Quantitative Image Analysis in Cardiac CT Angiography

Michiel Schaap

Quantitative Image Analysis in Cardiac CT Angiography
Michiel Schaap
PhD thesis, Erasmus University Rotterdam, the Netherlands.

The work in this thesis was conducted at the Departments of Radiology and Medical Informatics at the Erasmus MC University Medical Center, Rotterdam, The Netherlands. The research was supported by the Stichting voor Technische Wetenschappen (STW) of The Netherlands Organization for Scientific Research (NWO). Financial support for the publication of this thesis was provided by the Department of Radiology of the Erasmus MC University Medical Center and the Erasmus University Rotterdam.

ISBN: 978-90-6464-441-2
Cover design by Studio Harry Zoutendijk and Michiel Schaap
Printed by GVO drukkers & vormgevers B.V. | Ponsen & Looijen

Quantitative Image Analysis in Cardiac CT Angiography

Kwantitatieve beeldanalyse in CT angiografie van het hart

Proefschrift

ter verkrijging van de graad van doctor aan de Erasmus Universiteit Rotterdam op gezag van de rector magnificus

Prof.dr. H.G. Schmidt

en volgens besluit van het College voor Promoties.

De openbare verdediging zal plaatsvinden op woensdag, 15 december 2010 om 11:30 uur

door

Michiel Schaap

geboren te Hilversum.

Promotiecommissie

Promotor:	Prof.dr. W.J. Niessen
Overige leden:	Prof.dr. P.J. de Feyter Dr. B.P.F. Lelieveldt Dr. B. van Ginneken
Copromotor:	Dr.ir. Th. van Walsum

Contents

Chapter 1

Introduction

1.1 Introduction

Despite advances in diagnosis, treatment and interventions, cardiovascular disease remains a major health problem in the western world. Owing to the ageing society and changing lifestyle, cardiovascular disease affects an increasing number of individuals [92, 110]. Often, there are no symptoms prior to clinical cardiovascular disease related events such as myocardial infarction (heart attack), and irreversible damage or death may occur prior to any medical intervention [183]. Advances in imaging technology have great potential to improve the management of cardiovascular disease through early detection, improved diagnosis and treatment monitoring [30]. This requires not only advanced techniques to image cardiovascular anatomy and function, but also robust methods to extract relevant (quantitative) parameters from these imaging data and novel techniques that effectively visualize diagnostically relevant image information.

This thesis describes the development and validation of quantitative image analysis and visualization methods, which are essential steps towards improved image-based diagnosis and prognosis of cardiovascular disease. In the subsequent sections we will introduce the cardiac anatomy, coronary artery disease, CT angiography (CTA) imaging of the coronary arteries, and methods for quantitative CTA image analysis. Following the introduction, we will provide an outline of this thesis in Section 1.2.

1.1.1 The heart and the coronary arteries

The heart is a vital organ that pumps oxygen-deficient blood from the body to the lungs and oxygen-rich blood from the lungs to the body. It consists of four chambers, the left and right atrium and the left and right ventricle, of which the left ventricle pumps the oxygen-rich blood through the body (see Figure 1.1). Approximately 5% of the blood is pumped through the coronary arteries to the myocardium (the heart muscle) to feed it with oxygen and nutrients [100]. The coronary arteries are composed of two vessel trees that cover the complete myocardium (see Figure 1.2). It is of key importance that the flow of blood through the coronaries is maintained, as the heart will lose its pumping function in case of insufficient blood supply. The left coronary tree is responsible for most of the oxygen supply to the left ventricle

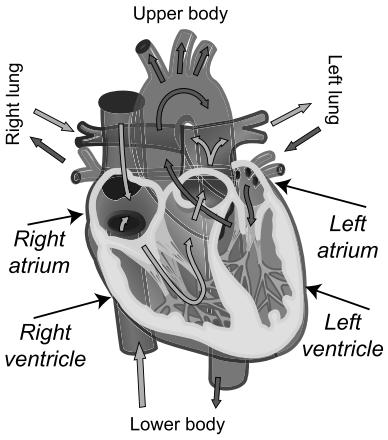

Figure 1.1: *A schematic visualization of the human heart.*

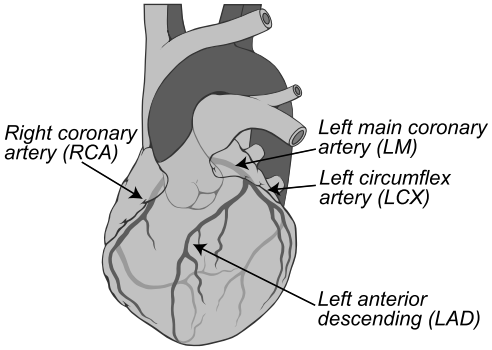

Figure 1.2: *A schematic visualization of the four main coronary arteries.*

and the right coronary arteries mainly feed the right ventricle. Both trees bifurcate into many small branches and the topology of these trees can differ significantly between patients. Four main branches can be differentiated in virtually all people: the left main coronary artery (LM), the left ascending artery (LAD), the left circumflex (LCX), and the right coronary artery (RCA).

1.1.2 Coronary artery disease

Coronary artery disease (CAD), or coronary atherosclerosis, is a major cause of death. Worldwide, approximately 1 out of every 5 people die of coronary artery disease [182]. Risk factors for atherosclerosis include hypertension (high blood pressure), hypercholesterolemia (high cholesterol), diabetes mellitus, obesity, and tobacco use. Atherosclerosis is a complex disease that affects the vessel wall, and results in a thickening or hardening of the vessel wall that eventually may compromise blood flow through the coronaries. Atherosclerosis is thought to be caused by an inflammation between the two inner layers of the vessel wall, causing a buildup of plaque, consisting of lipids, such as cholesterol, and fibrous tissue, in the vessel wall [90]. Calcifications (calcified parts of atherosclerotic plaque in the vessel wall) may develop in these plaque regions, and traditionally these calcifications have been

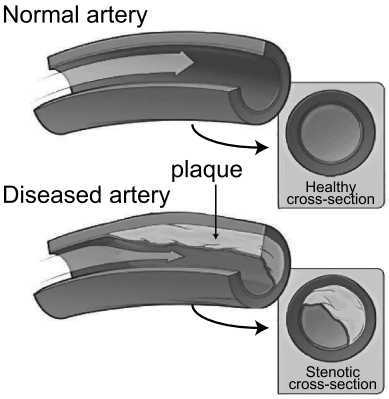

Figure 1.3: *A schematic illustration of coronary artery disease. The top figure shows an artery with normal blood flow. The bottom figure shows an artery after plaque buildup.*

an important marker for CAD. Plaque build-up and associated thickening of the coronary vessel wall can cause a stenosis (narrowing of the vessel lumen; the area where blood flows through the arteries), resulting in insufficient blood supply to the myocardium (see Figure 1.3). Moreover, after plaque build-up, even without vessel lumen narrowing, the vessel wall can rupture, causing release of thrombogenic plaque into the vessel lumen and subsequent occlusion of the vessel, resulting, in most cases, in acute myocardial infarction [90] .

1.1.3 Computed Tomography Angiography

Currently image-based diagnosis of CAD focuses mainly on stenotic (i.e. narrowed) vessels. The gold standard for the assessment of the coronary lumen is conventional coronary angiography (CCA) [29], an invasive catheter-based X-ray medical image modality. However, because of its invasive nature, CCA has a low, but non-negligible, risk of procedure related complications [189]. Moreover, it is a two-dimensional modality and only provides information on the coronary lumen and can not be used to visualize the coronary plaque. Figure 1.5 shows an example of a CCA image.

Figure 1.4: *A CT scanner (SOMATOM Definition Flash, Siemens Healthcare). A patient is placed on the table that moves through the gantry aperture around which X-ray source and detectors rotate.*

Computed Tomography (CT) and Computed Tomography Angiography (CTA) are potential alternatives for CCA [109]. CT images the attenuation of X-rays by the human body using attenuation measurements made with an X-ray source and detector array that rotate around the patient (see Figure 1.4). After measuring the X-ray attenuation along a large number of paths dedicated software calculates the attenuation in a large grid of small volumetric picture elements (voxels), thereby providing a detailed 3-dimensional picture of the human heart. Since the first commercial deployment of a CT scanner, which was designed to image the human brain, in the early 1970s [67], technological advances in CT scanner hard- and software have been tremendous. The number of axial slices acquired simultaneously has increased from 1 to 320, the resolution has increased significantly with voxel sizes being reduced from approximately 3 cm to less than half a millimeter, the radiation dose has been decreased significantly, and the rotation time has decreased from approximately 2 minutes to 0.3 seconds [3, 138]. These developments have made CT the imaging modality of choice for visualizing the rapidly moving coronary arteries with high resolution. Moreover, recently dual-source CT scanners have been clinically

introduced, which has further increased the temporal resolution of the CT acquisition and provides dual-energy possibilities, such as better material differentiation [69]. All these developments have made CT and CTA very promising imaging modalities for three-dimensional minimally-invasive imaging of the heart [42]. With CT calcified plaques can be visualized. Coronary calcifications can be appreciated as small, high density spots in the vessel wall in CT scans. Several studies have demonstrated the prognostic value of the amount of calcium quantified with CT for future cardiac events [118, 122].

In CTA, intravenously injected iodine-based contrast fluid is used to highlight the coronary lumen. This imaging modality is expected to allow, next to the assessment of the coronary lumen morphology, the evaluation of the presence, extent, and type (non-calcified or calcified) of coronary plaque [2, 82, 148]. This may be relevant for improving risk stratification when combined with current non-image based risk factors, such as hypertension, hypercholesterolemia, traditional image-based measures such as the severity of stenosis and the amount of calcium [29], and genetic risk factors [151].

A disadvantage of CTA is that the resolution, although 3D, is lower than the 2D resolution of CCA. In Figure 1.5 we show an example of CCA, CT and CTA of the same lesion in left circumflex artery (LCX).

1.1.4 Advanced visualization and quantitative image analysis

On average, a modern clinical cardiac CTA image contains approximately $512 \times 512 \times 300$ voxels (i.e. approximately 75 megavoxels). Inspecting these images can be labour-intensive for a radiologist or physician. Therefore, advanced post-processing techniques are being developed to support effective diagnosis with CTA data. For example, techniques have been developed to reformat vascular data such that the lumen, wall, and surrounding tissue of a vessel can be appreciated in one plane [73] (see Figure 1.5(d)).

Besides techniques for improved visual inspection, a major step towards improved diagnosis and prognosis of CAD is a more objective and automated assessment of the state and presence of CAD. This has lead to development of quantitative image analysis (QIA) techniques. These techniques allow the automated extraction of quantitative measurements, such as absolute measures of morphology and function, from images with computerized image processing techniques [50]. The robust extraction of accurate and objective image-based measurements indicating the presence or severity of disease (the so-called 'Quantitative Imaging Biomarkers'), is expected to have significant influence on the clinical decision making process [162].

Quantitative image analysis finds its application in two major areas. First, QIA approaches are continuously improving the functionality of computer aided diagnosis systems because QIA allows more objective, cost efficient establishment of the presence and severity of disease, potentially with higher sensitivity and specificity

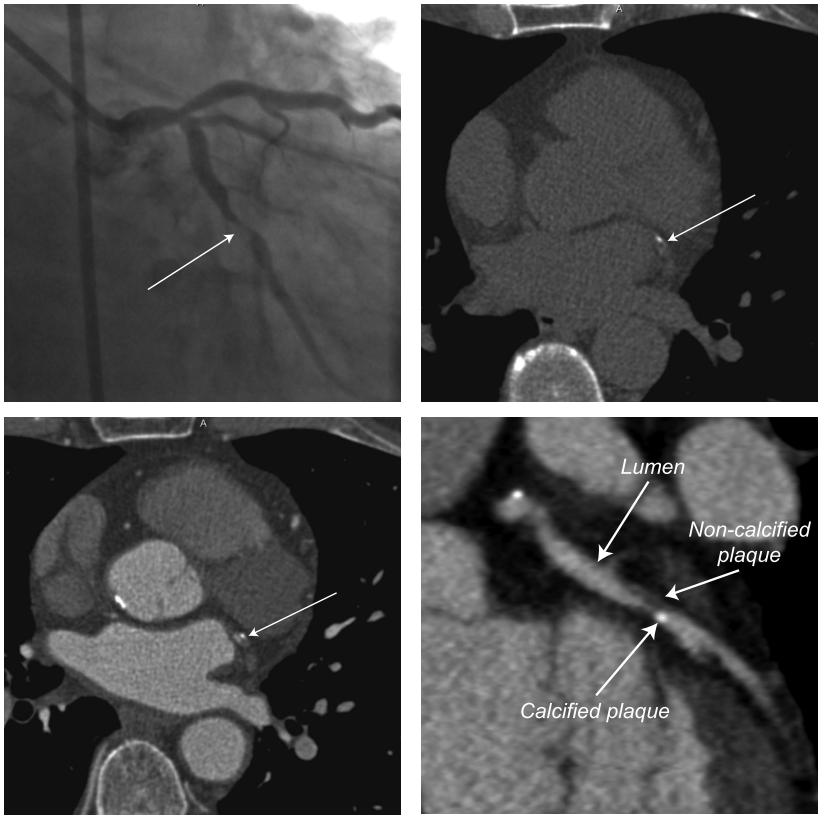

Figure 1.5: *Examples of different techniques for the visualization of the same coronary lesion in the left circumflex (LCX) artery. (a) conventional coronary angiography (CCA), (b) Computed tomography (CT), (c) Computed tomography angiography (CTA), and (d) a planar reformation of the CTA dataset in (c) showing a magnification of the lesion and its components.*

than human experts [50]. For example, in the diagnosis of coronary artery disease already for over two decades QIA methods are being used to quantify the coronary lumen morphology in conventional coronary angiography (CCA) [128] and new and improved systems are continuously being proposed for this application [124, 159]. In addition, an increasing interest in quantitative imaging biomarkers for treatment monitoring, and in the discovery and development of novel drugs is foreseen, as potentially they can be used as surrogate end points [125, 131, 161, 162].

