar meu método, Dr. Tournoy apareceu não só como o provedor de tais imagens, mas também sempre esteve disposto à colaborar, sendo uma importante fonte de informações médicas.

Mas antes de vir para o VisionLab, eu precisei sair do meu emprego no Brasil. Sair de um emprego nunca é uma tarefa fácil, especialmente de um do qual se gosta. Mas quando se tem um chefe como o Sr. Sílvia Pereira da

TV Globo, a tensão do momento se transforma em apoio imensurável. Não é por acaso que tenho o Sílvio como um dos meus mais queridos amigos hoje.

Querido também é o Prof. Dr. Alexandre Plastino, professor exemplar da minha época de graduação, na Universidade Federal Fluminense (UFF). Alexandre demonstrou enorme disponibilidade quando lhe pedi cartas de referências para a vaga de Doutorado. Também não é por acaso que ele é um dos grandes nomes na UFF.

Outras pessoas fizeram parte indiretamente deste projeto Doutorado. Pessoas como Sabine Covent, que hospedou à minha esposa e a mim em sua pousada durante nossas primeiras semanas na Bélgica e nos deu muita ajuda (e até uma cama) enquanto estávamos nos alojando. Por último, Armand e Marília Janssens, muito mais do que ótimos senhorios, foram pessoas de uma generosidade ilimitada, que levaremos em nossos corações onde quer que estejamos.

Apesar de ser uma pessoa bastante racional, eu acredito que há momentos na vida que transcendem as explicações dadas pela matemática, física, estatística, lógica, acaso, sorte, mérito etc. O caminho deste projeto Doutorado acabou indo além de qualquer plano ou estratégia, algo que só uma força superior é capaz de controlar. De maneiras que eu jamais serei capaz de compreender, esta força trilhou o caminho (às vezes tortuoso) que me guiou até este momento muito especial e colocou nele as pessoas e as coisas que me ajudaram a fazer este sonho de Doutorado se tornar realidade. Portanto, eu ofereço a Você, meu Deus, duas palavras humildes, mas sinceras:

MUITO OBRIGADO.