1.1.5 Coronary artery lumen segmentation in CTA

Nowadays, the quantification of cardiovascular disease in CTA is still very much under investigation. Sophisticated software tools are available to manually outline the coronary arteries in CTA, but this process is too time-consuming for clinical practice and prone to inter- and intra-observer variability. Several automated methods are available for the quantification of the coronary lumen, but only very few of these systems have been evaluated thoroughly [142]. Recently, commercial CT-based computer aided detection systems have been presented with promising results (see e.g. [6]), but extensive validation of these systems is often lacking and quantification results differ significantly between commercial systems [59]. Moreover, at this point no system has been presented which includes a well-validated accurate plaque differentiation and quantification method.

The quantification of cardiovascular disease in cardiac CTA images in general starts with the delineation, or segmentation, of the coronary lumen. Subsequently the topology and morphology of the coronary lumen can be quantified, plaque can be delineated, and lumen and plaque parameters can be quantified [37, 75, 178, 187].

The segmentation of the vascular lumen from the image is often tackled in a two-step coarse-to-fine fashion (see e.g. [75, 101, 178]): first a method is used to locate the coronary arteries in the CTA image and then the vessel is segmented given its location. In the first step the algorithm uses the complete image to extract a set of curves running through the center of the vascular lumen. This step is referred to as centerline extraction. In the second step an algorithm uses local image information around these curves to segment the vessels. Especially centerline extraction in medical images has received considerable attention, as it can be used as a preprocessing step for vessel segmentation, but also to visualize cardiac CTA data with multi-planar reformatting (MPR) and curved planar reformatting (CPR) visualization techniques [73]. These reformatted images can for example be used to diagnose CAD, but also in interventional cardiology, to plan the type of intervention and size of stents [57]. Although several commercial systems exist for coronary segmentation, up to now no well-validated systems are described in scientific literature. For an extensive review of existing literature on coronary centerline extraction and segmentation we would like to refer to Chapters 2 and 5.

Coronary centerline extraction and lumen segmentation are both challenging tasks owing to the small size of the coronary arteries visible in cardiac CTA (their diameter ranges from less than 1 mm to approximately 5 mm in diameter), the limited spatial resolution of CT (approximately 0.7 mm to 1.4 mm [132]), motion induced blurring, high intensity calcium close to the coronary lumen, and the presence of severe stenoses.

1.2 This thesis

The purpose of the work described in this thesis was the development and evaluation of techniques for the quantification and improved visualization of cardiovascular disease in CTA images. In this thesis, we focus on four aspects of cardiovascular CTA image processing: 1) we introduce methods for the objective evaluation of different coronary centerline extraction techniques; 2) we propose a probabilistic centerline tracking technique; 3) we propose two novel methods for coronary segmentations; 4) we propose an efficient method for noise reduction in CTA for improved 3D visualization.

In **Chapter 2** we focus on the objective evaluation of different approaches for (semi-)automatically extracting coronary centerlines from cardiac CTA data. Many methods have been proposed for this task, but almost none of these have been thoroughly evaluated. It is therefore difficult to assess the performance of these methods. We present a publically available standardized evaluation framework for these algorithms. The framework consists of a representative set of cardiac CTA datasets, an accurate manually annotated reference standard, an appropriate set of evaluation measures, and an accessible web-based framework for easy comparison of different algorithms.

Probabilistic, or multiple hypotheses, centerline extraction methods have recently been presented in order to circumvent problems common with standard path tracking methods [45, 48]. **Chapter 3** presents several improvements in the field of probabilistic centerline extraction. These methods allow easy incorporation of prior information about the geometry and appearance of the vessel of interest in the image. By jointly optimizing these probabilities the path is found that is optimal according to both the geometrical and appearance prior information.

Because of their probabilistic optimization strategy, these algorithms, in contrast to many existing algorithms, can find optimal solutions to the objective function described by advanced geometrical and appearance priors. Of course, the geometric and appearance model are essential for these approaches: the structure of interest will be incorrectly segmented if the assumptions underlying these models are violated. In this chapter we present a new appearance model that can better handle situations when the vessel is running through areas with significantly changing background intensities. Furthermore, because probabilistic centerline extraction algorithms are computationally expensive, steps towards a computationally efficient implementation are suggested in this chapter.

After centerline extraction, the second step of coronary artery disease quantification in cardiac CTA is often the segmentation of the coronary arteries. We present two different approaches for the segmentation of the coronary lumen. **Chapter 4** presents a method based on a combination of graph-cut optimization and robust kernel regression. Graph-cut optimization is used to find a surface surrounding an initial coronary centerline that is most probable according to the intensity change

perpendicular to the surface. Then we remove any possible outliers from the resulting segmentation with a robust kernel regression technique.

In the next chapter we propose a supervised approach that can learn plausible vessel shapes and their appearance in annotated medical image data and subsequently apply this knowledge for the segmentation of unseen data. This machine learning-based segmentation approach, presented in **Chapter 5**, is more generic than the method presented in Chapter 4 and is expected to have a wider applicability.

Chapter 6 presents a noise reduction method that can be used to improve the diagnostic quality of advanced 3D visualizations of CTA scans. The method is based on the well-known anisotropic diffusion equation. Anisotropic diffusion has the disadvantage of being computationally complex and thereby relatively slow, which hampers its introduction into clinical practice. In order to overcome this problem we present a new speed optimized implementation for anisotropic diffusion. This optimized anisotropic diffusion implementation is used in a novel image enhancement method that improves the diagnostic quality of reconstructed CT image data.

The thesis is concluded with a summary in **Chapter 7** and a short description of several additional CT angiography image processing contributions in **Chapter 8**.

Chapter 2

Evaluation of Coronary Artery Centerline Extraction Algorithms.

Based on:

M. Schaap, C.T. Metz, T. van Walsum, A.G. van der Giessen, A.C. Weustink, N.R. Mollet, C. Bauer, H. Bogunović, C. Castro, X. Deng, E. Dikici, T. O'Donnell, M. Frenay, O. Friman, M. Hernández Hoyos, P.H. Kitslaar, K. Krissian, C. Kühnel, M.A. Luengo-Oroz, M. Orkisz, Ö. Smedby, M. Styner, A. Szymczak, H. Tek, C. Wang, S.K. Warfield, S. Zambal, Y. Zhang, G.P. Krestin, W.J. Niessen
Standardized Evaluation Methodology and Reference Database for Evaluating Coronary Artery Centerline Extraction Algorithms, Medical Image Analysis, 2009.

Efficiently obtaining a reliable coronary artery centerline from computed tomography angiography data is relevant in clinical practice. Whereas numerous methods have been presented for this purpose, up to now no standardized evaluation methodology has been published to reliably evaluate and compare the performance of the existing or newly developed coronary artery centerline extraction algorithms. This chapter describes a standardized evaluation methodology and reference database for the quantitative evaluation of coronary artery centerline extraction algorithms. The contribution of this work is fourfold: 1) a method is described to create a consensus centerline with multiple observers, 2) well-defined measures are presented for the evaluation of coronary artery centerline extraction algorithms, 3) a database containing thirty-two cardiac CTA datasets with corresponding reference standard is described and made available, and 4) thirteen coronary artery centerline extraction algorithms, implemented by different research groups, are quantitatively evaluated and compared. The presented evaluation framework is made available to the medical imaging community for benchmarking existing or newly developed coronary centerline extraction algorithms.

2.1 Introduction

Coronary artery disease (CAD) is currently the primary cause of death among American males and females [134] and one of the main causes of death in the world [183]. The gold standard for the assessment of CAD is conventional coronary angiography (CCA) [29]. However, because of its invasive nature, CCA has a low, but non-negligible, risk of procedure related complications [189]. Moreover, it only provides information on the coronary lumen.

Computed Tomography Angiography (CTA) is a potential alternative for CCA [109]. CTA is a non-invasive technique that allows, next to the assessment of the coronary lumen, the evaluation of the presence, extent, and type (non-calcified or calcified) of coronary plaque [82]. Such non-invasive, comprehensive plaque assessment may be relevant for improving risk stratification when combined with current risk measures: the severity of stenosis and the amount of calcium [29]. A disadvantage of CTA is that the current imaging protocols are associated with a higher radiation dose exposure than CCA [40].

Several techniques to visualize CTA data are used in clinical practice for the diagnosis of CAD. Besides evaluating the axial slices, other visualization techniques such as maximum intensity projections (MIP), volume rendering techniques, multiplanar reformatting (MPR), and curved planar reformatting (CPR) are used to review CTA data [29]. CPR and MPR images of coronary arteries are based on the CTA image and a central lumen line (for convenience referred to as centerline) through the vessel of interest [73]. These reformatted images can also be used during procedure planning for, among other things, planning the type of intervention and size of stents [57]. Efficiently obtaining a reliable centerline is therefore relevant in clinical practice. Furthermore, centerlines can serve as a starting point for lumen segmentation, stenosis grading, and plaque quantification [75, 101, 178].

This chapter introduces a framework for the evaluation of coronary artery centerline extraction methods. The framework encompasses a publicly available database of coronary CTA data with corresponding reference standard centerlines derived from manually annotated centerlines, a set of well-defined evaluation measures, and an on-line tool for the comparison of coronary CTA centerline extraction techniques. We demonstrate the potential of the proposed framework by comparing thirteen coronary artery centerline extraction methods, implemented by different authors as part of a segmentation challenge workshop at the Medical Image Computing and Computer-Assisted Intervention (MICCAI) conference [105].

In the next two sections we will respectively describe our motivation of the study presented in this chapter and discuss previous work on the evaluation of coronary segmentation and centerline extraction techniques. The evaluation framework will then be outlined by discussing the data, reference standard, evaluation measures, evaluation categories, and web-based framework. The chapter will be concluded by presenting the comparative results of the thirteen centerline extrac-

tion techniques, a discussion of these results, and a conclusion about the work presented.

2.2 Motivation

The value of a standardized evaluation methodology and a publicly available image repository has been shown in a number of medical image analysis and general computer vision applications, for example in the Retrospective Image Registration Evaluation Project [179], the Digital Retinal Images for Vessel Extraction database [153], the Lung Image Database project [5], the Middlebury Stereo Vision evaluation [145], the Range Image Segmentation Comparison [66], the Berkeley Segmentation Dataset and Benchmark [102], and a workshop and on-line evaluation framework for liver and caudate segmentation [164].

Similarly, standardized evaluation and comparison of coronary artery centerline extraction algorithms has scientific and practical benefits. A benchmark of state-of-the-art techniques is a prerequisite for continued progress in this field: it shows which of the popular methods are successful and researchers can quickly apprehend where methods can be improved.

It is also advantageous for the comparison of new methods with the state-of-the-art. Without a publicly available evaluation framework, such comparisons are difficult to perform: the software or source code of existing techniques is often not available, articles may not give enough information for re-implementation, and if enough information is provided, re-implementation of multiple algorithms is a laborious task.

The understanding of algorithm performance that results from the standardized evaluation also has practical benefits. It may, for example, steer the clinical implementation and utilization, as a system architect can use objective measures to choose the best algorithm for a specific task.

Furthermore, the evaluation could show under which conditions a particular technique is likely to succeed or fail, it may therefore be used to improve the acquisition methodology to better match the post-processing techniques.

It is therefore our goal to design and implement a standardized methodology for the evaluation and comparison of coronary artery centerline extraction algorithms and publish a cardiac CTA image repository with associated reference standard. To this end, we will discuss the following tasks below:

- Collection of a representative set of cardiac CTA datasets, with a manually annotated reference standard, available for the entire medical imaging community;
- Development of an appropriate set of evaluation measures for the evaluation of coronary artery centerline extraction methods;

- Development of an accessible framework for easy comparison of different algorithms;
- Application of this framework to compare several coronary CTA centerline extraction techniques;
- Public dissemination of the results of the evaluation.

2.3 Previous work

Approximately thirty papers have appeared that present and/or evaluate (semi-)automatic techniques for the segmentation or centerline extraction of human coronary arteries in cardiac CTA datasets. The proposed algorithms have been evaluated by a wide variety of evaluation methodologies.

A large number of methods have been evaluated qualitatively [12, 22, 31, 43, 45, 58, 81, 93, 96, 116, 129, 141, 156, 167, 177, 185, 187]. In these articles detection, extraction, or segmentation correctness have been visually determined. An overview of these methods is given in Table 2.1. Other articles include a quantitative evaluation of the performance of the proposed methods [27, 28, 37, 75, 80, 84, 88, 101, 104, 121, 178, 186]. See Table 2.2 for an overview of these methods.

None of the abovementioned algorithms has been compared to another and only three methods were quantitatively evaluated on both the extraction ability (i.e. how much of the real centerline can be extracted by the method?) and the accuracy (i.e. how accurately can the method locate the centerline or wall of the vessel?). Moreover, only one method was evaluated using annotations from more than one observer [104].

Four methods were assessed on their ability to quantify clinically relevant measures, such as the degree of stenosis and the number of calcium spots in a vessel [37, 75, 178, 187]. These clinically oriented evaluation approaches are very appropriate for assessing the performance of a method for a possible clinical application, but the performance of these methods for other applications, such as describing the geometry of coronary arteries [94, 192], can not easily be judged.

Two of the articles (Dewey et al. [37] and Busch et al. [28]) evaluate a commercially available system (respectively Vitrea 2, Version 3.3, Vital Images and Syngo Circulation, Siemens). Several other commercial centerline extraction and stenosis grading packages have been introduced in the past years, but we are not aware of any scientific publication containing a clinical evaluation of these packages.

2.4 Evaluation framework

In this section we will describe our framework for the evaluation of coronary CTA centerline extraction techniques.

Table 2.1: *An overview of CTA coronary artery segmentation and centerline extraction algorithms that were qualitatively evaluated. The column 'Time' indicates if information is provided about the computational time of the algorithm.*

Article	Patients / observers	Vessels	Evaluation details	Time
Bartz et al. [12]	1/1	Complete tree	Extraction was judged to be satisfactory.	Yes
Bouraoui et al. [22]	40/1	Complete tree	Extraction was scored satisfactory or not.	No
Carrillo et al. [31]	12/1	Complete tree	Extraction was scored with the number of extracted small branches.	Yes
Florin et al. [43]	1/1	Complete tree	Extraction was judged to be satisfactory.	Yes
Florin et al. [45]	34/1	6 vessels	Scored with the number of correct extractions.	No
Hennemuth et al. [58]	61/1	RCA, LAD	Scored with the number of extracted vessels and categorized on the dataset difficulty.	Yes
Lavi et al. [81]	34/1	3 vessels	Scored qualitatively with scores from 1 to 5 and categorized on the image quality.	Yes
Lorenz et al. [93]	3/1	Complete tree	Results were visually analyzed and criticized.	Yes
Luengo-Oroz et al. [96]	9/1	LAD & LCX	Scored with the number of correct vessel extractions. The results are categorized on the image quality and amount of disease.	Yes
Nain et al. [116]	2/1	Left tree	Results were visually analyzed and criticized.	No
Renard et al. [129]	2/1	Left tree	Extraction was judged to be satisfactory.	No
Schaap et al. [141]	2/1	RCA	Extraction was judged to be satisfactory.	No
Szymczak et al. [156]	5/1	Complete tree	Results were visually analyzed and criticized.	Yes
Wang et al. [167]	33/1	Complete tree	Scored with the number of correct extractions.	Yes
Wesarg et al. [177]	12/1	Complete tree	Scored with the number of correct extractions.	Yes
Yang et al. [187]	2/1	Left tree	Extraction was judged to be satisfactory.	Yes
Yang et al. [185]	2/1	4 vessels	Scored satisfactory or not. Evaluated in ten ECG gated reconstructions per patient.	Yes

Table 2.2: *An overview of the quantitatively evaluated CTA coronary artery segmentation and centerline extraction algorithms. With 'centerline' and 'reference' we respectively denote the (semi-)automatically extracted centerline and the manually annotated centerline. The column 'Time' indicates if information is provided about the computational time of the algorithm. 'Method eval.' indicates that the article evaluates an existing technique and that no new technique has been proposed.*

Article	Patients / observers	Vessels	Used evaluation measures and details	Time	Method eval.
Bulow et al. [27]	9/1	3-5 vessels	**Overlap**: Percentage reference points having a centerline point within 2 mm.	No	
Busch et al. [28]	23/2	Complete tree	**Stenoses grading**: Compared to human performance with CCA as ground truth.	No	×
Dewey et al. [37]	35/1	3 vessels	**Length difference**: Difference between reference length and centerline length. **Stenoses grading**: Compared to human performance with CCA as ground truth.	Yes	×
Khan et al. [75]	50/1	3 vessels	**Stenoses grading**: Compared to human performance with CCA as ground truth.	No	×
Larralde et al. [80]	6/1	Complete tree	**Stenoses grading and calcium detection**: Compared to human performance.	Yes	
Lesage et al. [84]	19/1	3 vessels	Same as Metz et al. [104]	Yes	
Li et al. [88]	5/1	Complete tree	**Segmentation**: Voxel-wise similarity indices.	No	
Marquering et al. [101]	1/1	LAD	**Accuracy**: Distance from centerline to reference standard.	Yes	
Metz et al. [104]	6/3	3 vessels	**Overlap**: Segments on the reference standard and centerline are marked as true positives, false positives or false negatives. This scoring was used to construct similarity indices. **Accuracy**: Average distance to the reference standard for true positive sections.	No	
Olabarriaga et al. [121]	5/1	3 vessels	**Accuracy**: Mean distance from the centerline to the reference.	No	
Wesarg et al. [178]	10/1	3 vessels	**Calcium detection**: Performance compared to human performance.	No	×
Yang et al. [186]	2/1	3 vessels	**Overlap**: Percentage of the reference standard detected. **Segmentation**: Average distance to contours.	No	

2.4.1 Cardiac CTA data

The CTA data was acquired in the Erasmus MC, University Medical Center Rotterdam, The Netherlands. Thirty-two datasets were randomly selected from a series of patients who underwent a cardiac CTA examination between June 2005 and June 2006. Twenty datasets were acquired with a 64-slice CT scanner and twelve datasets with a dual-source CT scanner (Sensation 64 and Somatom Definition, Siemens Medical Solutions, Forchheim, Germany).

A tube voltage of 120 kV was used for both scanners. All datasets were acquired with ECG-pulsing [180]. The maximum current (625 mA for the dual-source scanner and 900 mA for the 64-slice scanner) was used in the window from 25% to 70% of the R-R-interval and outside this window the tube current was reduced to 20% of the maximum current.

Both scanners operated with a detector width of 0.6 mm. The image data was acquired with a table feed of 3.8 mm per rotation (64-slice datasets) or 3.8 mm to 10 mm, individually adapted to the patient's heart rate (dual-source datasets).

Diastolic reconstructions were used, with reconstruction intervals varying from 250 ms to 400 ms before the R-peak. Three datasets were reconstructed using a sharp (B46f) kernel, all others were reconstructed using a medium-to-smooth (B30f) kernel. The mean voxel size of the datasets is $0.32 \times 0.32 \times 0.4\text{mm}^3$.

2.4.1.1 Training and test datasets

To ensure representative training and test sets, the image quality of and presence of calcium in each dataset was visually assessed by a radiologist with three years experience in cardiac CT.

Image quality was scored as poor (defined as presence of image-degrading artifacts and evaluation only possible with low confidence), moderate (presence of artifacts but evaluation possible with moderate confidence) or good (absence of any image-degrading artifacts related to motion and noise). Presence of calcium was scored as absent, modest or severe. Based on these scorings the data was distributed equally over a group of 8 and a group of 24 datasets. The patient and scan parameters were assessed by the radiologist to be representative for clinical practice. Table 2.3 and 2.4 describe the distribution of respectively the image quality and calcium scores in the datasets.

The first group of 8 datasets can be used for training and the other 24 datasets are used for performance assessment of the algorithms. All the thirty-two cardiac CTA datasets and the corresponding reference standard centerlines for the training data are made publicly available.

2.4.2 Reference standard

In this work we define the centerline of a coronary artery in a CTA scan as the curve that passes through the center of gravity of the lumen in each cross-section.

Table 2.3: *Image quality of the training and test datasets.*

	Poor	Moderate	Good	Total
Training	2	3	3	8
Testing	4	8	12	24

Table 2.4: *Presence of calcium in the training and test datasets.*

	Low	Moderate	Severe	Total
Training	3	4	1	8
Testing	9	12	3	24

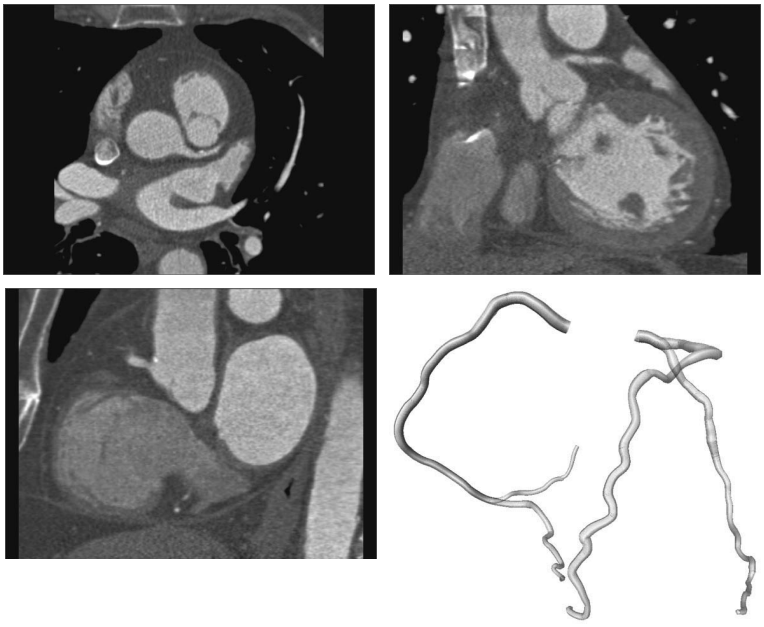

Figure 2.1: *An example of the data with corresponding reference standard. Top-left: axial view of data. Top-right: coronal view. Bottom-left: sagittal view. Bottom-right: a 3D rendering of the reference standard.*

We define the start point of a centerline as the center of the coronary ostium (i.e. the point where the coronary artery originates from the aorta), and the end point as the most distal point where the artery is still distinguishable from the background. The centerline is smoothly interpolated if the artery is partly indistinguishable from the background, e.g. in case of a total occlusion or imaging artifacts.

This definition was used by three trained observers to annotate centerlines in the selected cardiac CTA datasets. Four vessels were selected for annotation by one of the observers in all 32 datasets, yielding $32 \times 4 = 128$ selected vessels. The first three vessels were always the right coronary artery (RCA), left anterior descending artery (LAD), and left circumflex artery (LCX). The fourth vessel was selected from the large side branches of these main coronary arteries and the selection was as follows: first diagonal branch (14x), second diagonal branch (6x), optional diagonal coronary artery (6x), first obtuse marginal branch (2x), posterior descending artery (2x), and acute marginal artery (2x). This observer annotated for all the four selected vessels points close to the selected vessels. These points (denoted with 'point A') unambiguously define the vessels, i.e. the vessel of interest is the vessel closest to the point and no side-branches can be observed after this point.

After the annotation of these 128 points, the three observers used these points to independently annotate the centerlines of the same four vessels in the 32 datasets. The observers also specified the radius of the lumen at least every 5 mm, where the radius was chosen such that the enclosed area of the annotated circle matched the area of the lumen. The radius was specified after the complete central lumen line was annotated (see Figure 2.4).

The paths of the three observers were combined to one centerline per vessel using a Mean Shift algorithm for open curves: The centerlines are averaged while taking into account the possibly spatially varying accuracy of the observers by iteratively estimating the reference standard and the accuracy of the observers. Each point of the resulting reference standard is a weighted average of the neighboring observer centerline points, with weights corresponding to the locally estimated accuracy of the observers [165].

After creating this first weighted average, a consensus centerline was created with the following procedure: The observers compared their centerlines with the average centerline to detect and subsequently correct any possible annotation errors. This comparison was performed utilizing curved planar reformatted images displaying the annotated centerline color-coded with the distance to the reference standard and vice-versa (see Figure 2.2). The three observers needed in total approximately 300 hours for the complete annotation and correction process.

After the correction step the centerlines were used to create the reference standard, using the same Mean Shift algorithm. Note that the uncorrected centerlines were used to calculate the inter-observer variability and agreement measures (see section 2.4.5).

The points where for the first time the centerlines of two observers lie within

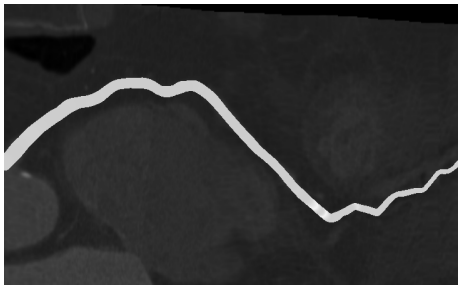

Figure 2.2 (see page 145 for a color-version): *An example of a color-coded curved planar reformatted images used to detect possible annotation errors.*

the radius of the reference standard when traversing over this centerline from respectively the start to the end or vice versa were selected as the start- and end point of the reference standard. Because the observers used the abovementioned centerline definition it is assumed that the resulting start points of the reference standard centerlines lie within the coronary ostium.

The corrected centerlines contained on average 44 points and the average distance between two successive annotated points was 3.1 mm. The 128 resulting reference standard centerlines were on average 138 mm (std. dev. 41 mm, min. 34 mm, max. 249 mm) long.

The radius of the reference standard was based on the radii annotated by the observers and a point-to-point correspondence between the reference standard and the three annotated centerlines. The reference standard centerline and the corrected observer centerlines were first resampled equidistantly using a sampling distance of 0.03 mm. Dijkstra's graph searching algorithm was then used to associate each point on the reference standard with one or more points on each annotated centerline and vice versa. Using this correspondence, the radius at each point of the reference standard was determined by averaging the radius of all the connected points on the three annotated centerlines (see also Figure 2.3 and Figure 2.4). An example of annotated data with corresponding reference standard is shown in Figure 2.1. Details about the connectivity algorithm are given in section 2.4.3.

2.4.3 Correspondence between centerlines

All the evaluation measures are based on a point-to-point correspondence between the reference standard and the evaluated centerline. This section explains the mechanism for determining this correspondence.

Before the correspondence is determined the centerlines are first sampled equidistantly using a sampling distance of 0.03 mm, enabling an accurate comparison. The evaluated centerline is then clipped with a disc that is positioned at the start of

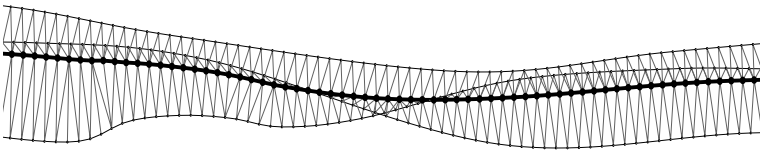

Figure 2.3: *An illustrative example of the Mean Shift algorithm showing the annotations of the three observers as a thin black line, the resulting average as a thick black line, and the correspondence that are used during the last Mean Shift iteration in light-gray.*

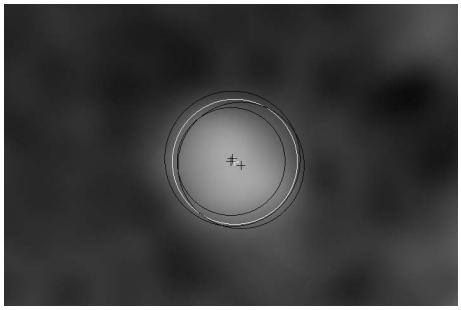

Figure 2.4: *An example of the annotations of the three observers in black and the resulting reference standard in white. The crosses indicate the centers and the circles indicate the radii.*

the reference standard centerline (i.e. in or very close to the coronary ostium). The centerlines are clipped because we define the start point of a coronary centerline at the coronary ostium and because for a variety of applications the centerline can start somewhere in the aorta. The radius of the disc is twice the annotated vessel radius and the disc normal is the tangential direction at the beginning of the reference standard centerline. Every point before the first intersection of a centerline and this disc is not taken into account during evaluation.

The correspondence is then determined by finding the minimum of the sum of the Euclidean lengths of all point-point connections that are connecting the two centerlines over all valid correspondences. A valid correspondence for centerline I, consisting of an ordered set of points p_i ($0 \leq i < n$, p_0 is the most proximal point of the centerline), and centerline II, consisting of an ordered set of points q_j ($0 \leq j < m$, q_0 is the most proximal point of the centerline), is defined as the ordered set of connections $C = \{c_0, \ldots, c_{n+m-1}\}$, where c_k is a tuple $[p_a, q_b]$ that represents a connection from p_a to q_b, which satisfies the following conditions:

- The first connection c_0 connects the start points: $c_0 = [p_0, q_0]$.
- The last connection c_{n+m-1} connects the end points: $c_{n+m-1} = [p_{n-1}, q_{m-1}]$.

- If connection $c_k = [p_a, q_b]$ then connection c_{k+1} equals either $[p_{a+1}, q_b]$ or $[p_a, q_{b+1}]$.

These conditions guarantee that each point of centerline I is connected to at least one point of centerline II and vice versa.

Dijkstra's graph search algorithm is used on a matrix with connection lengths to determine the minimum Euclidean length correspondence. See Figure 2.3 for an example of a resulting correspondence.

2.4.4 Evaluation measures

Coronary artery centerline extraction may be used for different applications, and thus different evaluation measures may apply. We account for this by employing a number of evaluation measures. With these measures we discern between extraction capability and extraction accuracy. Accuracy can only be evaluated when extraction succeeded; in case of a tracking failure the magnitude of the distance to the reference centerline is no longer relevant and should not be included in the accuracy measure.

2.4.4.1 Definition of true positive, false positive and false negative points

All the evaluation measures are based on a labeling of points on the centerlines as true positive, false negative or false positive. This labeling, in its turn, is based on a correspondence between the points of the reference standard centerline and the points of the centerline to be evaluated. The correspondence is determined with the algorithm explained in section 2.4.3.

A point of the reference standard is marked as true positive $\mathrm{TPR_{ov}}$ if the distance to at least one of the connected points on the evaluated centerline is less than the annotated radius and false negative $\mathrm{FN_{ov}}$ otherwise.

A point on the centerline to be evaluated is marked as true positive $\mathrm{TPM_{ov}}$ if there is at least one connected point on the reference standard at a distance less than the radius defined at that reference point, and it is marked as false positive $\mathrm{FP_{ov}}$ otherwise. With $\|.\|$ we denote the cardinality of a set of points, e.g. $\|\mathrm{TPR_{ov}}\|$ denotes the number of reference points marked true positive. See also Figure 2.5 for a schematic explanation of these terms and the terms mentioned in the next section.

2.4.4.2 Overlap measures

Three different overlap measures are used in our evaluation framework.

Overlap (OV) represents the ability to track the complete vessel annotated by the human observers and this measure is similar to the well-known Dice coefficient. It is defined as:

$$\mathrm{OV} = \frac{\|\mathrm{TPM_{ov}}\| + \|\mathrm{TPR_{ov}}\|}{\|\mathrm{TPM_{ov}}\| + \|\mathrm{TPR_{ov}}\| + \|\mathrm{FN_{ov}}\| + \|\mathrm{FP_{ov}}\|}.$$

Overlap until first error (OF) determines how much of a coronary artery has been extracted before making an error. This measure can for example be of interest for image guided intra-vascular interventions in which guide wires are advanced based on pre-operatively extracted coronary geometry [127]. The measure is defined as the ratio of the number of true positive points on the reference before the first error (TPR_{of}) and the total number of reference points (TPR_{of} + FN_{of}):

$$OF = \frac{\|TPR_{of}\|}{\|TPR_{of}\| + \|FN_{of}\|}.$$

The first error is defined as the first FN_{ov} point when traversing from the start of the reference standard to its end while ignoring false negative points in the first 5 mm of the reference standard. Errors in the first 5 mm are not taken into account because of the strictness of this measure and the fact that the beginning of a coronary artery centerline is sometimes difficult to define and for some applications not of critical importance. The threshold of five millimeters is equal to the average diameter annotated at the beginning of all the reference standard centerlines.

Overlap with the clinically relevant part of the vessel (OT) gives an indication of how well the method is able to track the section of the vessel that is assumed to be clinically relevant. Vessel segments with a diameter of 1.5 mm or larger, or vessel segments that are distally from segments with a diameter of 1.5 mm or larger are assumed to be clinically relevant [87, 133]. The point closest to the end of the reference standard with a radius larger than or equal to 0.75 mm is determined. Only points on the reference standard between this point and the start of the reference standard and points on the (semi-)automatic centerline connected to these reference points are used when defining the true positives (TPM_{ot} and TPR_{ot}), false negatives (FN_{ot}) and false positives (FP_{ot}). The OT measure is calculated as follows:

$$OT = \frac{\|TPM_{ot}\| + \|TPR_{ot}\|}{\|TPM_{ot}\| + \|TPR_{ot}\| + \|FN_{ot}\| + \|FP_{ot}\|}.$$

2.4.4.3 Accuracy measure

In order to discern between tracking ability and tracking accuracy we only evaluate the accuracy within sections where tracking succeeded.

Average inside (AI) is the average distance of all the connections between the reference standard and the automatic centerline given that the connections have a length smaller than the annotated radius at the connected reference point. The measure represents the accuracy of centerline extraction, provided that the evaluated centerline is inside the vessel.

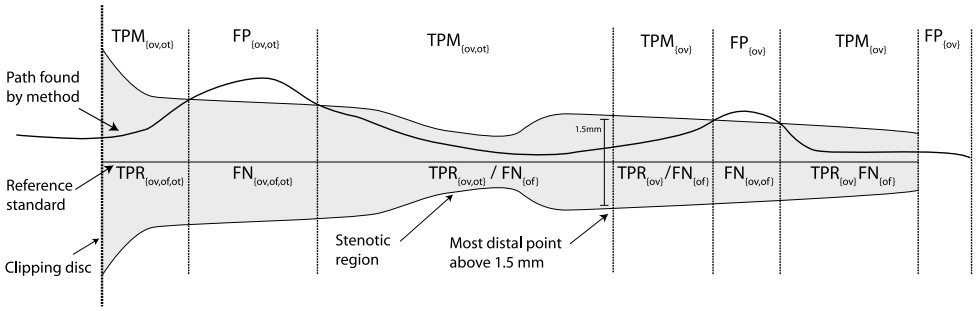

Figure 2.5: *An illustration of the terms used in the evaluation measures (see section 2.4.4). The reference standard with annotated radius is depicted in gray. The terms on top of the figure are assigned to points on the centerline found by the evaluated method. The terms below the reference standard line are assigned to points on the reference standard.*

2.4.5 Observer performance and scores

Each of the evaluation measures is related to the performance of the observers by a relative score. A score of 100 points implies that the result of the method is perfect, 50 points implies that the performance of the method is similar to the performance of the observers, and 0 points implies a complete failure. This section explains how the observer performance is quantified for each of the four evaluation measures and how scores are created from the evaluation measures by relating the measures to the observer performance.

2.4.5.1 Overlap measures

The inter-observer agreement for the overlap measures is calculated by comparing the uncorrected paths with the reference standard. The three overlap measures (OV, OF, OT) were calculated for each uncorrected path and the true positives, false positives and false negatives for each observer were combined into inter-observer agreement measures per centerline as follows:

$$\mathrm{OV_{ag}} = \frac{\sum \left(\|\mathrm{TPR}^i_{\mathrm{ov}}\| + \|\mathrm{TPM}^i_{\mathrm{ov}}\| \right)}{\sum \left(\|\mathrm{TPR}^i_{\mathrm{ov}}\| + \|\mathrm{TPM}^i_{\mathrm{a}}\| + \|\mathrm{FP}^i_{\mathrm{ov}}\| + \|\mathrm{FN}^i_{\mathrm{ov}}\| \right)}$$

$$\mathrm{OF_{ag}} = \frac{\sum \|\mathrm{TPR}^i_{\mathrm{of}}\|}{\sum \left(\|\mathrm{TPR}^i_{\mathrm{of}}\| + \|\mathrm{FN}^i_{\mathrm{of}}\| \right)}$$

$$\mathrm{OT_{ag}} = \frac{\sum \left(\|\mathrm{TPR}^i_{\mathrm{ot}}\| + \|\mathrm{TPM}^i_{\mathrm{ot}}\| \right)}{\sum \left(\|\mathrm{TPR}^i_{\mathrm{ot}}\| + \|\mathrm{TPM}^i_{\mathrm{ot}}\| + \|\mathrm{FP}^i_{\mathrm{ot}}\| + \|\mathrm{FN}^i_{\mathrm{ot}}\| \right)},$$

where $i = \{0, 1, 2\}$ indicates the observer.

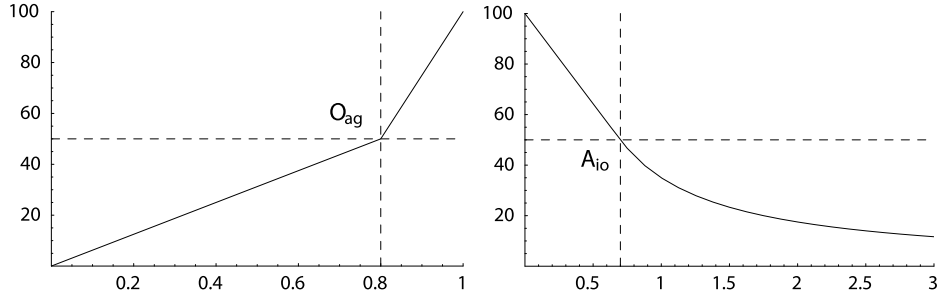

(a) Overlap to score (b) Accuracy to score

Figure 2.6: *(a) shows an example of how overlap measures are transformed into scores. (b) shows this transformation for the accuracy measures.*

After calculating the inter-observer agreement measures, the performance of the method is scored. For methods that perform better than the observers the OV, OF, and OT measures are converted to scores by linearly interpolating between 100 and 50 points, respectively corresponding to an overlap of 1.0 and an overlap similar to the inter-observer agreement value. If the method performs worse than the inter-observer agreement the score is obtained by linearly interpolating between 50 and 0 points, with 0 points corresponding to an overlap of 0.0:

$$\text{Score}_{\text{O}} = \begin{cases} (\text{O}_{\text{m}}/\text{O}_{\text{ag}}) * 50 & \text{O}_{\text{m}} \leq \text{O}_{\text{ag}} \\ 50 + 50 * \frac{\text{O}_{\text{m}} - \text{O}_{\text{ag}}}{1 - \text{O}_{\text{ag}}} & \text{O}_{\text{m}} > \text{O}_{\text{ag}}, \end{cases}$$

where O_{m} and O_{ag} define the OV, OF, or OT performance of respectively the method and the observer. An example of this conversion is shown in Figure 2.6(a).

2.4.5.2 Accuracy measures

The inter-observer variability for the accuracy measure AI is defined at every point of the reference standard as the expected error that an observer locally makes while annotating the centerline. It is determined at each point as the root mean squared distance between the uncorrected annotated centerline and the reference standard:

$$\text{A}_{\text{io}}(x) = \sqrt{1/n \sum (d(p(x), p_i))^2},$$

where $n = 3$ (three observers), and $d(p(x), p_i)$ is the average distance from point $p(x)$ on the reference standard to the connected points on the centerline annotated by observer i.

The extraction accuracy of the method is related per connection to the inter-observer variability. A connection is worth 100 points if the distance to the reference standard is 0 mm and it is worth 50 points if the distance is equal to the

inter-observer variability at that point. Methods that perform worse than the inter-observer variability get a decreasing amount of points if the distance increases. They are rewarded per connection 50 points times the fraction of the inter-observer variability and the method accuracy:

$$\text{Score}_{\text{A}}(x) = \begin{cases} 100 - 50(\text{A}_{\text{m}}(x)/\text{A}_{\text{io}}(x)) & \text{A}_{\text{m}}(x) \leq \text{A}_{\text{io}}(x) \\ (\text{A}_{\text{io}}(x)/\text{A}_{\text{m}}(x)) * 50 & \text{A}_{\text{m}}(x) > \text{A}_{\text{io}}(x), \end{cases}$$

where $\text{A}_{\text{m}}(x)$ and $\text{A}_{\text{io}}(x)$ define the distance from the method centerline to the reference centerline and the inter-observer accuracy variability at point x. An example of this conversion is shown in Figure 2.6(b).

The average score over all connections that connect TPR and TPM points yields the AI observer performance score. Because the average accuracy score is a non-linear combination of all the distances, it can happen that a method has a lower average accuracy in millimeters and a higher score in points than another method, or vice versa.

Note that because the reference standard is constructed from the observer centerlines, the reference standard is slightly biased towards the observer centerlines, and thus a method that performs similar as an observer according to the scores probably performs slightly better. Although more sophisticated methods for calculating the observer performance and scores would have been possible, we decided because of simplicity and understandability for the approach explained above.

2.4.6 Ranking the algorithms

In order to rank the different coronary artery centerline extraction algorithms the evaluation measures have to be combined. We do this by ranking the resulting scores of all the methods for each measure and vessel. Each method receives for each vessel and measure a rank ranging from 1 (best) to the number of participating methods (worst). A user of the evaluation framework can manually mark a vessel as failed. In that case the method will be ranked last for the flagged vessel and the absolute measures and scores for this vessel will not be taken into account in any of the statistics.

The tracking capability of a method is defined as the average of all the 3(overlap measures) $\times$ 96 (vessels) $=$ 288 related ranks. The average of all the 96 accuracy measure ranks defines the tracking accuracy of each method. The average overlap rank and the accuracy rank are averaged to obtain the overall quality of each of the methods and the method with the best (i.e. lowest) average rank is assumed to be the best.

2.5 Algorithm categories

We discern three different categories of coronary artery centerline extraction algorithms: automatic extraction methods, methods with minimal user interaction and interactive extraction methods.

2.5.1 Category 1: automatic extraction

Automatic extraction methods find the centerlines of coronary arteries without user interaction. In order to evaluate the performance of automatic coronary artery centerline extraction, two points per vessel are provided to extract the coronary artery of interest:

- Point A: a point inside the distal part of the vessel; this point unambiguously defines the vessel to be tracked;
- Point B: a point approximately 3 cm (measured along the centerline) distal of the start point of the centerline.

Point A should be used for selecting the appropriate centerline. If the automatic extraction result does not contain centerlines near point A, point B can be used. Point A and B are only meant for selecting the right centerline and it is not allowed to use them as input for the extraction algorithm.

2.5.2 Category 2: extraction with minimal user interaction

Extraction methods with minimal user interaction are allowed to use one point per vessel as input for the algorithm. This can be either one of the following points:

- Point A or B, as defined above;
- Point S: the start point of the centerline;
- Point E: the end point of the centerline;
- Point U: any manually defined point.

Points A, B, S and E are provided with the data. Furthermore, in case the method obtains a vessel tree from the initial point, point A or B may be used after the centerline determination to select the appropriate centerline.

2.5.3 Category 3: interactive extraction

All methods that require more user-interaction than one point per vessel as input are part of category 3. Methods can use e.g. both points S and E from category 2, a series of manually clicked positions, or one point and a user-defined threshold.

2.6 Web-based evaluation framework

The proposed framework for the evaluation of CTA coronary artery centerline extraction algorithms is made publicly available through a web-based interface at (*http://coronary.bigr.nl*). The thirty-two cardiac CTA datasets, and the corresponding reference standard centerlines for the training data, are available for download for anyone who wishes to validate their algorithm. Extracted centerlines can be submitted and the obtained results can be used in a publication. Furthermore, the website provides several tools to inspect the results and compare the algorithms.

2.7 MICCAI 2008 workshop

This study started with the workshop '3D Segmentation in the Clinic: A Grand Challenge II' at the 11th International Conference on Medical Image Computing and Computer Assisted Intervention (MICCAI) in September 2008 [105]. Approximately 100 authors of related publications, and the major medical imaging companies, were invited to submit their results on the 24 test datasets. Fifty-three groups showed their interest by registering for the challenge, 36 teams downloaded the training and test data, and 13 teams submitted results: five fully automatic methods, three minimally interactive methods, and five interactive methods. A brief description of the thirteen methods is given below.

During the workshop we used two additional measures: the average distance of all the connections (AD) and the average distance of all the connections to the clinical relevant part of the vessel (AT). In retrospect we found that these accuracy measures were too much biased towards methods with high overlap and therefore we do not use them anymore in the evaluation framework. This resulted in a slightly different ranking than the ranking published during the MICCAI workshop [105]. Please note that the two measures that were removed are still calculated for all the evaluated methods and they can be inspected using the web-based interface.

2.7.1 Fully automatic methods

- AutoCoronaryTree [52, 157]: The full centerline tree of the coronary arteries is extracted via a multi-scale medialness-based vessel tree extraction algorithm which starts a tracking process from the ostia locations until all coronary branches are reached.
- CocomoBeach [76]: This method starts by segmenting the ascending aorta and the heart. Candidate coronary regions are obtained using connected component analysis and the masking of large structures. Using these components a region growing scheme, starting in the aorta, segments the complete tree. Finally, centerlines within the pre-segmented tree are obtained using the WaveProp [101] method.

- DepthFirstModelFit [188]: Coronary centerline extraction is accomplished by fitting models of shape and appearance. A large-scale model of the complete heart in combination with symmetry features is used for detecting coronary artery seeds. To fully extract the coronary artery tree, two small-scale cylinder-like models are matched via depth-first search.

- GVFTube'n'Linkage [14]: This method uses a Gradient Vector Flow [184] based tube detection procedure for identification of vessels surrounded by arbitrary tissues [13, 15]. Vessel centerlines are extracted using ridge-traversal and linked to form complete tree structures. For selection of coronary arteries gray value information and centerline length are used.

- VirtualContrast [168]: This method segments the coronary arteries based on the connectivity of the contrast agent in the vessel lumen, using a competing fuzzy connectedness tree algorithm [167]. Automatic rib cage removal and ascending aorta tracing are included to initialize the segmentation. Centerline extraction is based on the skeletonization of the tree structure.

2.7.2 Semi automatic methods

- AxialSymmetry [38]: This method finds a minimum cost path connecting the aorta to a user supplied distal endpoint. Firstly, the aorta surface is extracted. Then, a two-stage Hough-like election scheme detects the high axial symmetry points in the image. Via these, a sparse graph is constructed. This graph is used to determine the optimal path connecting the user supplied seed point and the aorta.

- CoronaryTreeMorphoRec [32]: This method generates the coronary tree iteratively from point S. Pre-processing steps are performed in order to segment the aorta, remove unwanted structures in the background and detect calcium. Centerline points are chosen in each iteration depending on the previous vessel direction and a local gray scale morphological 3D reconstruction.

- KnowledgeBasedMinPath [79]: For each voxel, the probability of belonging to a coronary vessel is estimated from a feature space and a vesselness measure is used to obtain a cost function. The vessel starting point is obtained automatically, while the end point is provided by the user. Finally, the centerline is obtained as the minimal cost path between both points.

2.7.3 Interactive methods

- 3DInteractiveTrack [190]: This method calculates a local cost for each voxel based on eigenvalue analysis of the Hessian matrix. When a user selects a point, the method calculates the cost linking this point to all other voxels.

If a user then moves to any voxel, the path with minimum overall cost is displayed. The user is able to inspect and modify the tracking to improve performance.

- ElasticModel [61]. After manual selection of a background-intensity threshold and one point per vessel, centerline points are added by prediction and refinement. Prediction uses the local vessel orientation, estimated by eigen-analysis of the inertia matrix. Refinement uses centroid information and is restricted by continuity and smoothness constraints of the model [60].

- MHT [49]: Vessel branches are in this method found using a Multiple Hypothesis Tracking (MHT) framework. A feature of the MHT framework is that it can traverse difficult passages by evaluating several hypothetical paths. A minimal path algorithm based on Fast Marching is used to bridge gaps where the MHT terminates prematurely.

- Tracer [155]: This method finds the set of core points (centers of intensity plateaus in 2D slices) that concentrate near vessel centerlines. A weighted graph is formed by connecting nearby core points. Low weights are given to edges of the graph that are likely to follow a vessel. The output is the shortest path connecting point S and point E.

- TwoPointMinCost [106]: This method finds a minimum cost path between point S and point E using Dijkstra's algorithm. The cost to travel through a voxel is based on Gaussian error functions of the image intensity and a Hessian-based vesselness measure [47], calculated on a single scale.

2.8 Results

The results of the thirteen methods are shown in Table 2.5, 2.6, and 2.7. Table 2.6 shows the results for the three overlap measures, Table 2.7 shows the accuracy measures, and Table 2.5 shows the final ranking, the approximate processing time, and amount of user-interaction that is required to extract the four vessels. In total 10 extractions ($< 1\%$) where marked as failed (see section 2.4.6).

We believe that the final ranking in Table 2.5 gives a good indication of the relative performance of the different methods, but one should be careful to judge the methods on their final rank. A method ranked first does not have to be the method of choice for a specific application. For example, if a completely automatic approximate extraction of the arteries is needed one could choose GVFTube'n'Linkage [14] because it has the highest overlap with the reference standard (best OV result). But if one wishes to have a more accurate automatic extraction of the proximal part of the coronaries the results point you toward DepthFirstModelFit [188] because this

method is highly ranked in the OF measure and is ranked first in the automatic methods category with the AI measure.

The results show that on average the interactive methods perform better on the overlap measures than the automatic methods (average rank of 6.30 vs. 7.09) and vice versa for the accuracy measures (8.00 vs. 6.25). The better overlap performance of the interactive methods can possibly be explained by the fact that the interactive methods use the start- and/or end point of the vessel. Moreover, in two cases (MHT [49] and 3DInteractiveTrack [190]) additional manually annotated points are used, which can help the method to bridge difficult regions.

When vessels are correctly extracted, the majority of the methods are accurate to within the image voxel size (AI < 0.4mm). The two methods that use a tubular shape model (MHT [49] and DepthFirstModelFit [188]) have the highest accuracy, followed by the multi-scale medialness-based AutoCoronaryTree [52, 157] method and the CocomoBeach [76] method.

Overall it can be observed that some of the methods are highly accurate and some have great extraction capability (i.e. high overlap). Combining a fully automatic method with high overlap (e.g. GVFTube'n'Linkage [14]) and a, not necessarily fully automatic, method with high accuracy (e.g. MHT [49]) may result in an fully automatic method with high overlap and high accuracy.

2.8.1 Results categorized

Separate rankings are made for each group of datasets with corresponding image quality and calcium rating to determine if the image quality or the amount of calcium has influence on the rankings.
Separate rankings are also made for each of the four vessel types. These rankings are presented in Table 2.8. It can be seen that some of the methods perform relatively worse when the image quality is poor or an extensive amount of calcium is present (e.g. CocomoBeach [76] and DepthFirstModelFit [188]) and vice versa (e.g. KnowledgeBasedMinPath [79] and VirtualContrast [168]).

Table 2.8 also shows that on average the automatic methods perform relatively worse for datasets with poor image quality (i.e. the ranks of the automatic methods in the P-column are on average higher compared to the ranks in the M- and G-column). This is also true for the extraction of the LCX centerlines. Both effects can possibly be explained by the fact that centerline extraction from poor image quality datasets and centerline extraction of the (on average relatively thinner) LCX is more difficult to automate.

2.8.2 Algorithm performance with respect to ostium distance

For a number of coronary artery centerline extraction applications it is not important to extract the whole coronary artery; only extraction up to a certain distance from the coronary ostium is required (see e.g. [65, 170]).

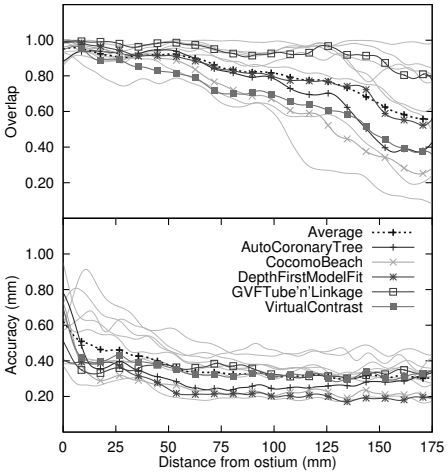

(a) Fully automatic coronary artery centerline extraction methods

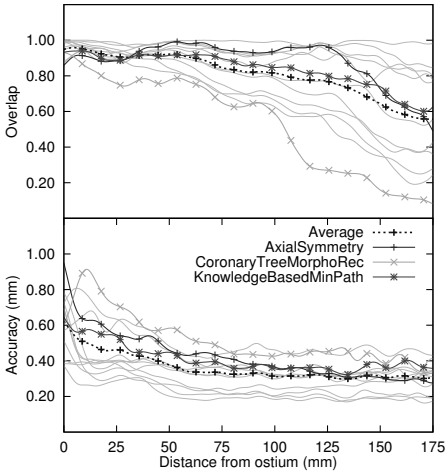

(b) Semi automatic coronary artery centerline extraction methods

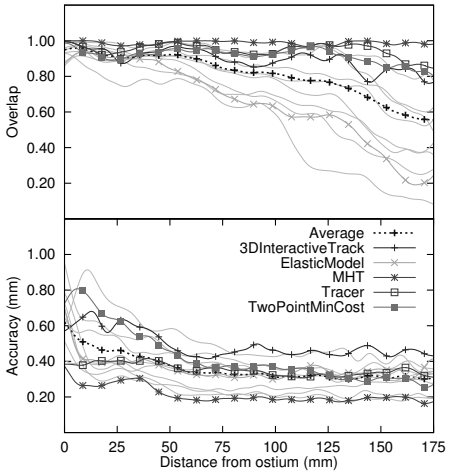

(c) Interactive coronary artery centerline extraction methods

Figure 2.7 (see page 146 for a color-version): *The algorithm performance of each method with respect to the distance from the ostium averaged over all 96 evaluated vessels over the first 175mm. Overlap: the fraction of points on the reference standard marked as true positive. Accuracy: the average distance to the centerline if the point is marked true positive. Each of the three graphs shows in light-gray the results of all the thirteen evaluated methods and in color the results of the respective algorithm category. The graphs also show in black the average accuracy and overlap for all thirteen evaluated methods.*

Table 2.5: *The overall ranking of the thirteen evaluated methods. The average overlap rank, accuracy rank and the average of these two is shown together with an indication of the computation time and the required user-interaction.*

Method	Cat.	Avg. Ov. rank	Avg. Acc. rank	Avg. rank	Computation time	User-interaction
MHT [49]	3	2.07	1.58	1.83	6 min.	2 to 5 pts.
Tracer [155]	3	4.21	2.52	3.37	30 min.	Pt. S and pt. E
DepthFirstModelFit [188]	1	6.17	3.33	4.75	4-8 min.	
KnowledgeBasedMinPath [79]	2	4.31	8.36	6.34	7 hours	Pt. E
AutoCoronaryTree [157]	1	7.69	5.18	6.44	$<$ 30 seconds	
GVFTube'n'Linkage [14]	1	5.39	8.02	6.71	10 min.	
CocomoBeach [76]	1	8.56	5.04	6.80	70 seconds	
TwoPointMinCost [106]	3	5.30	8.80	7.05	12 min.	Pt. S and pt. E
VirtualContrast [168]	1	8.71	7.74	8.23	5 min.	
AxialSymmetry [38]	2	6.95	9.60	8.28	5 min.	Point E
ElasticModel [61]	3	9.05	8.29	8.67	2-6 min.	Global thresh. + 1 pt. per axis
3DInteractiveTrack [190]	3	7.52	10.91	9.22	3-6 min.	3 to 10 pts.
CoronaryTreeMorphoRec [32]	2	10.42	11.59	11.01	30 min.	Pt. S

In order to evaluate the performance of the methods with respect to the distance from the ostium, charts are generated that demonstrate the average performance over all 96 evaluated centerlines for each of the methods at a specific distance from the ostium (measured along the reference standard). Figure 2.7(a) shows these results for the automatic methods, Figure 2.7(b) shows the results for the methods with minimal user-interaction, and Figure 2.7(c) shows the results for the semi automatic methods.

The graphs show that all the evaluated methods are better able to extract the proximal part of the coronaries than the more distal part of the vessels. Moreover, they show that after approximately 5 cm the accuracy of almost all the methods is relatively constant. Furthermore, the graphs again demonstrate the fact that the automatic methods are on average more accurate than the semi automatic or interactive methods.

2.8.3 More statistics available online

Space limitations prevent us to incorporate more statistics here, but the on-line evaluation framework (*http://coronary.bigr.nl*) provides the possibilities to rank the methods based on different measures or scores, create statistics on a subset of the data and create overview tables for specific measures, categorized on image quality or score. It is for example possible to create Table 2.5, 2.6, and 2.7 for a specific subset of the data or to create Table 2.8 with a measure or score of choice, instead of the overall ranks. The website also contains the most recent version of the results. The on-line results are different from the results reported in this chapter because of recent submissions and improvements in implementation of the different methods.

Table 2.6: *The resulting overlap measures for the thirteen evaluated methods. The average overlap, score and rank is shown for each of the three overlap measures.*

Method	Cat.	OV			OF			OT		
		%	score	rank	%	score	rank	%	score	rank
MHT [49]	3	98.5	84.0	1.74	83.1	72.8	2.64	98.7	84.5	1.83
Tracer [155]	3	95.1	71.0	3.60	63.5	52.0	5.22	95.5	70.2	3.81
DepthFirstModelFit [188]	1	84.7	48.6	7.29	65.3	49.2	5.32	87.0	60.1	5.90
KnowledgeBasedMinPath[79]	2	88.0	67.4	4.46	74.2	61.1	4.27	88.5	70.0	4.21
AutoCoronaryTree [157]	1	84.7	46.5	8.13	59.5	36.1	7.26	86.2	50.3	7.69
GVFTube'n'Linkage [14]	1	92.7	52.3	6.20	71.9	51.4	5.32	95.3	67.0	4.66
CocomoBeach [76]	1	78.8	42.5	9.34	64.4	40.0	7.39	81.2	46.9	8.96
TwoPointMinCost [106]	3	91.9	64.5	4.70	56.4	45.6	6.22	92.5	64.5	4.97
VirtualContrast [168]	1	75.6	39.2	9.74	56.1	34.5	7.74	78.7	45.6	8.64
AxialSymmetry [38]	2	90.8	56.8	6.17	48.9	35.6	7.96	91.7	55.9	6.71
ElasticModel [61]	3	77.0	40.5	9.60	52.1	31.5	8.46	79.0	45.3	9.09
3DInteractiveTrack [190]	3	89.6	51.1	7.04	49.9	30.5	8.36	90.6	52.4	7.15
CoronaryTreeMorphoRec [32]	2	67.0	34.5	11.00	36.3	20.5	9.53	69.1	36.7	10.74

Table 2.7: *The accuracy of the thirteen evaluated methods. The average distance, score and rank of each is shown for the accuracy when inside (AI) measure.*

Method	Cat.	AI		
		mm	score	rank
MHT [49]	3	0.23	47.9	1.58
Tracer [155]	3	0.26	44.4	2.52
DepthFirstModelFit [188]	1	0.28	41.9	3.33
KnowledgeBasedMinPath [79]	2	0.39	29.2	8.36
AutoCoronaryTree [157]	1	0.34	35.3	5.18
GVFTube'n'Linkage [14]	1	0.37	29.8	8.02
CocomoBeach [76]	1	0.29	37.7	5.04
TwoPointMinCost [106]	3	0.46	28.0	8.80
VirtualContrast [168]	1	0.39	30.6	7.74
AxialSymmetry [38]	2	0.46	26.4	9.60
ElasticModel [61]	3	0.40	29.3	8.29
3DInteractiveTrack [190]	3	0.51	24.2	10.91
CoronaryTreeMorphoRec [32]	2	0.59	20.7	11.59

Table 2.8: *Ranks per image quality; poor(P), moderate(M) or good(G), calcium score; low(L), moderate(M) or severe(S) and vessel type. The numbers indicate the rank of each team if only the specified datasets or vessels would have been taken into account.*

Method	Cat.	Image quality			Calcium score			Vessel			
		P	M	G	L	M	S	RCA	LAD	LCX	4th
MHT [49]	3	1	1	1	1	1	1	1	1	1	1
Tracer [155]	3	2	2	2	2	2	2	2	2	2	2
DepthFirstModelFit [188]	1	5	4	3	4	3	5	5	4	5	3
KnowledgeBasedMinPath [79]	2	3	3	4	3	4	3	3	6	3	4
AutoCoronaryTree [157]	1	7	8	7	8	7	6	6	9	6	9
GVFTube'n'Linkage [14]	1	6	5	6	6	6	4	4	3	7	6
CocomoBeach [76]	1	12	7	9	9	8	8	8	8	10	8
TwoPointMinCost [106]	3	4	6	5	5	5	7	7	5	4	5
VirtualContrast [168]	1	9	11	12	12	10	9	11	10	12	10
AxialSymmetry [38]	2	8	9	8	7	9	12	10	11	8	7
ElasticModel [61]	3	11	12	10	10	12	11	9	12	11	12
3DInteractiveTrack [190]	3	10	10	11	11	11	10	12	7	9	11
CoronaryTreeMorphoRec [32]	2	13	13	13	13	13	13	13	13	13	13

2.9 Discussion

A framework for the evaluation of CTA coronary artery centerline extraction techniques has been developed and made available through a web-based interface at: (*http://coronary.bigr.nl*). Currently thirty-two cardiac CTA datasets with corresponding reference standard centerlines are available for anyone how wants to benchmark a coronary artery centerline extraction algorithm.

Although the benefits of a large-scale quantitative evaluation and comparison of coronary artery centerline extraction algorithms are clear, no previous initiatives have been taken towards such an evaluation. This is probably because creating a reference standard for many datasets is a laborious task. Moreover, in order to get a good reference standard, annotations are needed from multiple observers and combining annotations from multiple observers is known to be difficult [172] and until recently unexplored for three-dimensional curves [165]. Furthermore, an appropriate set of evaluation measures has to be developed and a representative set of clinical datasets have to be made available. By addressing these issues we were able to present and use the proposed framework.

A limitation of the current study is the point-based vessel selection step for fully automatic methods. Because the coronary artery tree contains more vessels than the four annotated vessels this selection step had to be included, but it introduced the problem that fully automatic methods can extract many false-positives but still obtain a good ranking. This fact combined with the presented results of the fully automatic methods for the four evaluated vessels makes us believe that a future evaluation framework for coronary artery extraction methods should focus on the complete coronary tree. An obvious approach for such an evaluation would be to annotate the complete coronary artery tree in all the 32 datasets, but this is very labor intensive. An alternative approach would be to use the proposed framework for the quantitative evaluation of the four vessels and to qualitatively evaluate the complete tree. In this qualitative evaluation an observer should score if any vessels are falsely extracted and if all vessels of interest are extracted.

A further limitation of this study is that all the data have been acquired on two CT scanners of the same manufacturer in one medical center. We aim to extend the collection of datasets with datasets from different manufacturers and different medical centers. Further studies based on this framework could extend the framework with the evaluation of coronary lumen segmentation methods, coronary CTA calcium quantification methods or methods that quantify the degree of stenosis.

2.10 Conclusion

A publicly available standardized methodology for the evaluation and comparison of coronary centerline extraction algorithms is presented in this chapter. The poten-

tial of this framework has successfully been demonstrated by thoroughly comparing thirteen different coronary CTA centerline extraction techniques.

Chapter 3

Probabilistic Tracking of Elongated Structures in 3D Images.

Based on:

1) M. Schaap, I. Smal, C.T. Metz, T. van Walsum, W.J. Niessen.
Bayesian tracking of elongated structures in 3D images, Proceedings of Information Processing In Medical Imaging, 2007.

2) M. Schaap, R. Manniesing, I. Smal, T. van Walsum, A. van der Lugt, W.J. Niessen
Bayesian Tracking of Tubular Structures and its Application to Carotid Arteries in CTA, Proceedings of the International Conference on Medical Image Computing and Computer Assisted Intervention, 2007.

Tracking of tubular elongated structures is an important goal in a wide range of biomedical imaging applications. This chapter presents a probabilistic tube extraction algorithm that can better overcome data deficiencies in the tracking process than deterministic approaches and allows to easily incorporate a priori knowledge about the geometry and appearance of the elongated structure of interest. A key element of our approach is a dedicated observation model for tubular structures in regions with varying intensities. Furthermore, because probabilistic tube tracking algorithms are computationally complex, steps towards a computationally efficient implementation are suggested in this chapter.

The algorithm is quantitatively evaluated on synthetic data and on clinical CTA data of the internal carotid artery in 14 patients (28 carotids). Tracking of the internal carotid artery through the skull base in CT angiography data is a challenging problem, owing to the close proximity of bone, overlap in intensity values of lumen voxels and (partial volume) bone voxels, and the tortuous vessel paths. The tracking was successful in 25 cases, and the extracted paths were found to be close (< 1.0mm) to manually traced paths by two observers. The approach also showed good performance on synthetic data with high levels of Gaussian noise.

3.1 Introduction

Tracking of tubular elongated structures is an important goal in a wide range of biomedical imaging applications [85]. For example, vessel tracking in medical images has received considerable attention, as it can be used as a preprocessing step towards stenosis evaluation and grading, by generating multi-planar reformatted images [33].

Vessel centerline extraction methods can roughly be divided into two groups; bottom-up approaches and tracking methods. Bottom-up approaches (e.g. [155, 168]) detect centerline segments in the image and connect them to obtain the vascular centerlines. Tracking methods find paths in the image from a certain start-point to a manually specified, or automatically determined, end-point of the vessel of interest. In deterministic tracking methods (e.g. [8]) the algorithm starts at a given seed point and extends the path in the most probable direction at each iteration of the method, according to local image information. Since in these methods only one assumption of the path configuration is taken into account, this may lead to incorrect results if the path direction cannot be derived from imaging data locally, e.g. owing to pathologies or corrupted imaging data. State-of-the-art centerline tracking methods circumvent this problem by taking into account multiple hypotheses during tracking. These methods can be sub-divided into dynamic programming methods and probabilistic methods. In dynamic programming methods, minimum cost paths $\hat{\mathbf{X}}^{\mathbf{S},\mathbf{E}}$ are found that optimally connect a specificed start region $\mathbf{S}$ with an end region $\mathbf{E}$, according to a cost function $C(\mathbf{X}(s), \mathbf{X}'(s))$:

$$\hat{\mathbf{X}}^{\mathbf{S},\mathbf{E}} = \underset{\mathbf{X}^{\mathbf{S},\mathbf{E}}}{\operatorname{argmin}} \int C(\mathbf{X}(s), \mathbf{X}'(s)) ds, \tag{3.1}$$

with $\mathbf{X}^{\mathbf{S},\mathbf{E}} : [0..L] \rightarrow \Re^d$ a path with length L running from the specified start region $\mathbf{S}$ to the end region $\mathbf{E}$ in the d-dimensional image space. Often, the dependency on the tangent of the path $\mathbf{X}'(s)$ is ignored, leading to an isotropic cost function that can be solved efficiently with the Fast Marching algorithm [121, 181] or discretely approximated with Dijkstra's algorithm [52, 105]. Recently, work has been presented on finding minimum cost paths with an anisotropical cost function [103]. This requires computationally more demanding optimization algorithms, such as the Fast Sweeping algorithm, but can lead to significantly better results [103]. For all minimum cost path methods the cost function should be low on (and for anisotropic cost function, in the direction of) vessel centerlines, and high in other areas. Examples of cost functions include the reciprocal of Frangi's well-known vesselness measure [107, 181] and anisotropic oriented medial-based measures [52, 86]. In recent work attempts have been made to find minimal cost paths in a 4-dimensional space describing the 3-dimensional position and radius of a vessel of interest [18, 88], thereby obtaining next to the position of the vessels also their radius. The combination of 4D minimum cost paths and anisotropic costs has received only little

attention [18].

A disadvantage of minimum cost path methods is that often a large amount of locations in the d-dimension space have to be processed, and, because of Bellman's principal of optimality [17], the cost function can only depend on the position and tangent in the d-dimensional space. If, for example, the curvature of the path has influence on the likelihood of a path, the direction of the path should be added to the d-dimensional state vector $\mathbf{X}(s)$, thereby increasing its dimensionality and often leading to computationally intractable situations.

In probabilistic, or multiple hypotheses, centerline extraction the d-dimensional space is sparsely sampled, requiring the evaluation of less points in the state vector space [45, 114]. Therefore, in these approaches it is possible to add additional variables to the state vector in order to incorporate advanced prior information, which can result in more robust results. In this work we present a probabilistic tracking approach in which we add the radius, intensity and direction of the path to the state vector.

One of the important considerations when developing a tubular structure tracking algorithm is an observation model that describes the appearance of the tubular structures in the image, given the parameters of the tube (e.g. position and orientation). Often, probabilistic centerline extraction techniques use an observation model that is based on the assumption that vessels have, relative to their surrounding tissue, a bright intensity [45, 114]. This model is not suited for situations where the background contains both lower and higher intensities than the tube, which is e.g. the case for carotid arteries in CT angiography, which are surrounded by low intensity soft tissue and high intensity bone.

We present a probabilistic iterative tracking approach for tracking of elongated structures in 2D and 3D images. The novelty of this work is twofold. First, we present a generic and intuitive observation model to overcome the problems associated with the bright vessel observational model. The new observation model is specifically tailored for tracking homogeneous tubular structures through backgrounds with varying intensities. Second, to overcome the computational complexity of probabilistic tracking algorithms, several computational optimizations are proposed.

The presented method is evaluated by tracking elongated structures in synthetic data and by tracking the internal carotid artery through the skull base in CT angiography data of 14 patients. This is a challenging clinical problem, owing to the close proximity of bone, overlap in intensity values of lumen voxels and (partial volume) bone voxels, and the tortuous path of the vessels. The success rate and accuracy of the tracking are assessed quantitatively by comparing results on 14 datasets (28 carotids) to tracings by two observers.

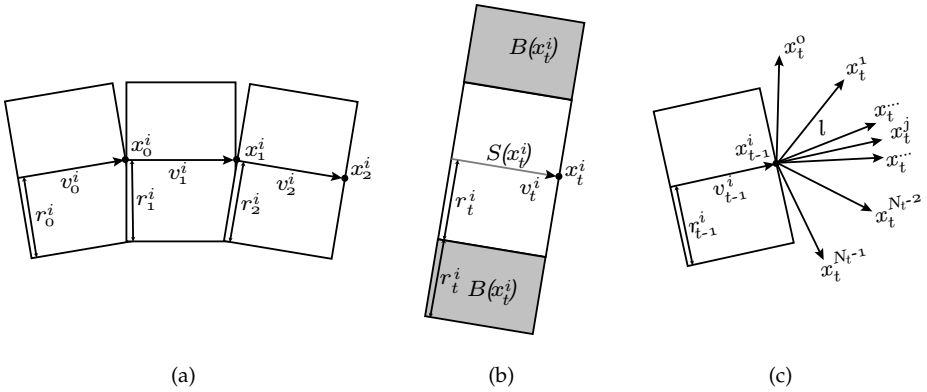

(a) (b) (c)

Figure 3.1: *Fig. (a) shows a part of the tube configuration $\mathbf{x}_{0:t}$. Fig. (b) visualizes the region in the tube $S(x_t^i)$ and the region in the band $B(x_t^i)$ around the tube. The prediction of new tubular segments, as explained in Section 3.2.5, is presented in Fig. (c).*

3.2 Method

The iterative probabilistic tracking approach searches for the track that is best explained by the data, given an observation model (which is based on our model of a tube), and prior information on the shape and appearance of the vessel (modeled via e.g. transition priors). The posterior probability density function (PDF) of a track, given the data, is recursively estimated. This is achieved by iteratively creating a set of currently most likely hypotheses. At the end, the maximum of the PDF determines the track that best represents the vessel of interest.

3.2.1 Tube model

Our tracking method considers an elongated structure as a series of straight tube segments (i.e. cylinders). A tube segment at iteration t is described by its location $\boldsymbol{p}_t = (x_t, y_t, z_t)^T$, orientation $\boldsymbol{v}_t = (\theta_t, \phi_t)$, radius r_t, and intensity I_t. Thus each tube segment is characterized by a state vector $\mathbf{x}_t = (\boldsymbol{p}_t, \boldsymbol{v}_t, r_t, I_t)^T$ and a complete tube is described by $\mathbf{x}_{0:t} \triangleq \{\mathbf{x}_0, \ldots, \mathbf{x}_t\}$, see Figure 3.1(a).

With every tube segment we associate a region of interest (ROI) $U(\mathbf{x}_t)$, defined by the components $\boldsymbol{p}_t$, r_t, and $\boldsymbol{v}_t$ of $\mathbf{x}_t$ (see Figure 3.1). Subsequently, we let $\mathbf{z}_t$ denote the image measurements (i.e. image intensities) within this ROI. Hence, all measurements corresponding to tube $\mathbf{x}_{0:t}$ are denoted with $\mathbf{z}_{0:t}$. $S(\mathbf{x}_t)$ defines the set of spatial coordinates that lie within the hypothesized tube and $B(\mathbf{x}_t) = U \backslash S$ defines the set of spatial coordinates in the band around the tube, see Figure 3.1(b).

3.2.2 Representing the posterior probability density function

Using Bayes' theorem, the probability density function (pdf) $p(\mathbf{x}_{0:t}|\mathbf{z}_{0:t})$, that describes the posterior probability of the tube configuration, having all the observations up to iteration t can be estimated with the following recursion [39]:

$$p(\mathbf{x}_{0:t}|\mathbf{z}_{0:t}) \propto p(\mathbf{x}_t|\mathbf{x}_{t-1})p(\mathbf{z}_t|\mathbf{x}_t)p(\mathbf{x}_{0:t-1}|\mathbf{z}_{0:t-1}), \tag{3.2}$$

where the transition prior $p(\mathbf{x}_t|\mathbf{x}_{t-1})$ is assumed to be Markovian ($\mathbf{x}_t$ only depends on $\mathbf{x}_{t-1}$ and not on any other past states). The observation model $p(\mathbf{z}_t|\mathbf{x}_t)$ relates the conditionally independent measurements at iteration t to the state $\mathbf{x}_t$.

The PDF can be represented with a set of N_t weighted states (also called particles):

$$p(\mathbf{x}_{0:t}|\mathbf{z}_{0:t}) = \sum_{i=1}^{N_t} w_t^i \delta(\mathbf{x}_{0:t} - \mathbf{x}_{0:t}^i), \tag{3.3}$$

where $\delta(\cdot)$ is the Dirac delta function.

Because the dimension of the joint space $p(\mathbf{x}_{0:t}|\mathbf{z}_{0:t})$ increases at every iteration the variance of the weights w_t^i also increases continuously [39]. In order to perform estimation in such high dimensional space a large number of particles is necessary, which would increase the computational complexity. Several techniques exist to counter-act this problem, of which the Sequential Importance Resampling particle filter (SIR) is well-known [7, 39, 51]. With SIR in each iteration a fixed number of particles is used which are efficiently redistributed in order to describe the pdf. In this work we propose a similar approach with the addition that we determine in each iteration the optimal number of particles (see section 3.2.6), requiring significantly less hypotheses in areas with conspicuous image information.

An additional optimization can be obtained by modeling the intensity in the marginal space, instead of in the joint space:

$$p(\mathbf{q}_{0:t}, I_t|\mathbf{z}_{0:t}) = p(I_t|\mathbf{q}_{0:t}, \mathbf{z}_{0:t})p(\mathbf{q}_{0:t}|\mathbf{z}_{0:t}). \tag{3.4}$$

This results in a lower dimension of the state space and thereby in less computational cost. Here, $p(\mathbf{q}_{0:t}|\mathbf{z}_{0:t})$ is calculated similarly to Equation 3.2. The marginal posterior intensity distribution $p(I_t|\mathbf{q}_{0:t}, \mathbf{z}_{0:t})$ is calculated as follows:

$$p(I_t|\mathbf{q}_{0:t}, \mathbf{z}_{0:t}) \propto \int p(I_t|I_{t-1})p(\mathbf{z}_t, \mathbf{q}_t|I_t)p(I_{t-1}|\mathbf{q}_{0:t-1}, \mathbf{z}_{0:t-1})dI_{t-1}. \tag{3.5}$$

Because this integral can be computationally difficult to calculate with the particle description in Equation 3.3 we assume that for each position, direction and radius $\mathbf{q}_t$ the intensities can be described with a normal distribution. By further assuming a linear Gaussian transition and observation model for the intensity component (see Section 3.2.3 and Section 3.2.3) the integral in Equation 3.5 can be calculated

analytically and very efficiently with the Kalman filter (see section 3.2.6). Thus, the PDF is described as follows:

$$p(\mathbf{q}_{0:t}, I_t|\mathbf{z}_{0:t}) = \sum_{i=1}^{N_t} w_t^i \delta(\mathbf{q}_{0:t} - \mathbf{q}_{0:t}^i)\mathcal{N}(I_t, P_t), \tag{3.6}$$

with $\mathcal{N}(.|\mu, \sigma^2)$ a real normal distribution with mean μ and variance σ^2, and P_t the modeled variance of the intensities.

The weights are normalized such that $\sum_{i=1}^{N_t} w_t^i = 1$. Each hypothesis $(\mathbf{q}_{0:t}^i, I_t)$ now represents a path $\mathbf{x}_{0:t}$, the directions $\mathbf{v}_{0:t}$, and radii $\mathbf{r}_{0:t}$, and an associated normal distribution for the intensity at the end of the path $\mathcal{N}(I_t, P_t)$. The corresponding weight w_t^i determines how likely this path-hypothesis is.

We will give two examples of different observation models in Section 3.2.3 and we describe our transition model in Section 3.2.4. Next, the method for generating new states, or hypotheses, is explained ('prediction') (Section 3.2.5), and lastly the mechanism to update the PDF is explained ('update') (Section 3.2.6).

3.2.3 Observation model

We present two different models for the observation model $p(\mathbf{z}_t|\mathbf{x}_t)$ in Equation 3.2. The first model is a common model designed for bright tubes in dark backgrounds, and the second is designed for tracking tubular structures through backgrounds with varying intensities, which can be both lower and higher than the tube of interest. In this last model we also incorporate an additional geometric prior.

For both models we define I_{in} as the average intensity of the intensities inside the tube (i.e. in the region $S(\mathbf{x}_t)$). Moreover, we assume that the observation model for the intensity I_t is a linear (conditional on the rest of the state parameters) Gaussian model given by:

$$I_{\text{in}} = I_t + \varsigma_t, \tag{3.7}$$

where $\varsigma_t \sim \mathcal{N}(\varsigma|0, \sigma_{\text{im}}^2)$, and σ_{im}^2 denotes the variance of the expected measurement error of the intensity I_{in}. In both observation models we use this model to determine what the conditional probability of the measured intensities is, given the predicted intensities I_t.

3.2.3.1 Bright-tube Observation Model

In this first observation model, we assume tubes to be bright relative to their background. Given the tube segment $\mathbf{x}_t$ and the measurements $\mathbf{z}_t$, the conditional probability of the measurements given the state $\mathbf{x}_t$, $p(\mathbf{z}_t|\mathbf{x}_t)$, is constructed as follows:

$$p(\mathbf{z}_t|\mathbf{x}_t) = p(\mathbf{z}_t|\boldsymbol{p}_t, \boldsymbol{v}_t, r_t, I_t) \propto p(I_{\text{in}}, I_{\text{out}})\mathcal{N}(I_{\text{in}}|I_t, \varsigma_t), \tag{3.8}$$

where

$$p(I_{\text{in}}, I_{\text{out}}) \propto \begin{cases} (\frac{I_{\text{in}} - I_{\text{out}}}{I_c})^s, & I_{\text{in}} > I_{\text{out}}, \\ 0, & I_{\text{in}} \leq I_{\text{out}}. \end{cases} \tag{3.9}$$

with I_{out} the average intensity of the intensities in the band outside the tube (i.e. in the region $B(\mathbf{x}_t^i)$), and I_c and s contrast regulating parameters. The term $p(I_{\text{in}}, I_{\text{out}})$ describes the probability that the observed intensities inside and outside the tube segment are representative for a bright tube. Finally, $\mathcal{N}(I_{\text{in}}|I_t, \sigma_{\text{im}}^2)$ describes the probability that the measured average inside intensity I_{in} resembles the modeled intensity.

3.2.3.2 Bhattacharyya Observation Model

The second observation model is similar to the first; the likelihood of the state vector given the observations depends on 1) how much the inside and outside intensities differ, and 2) how much the measured intensities differ from the predicted intensities.

This model is different from the first in how the intensities are compared. Implicitly, we assume in the first model that the distributions of the inside and outside intensities are mono-modal and that they have similar variation. In the second model we assume a Gaussian distribution for the inside intensity, but no specific distribution for the outside intensity. This makes it possible to use this model in situations when the only a priori appearance information about the tubes is that the intensities in the tube have a homogeneous intensity with additive Gaussian disturbance, which differs from the intensity of surrounding tissue.

With $p(I|S(\mathbf{x}_t))$ and $p(I|B(\mathbf{x}_t))$ we describe the normalized intensity histograms of the voxels at the inside and outside of the tube segment respectively. These distributions are constructed by sampling from the ROI defined by $\mathbf{x}_t$ using nearest neighbor interpolation. Given the tube segment $\mathbf{x}_t$ and the measurements $\mathbf{z}_t$, the conditional probability of the observation given the state $p(\mathbf{z}_t|\mathbf{x}_t)$ is given by:

$$p(\mathbf{z}_t|\mathbf{x}_t) = p(\mathbf{z}_t|\boldsymbol{p}_t, \boldsymbol{v}_t, r_t, I_t, \sigma_t) \propto D_{cp,t}\,(1 - D_{sb,t}), \tag{3.10}$$

where $D_{cp,t}$ and $D_{sb,t}$ denote respectively the similarity between the measured inside intensity and the modeled intensity I_t and the similarity between the intensity distribution inside and outside of the tube. These two intensity measures are calculated as follows:

$$D_{cp,t} = D(\mathcal{N}(I|\hat{I}_{s,t}, \hat{\sigma}_{s,t}^2), p(I|I_t, \sigma_{\text{im}}^2))^{c_1} \tag{3.11}$$

$$D_{sb,t} = D(\mathcal{N}(I|\hat{I}_{s,t}, \hat{\sigma}_{s,t}^2), p(I|B(\mathbf{x}_t)))^{c_2}, \tag{3.12}$$

where $\hat{I}_{s,t}$ and $\hat{\sigma}_{s,t}^2$ describe respectively the mean and variance of the intensity histogram $p(I|S(\mathbf{x}_t))$, and c_1 and c_2 are parameters to regulate the influence of the

different components. Several methods have been suggested in literature to calculate the similarity between two distributions. We use the Bhattacharyya Metric (BM) [158], defined by:

$$D(p_1, p_2) = BM(p_1, p_2) = \int \sqrt{p_1(x)p_2(x)}dx. \tag{3.13}$$

3.2.4 Transition model

The transition model $p(\mathbf{x}_t|\mathbf{x}_{t-1})$ in Equation 3.2 describes the conditional probability of a state $\mathbf{x}_t$ given its previous state $\mathbf{x}_{t-1}$. In our case we assume that the transition for the position, radius and intensity are indepedent. Therefore the transition model can be factorized as:

$$p(\mathbf{x}_t|\mathbf{x}_{t-1}) = p(\boldsymbol{p}_t, \boldsymbol{v}_t|\boldsymbol{p}_{t-1}, \boldsymbol{v}_{t-1})p(r_t|r_{t-1})p(I_t|I_{t-1}). \tag{3.14}$$

We define the transition density $p(\boldsymbol{p}_t, \boldsymbol{v}_t|\boldsymbol{p}^i_{t-1}, \boldsymbol{v}^i_{t-1})$ to be proportional to a Gaussian distribution with zero mean and standard deviation σ^2_φ on the angle between the tube axes, i.e.:

$$p(\boldsymbol{p}_t, \boldsymbol{v}_t|\boldsymbol{p}^i_{t-1}, \boldsymbol{v}^i_{t-1}) = \sum_{j=0}^{N^i_t-1} \tilde{\omega_j}\delta(\boldsymbol{p}_t - \boldsymbol{p}^j_t), \tag{3.15}$$

where the weight $\tilde{\omega_j}$ of a given enclosed angle ϕ_j between $\boldsymbol{v}_t$ and $\boldsymbol{v}_{t-1}$ is given by:

$$\tilde{\omega_j} = \frac{\omega(\phi_j)}{\sum_{k=0}^{N^i_t}\omega(\phi_k)} \tag{3.16}$$

$$\omega(\varphi) = N(\varphi|0, \sigma^2_\varphi), \tag{3.17}$$

with σ_φ a pre-defined parameter. We also model the transitions of r_t and I_t as linear Gaussian models:

$$r_t = r_{t-1} + \eta_t, \tag{3.18}$$

$$I_t = I_{t-1} + \xi_t, \tag{3.19}$$

where η_t and ξ_t are uncorrelated Gaussian random variables with variances σ^2_r and σ^2_I, respectively. Thus:

$$p(r_t|r_{t-1}) = \mathcal{N}(r_t|r_{t-1}, \sigma^2_r) \tag{3.20}$$

$$p(I_t|I_{t-1}) = \mathcal{N}(I_t|I_{t-1}, \sigma^2_I) \tag{3.21}$$

Additionally we propose an optional spatial term $p_M(\mathbf{q}_t|\mathbf{q}_{0:t-1})$ in the transition prior in order to prevent loops in the tracked tube. For this we do not model

the transition prior as being Markovian. Instead, we determine the posterior probability $p_M(\mathbf{q}_t|\mathbf{q}_{0:t-1})$ of a tube segment $\mathbf{q}_t$ given all previous tube segments and we incorporate this term in the transition prior, thereby obtaining the following modified (non-Markovian) transition prior:

$$p(\mathbf{x}_t|\mathbf{x}_{0:t-1}) = p_M(\mathbf{q}_t|\mathbf{q}_{0:t-1})p(\mathbf{x}_t|\mathbf{x}_{t-1}). \tag{3.22}$$

For the term $p_M(\mathbf{q}_t|\mathbf{q}_{0:t-1})$ we create a spatial map $M_t(x)$ which stores for each image voxel the probability that the voxel has not been identified as being part of the tubular structure. This map is initialized with $M_0(x) = 1$. Then, in each iteration the map is updated according to the current hypotheses and their likelihood. This update step is described in section 3.2.7. In the calculation of the transition prior the term $p_M(\mathbf{q}_t|\mathbf{q}_{0:t-1})$ is constructed by averaging over all values of this map that fall within the inside of the tube ($S(\mathbf{q}_t)$):

$$p_M(\mathbf{q}_t|\mathbf{q}_{0:t-1}) = \frac{\sum_{\boldsymbol{p}\in S(\mathbf{q}_t)} M_t(\boldsymbol{p})}{|S(\mathbf{q}_t)|}, \tag{3.23}$$

where $|.|$ defines the set size operator.

3.2.5 Prediction

In the prediction step we create N_t^i new hypotheses $\mathbf{x}_t^j : j \in \{1..., N_{t-1}\}$ for each hypothesis $\mathbf{x}_{t-1}^i$ from the previous iteration $t-1$. In the update step both the transition model and observation model are used to calculate new weights and intensity distributions for these new hypotheses.

First, new positions $\boldsymbol{p}_t^j$ and directions $\boldsymbol{v}_t^j$ are determined by sampling for each hypothesis $\mathbf{x}_t^i$ a set of $w_{t-1}^i N_{t-1}$ new hypotheses on the half sphere (with radius l_t^i) in front of $\boldsymbol{p}_{t-1}^i$ oriented in the direction of $\boldsymbol{v}_{t-1}^i$. They are sampled according to:

$$\boldsymbol{p}_t^j = \boldsymbol{p}_{t-1}^i + l_t^i \boldsymbol{v}_t^i \tag{3.24}$$

$$\boldsymbol{v}_t^i = R_z(\theta_{t-1}^i) R_y(\phi_{t-1}^i) R_z(\vartheta_t^j) R_y(\varphi_t^j)(0,0,1)^T \tag{3.25}$$

where $R_z(.)$ and $R_y(.)$ are rotation matrices around the z- and y-axis [176]. The angles $(\vartheta_t^j, \varphi_t^j)$ describe a point in the local spherical coordinate system with the z-axis orientated in the direction of $\boldsymbol{v}_{t-1}^i$ and origin at $\boldsymbol{p}_{t-1}$. They are chosen with an algorithm that uniformly distributes points on the half sphere, as described by Saff and Kuijlaars [139]. Because the angle φ_t^j is equal to the enclosed angle between $\boldsymbol{v}_{t-1}^i$ and $\boldsymbol{v}_t^j$ the transition probability in Equation 3.15 can efficiently be calculated.

Depending on the application of choice, the length of a tube segment l_t^i can be set to a constant value or be dependent on the radius of the tube r_{t-1}^i and a predefined

parameter α:

$$l_t^i = \frac{r_{t-1}^i}{\tan(\alpha)}. \tag{3.26}$$

Figure 3.1(c) provides a schematic explanation of the position transition in eq. (3.25). The transition model for the radius and intensity will be discussed in the next section.

3.2.6 Update

In the update step the weights $\{w_t^i\}$ and the intensity descriptors (I_t^i, P_t^i) in iteration t are calculated according to the observation model and the observations $\mathbf{z}_t$.

First we associate with each position,direction, radius $(\boldsymbol{p}_{t-1}^i, \boldsymbol{v}_{t-1}^i, r_{t-1}^i)$ a discrete set of K radii $r_t^{i,k}$, sampled uniformly between $r_{t-1}^i - \Delta r$ and $r_{t-1}^i + \Delta r$. Then we collect for each radius $r_t^{i,k}$ image statistics $\mathbf{z}_t^{i,k}$, and we calculate weights according to the observation and transition models:

$$w_t^{i,k} = w_{t-1}^i p(\mathbf{z}_t^{i,k}|\mathbf{x}_t^i) p(\mathbf{x}_t^{i,k}|\mathbf{x}_t^i). \tag{3.27}$$

Because the transition and observation model for the intensity component are both modelled as linear Gaussian models (see Equation 3.19 and 3.7) the intensity prediction and update can be solved analytically using the Kalman filter[39]. First the intensity mean and variance are predicted:

$$\begin{aligned} I_{t|t-1}^i &= I_{t-1}^i \\ P_{t|t-1}^i &= P_{t-1}^i + \sigma_I^2 \end{aligned}$$

and then they are updated:

$$\begin{aligned} I_t^{i,k} &= I_{t|t-1}^i + K_t (I_{\text{in}}^{i,k} - I_{t|t-1})^i, \\ P_t^{i,k} &= (1 - K_t) P_{t|t-1}^i \end{aligned}$$

where K_t denotes the Kalman gain $K_t = \frac{P_{t|t-1}^i}{P_{t|t-1}^i + \sigma_{\text{im}}^2}$. In the first iteration, when no a priori knowledge about the image intensity distribution is available, the Kalman filter is not used and I_1^i is set to I_{in} and P_1^i to σ_{im}^2. After calculating all radius weights, the radius of the new particle r_t^i is updated by selecting the radius with the highest weight $w_t^{i,k}$:

$$r_t^i = r_t^{i,\hat{k}} \tag{3.28}$$

with $\hat{k} = \text{argmax}_k \, w_t^{j,k}$. Finally, we set the weight and intensity descriptors to the values associated with the radius with maximum weight: $w_t^i = w_t^{i,\hat{k}}$, $I_t^i = I_t^{i,\hat{k}}$, $P_t^i = P_t^{i,\hat{k}}$.

3.2.7 Tube probability update

The spatial prior $p_M(\mathbf{x}_t)$ uses the map $M_t(x)$ that describes for each position the probability that it (still) can be traversed. For the first iteration this map is set to $M_t(x) = 1$, meaning that hypotheses are plausible at all positions. Given this first initialization, the map is updated in each iteration as follows:

$$M_{t+1}(\boldsymbol{p}) = M_t(\boldsymbol{p})(1 - \sum_{\forall i: \boldsymbol{p} \in S(x_t^i)} w_t^i). \tag{3.29}$$

3.2.8 Keeping the effective hypotheses

For improved efficiency we use in each iteration the variance of the (normalized) weights w_t^i to determine how many and which hypotheses should be kept in the next iteration. The N_t most probable hypotheses are kept according to the weights $w_t^i, i \in \{1, ..., N_t\}$, where:

$$N_t = \frac{1}{\sum_i^{N_t} (w_t^i)^2}. \tag{3.30}$$

From each of these states, $N_t^i = \max(\text{nint}(w_t^i N), N_{max})$ new states are created, where nint(.) denotes a nearest integer round, N is pre-defined and describes the maximum total number of hypotheses created and $N_{\max}$ describes the maximum number of hypotheses created from one hypothesis. This approach will keep only the relevant hypotheses and effectively distributes them according to the described pdf.

3.3 Synthetic Evaluation

The developed probabilistic tracking algorithm has been evaluated on 2D and 3D synthetic data and simulated computed tomography angiography (CTA) data.

3.3.1 Parameters and initialization

For all the experiments on the synthetic data we used the bright tube model with tubes of variable tube lengths (see Section 3.2.5) and we used fixed parameter settings that were empirically selected based on experiments on synthetic data not belonging to the test set. In this empirical optimization study it was observed that the method was not sensitive to parameter changes.

The following parameter settings were used: maximum number of hypotheses $N = 500$, maximum number of hypotheses created from an hypothesis $N_{\max} = 50$, observation model parameters $I_c = 1.0$ and $s = 1.0$, and (in arbitrary units): direction transition variance $\sigma_\varphi^2 = 1.5$, radius transition variance $\sigma_r^2 = 2$, intensity transition variance $\sigma_I^2 = 1000$, intensity measurement error $\sigma_{\text{im}}^2 = 10$, variable tube length parameter $\alpha = 0.25$.

The algorithm is initialized with a seed point and a radius estimation. The initial set of state vectors $\mathbf{X}_0$ contains two hypotheses, with locations $\boldsymbol{p}_0^0$ and $\boldsymbol{p}_0^1$ equal to the seed point and radii r_0^0 and r_0^1 equal to the estimated radius. The orientations were initialized in two arbitrarily, and opposing, directions: $\boldsymbol{v}_0^0 = (0, 0)$ and $\boldsymbol{v}_0^1 = (\pi, 0)$.

3.3.2 2D synthetic data

In order to investigate the accuracy of the developed method as a function of image noise, a quantitative study on 2D synthetic data was carried out. In this synthetic study we use the distinguishability measure $\text{MD} = \text{CNR}\sqrt{A}$, where CNR is the ratio of the contrast between an object and its background and the standard deviation of the noise in the image and A denotes the amount of voxels of the object [135]. Four random tubes were created with radii and distinguishability values of respectively $r = \{2.5, 5.0, 7.5, 10\}$ pixels and $\text{MD} = \sqrt{2}\{0.5, 1.0, 1.5, 2.0\}$ pixels, resulting in 64 different test images.

These synthetic tubes were created by fitting 3rd order splines through five randomly selected points in a 2D grid of 512×512 pixels. Tubes were excluded if a part of the tube had a distance to other parts of the tube, not being direct neighbors, of less than 100 pixels, or if the maximum curvature of the tube centerline was higher than 0.04 pixels^{-1}. See Figure 3.3(a) for an example tube.

The tubes were both tracked automatically and manually which allows comparison between the obtained accuracy and the human capabilities of tracking tubes in noisy data. Four observers each annotated centerlines in 16 of the 64 images. The different radii and noise levels were equally distributed over the four observers.

In the automatic tracking, the radius was initialized at the average size of the synthetic tubes evaluated, being 6.25 pixels. It is believed that this simulates typical a priori knowledge of an expected tube size.

As an evaluation criterion we use the root mean squared distance (RMSD) between the automatically or manually tracked tube and the centerline of the golden standard. This number is divided by the radius of the tube in order to present the accuracy of tracking relative to the tube size (RMSD/r).

Only distances for paths where at least 95% of the track is found at a distance within twice the radius of the golden standard are taken into account. In this way the ability of tracking the tube and the accuracy of tracking is presented.

3.3.3 3D synthetic data

A 3D synthetic evaluation was carried out to evaluate the performance of the algorithm for tracking 3D tubes with varying radius and high image noise. A 3D synthetic tube with varying radius from $r = 4$ voxels to $r = 10$ voxels is created and Gaussian noise is added, resulting in a CNR value of 0.25 and thereby distinguishability measures of MD $= 1.8$ to MD $= 4.4$ voxels. The algorithm was initialized with a seed point at the beginning of the tube and radius was set to the average radius of 7 voxels. The performance of the algorithm is evaluated by measuring the RMSD to the centerline of the synthetic tube.

3.3.4 Simulated CTA data

To simulate clinical data with high noise, a computed tomography angiography (CTA) artery was manually segmented from CTA data. The segmentation was given a contrast of 200 intensity units to its background, similar to the contrast of arteries in CTA. This segmentation was smoothed with a Gaussian kernel with standard deviation of 1 voxel, similar to the point spread function of CT. Afterward Gaussian noise was added, with a standard deviations of 240 intensity units, approximately eight times higher than the clinical noise value of 30 intensity units. The algorithm was initialized with a seedpoint at the approximate beginning of the tube and the radius was set to approximately twice the average radius of the simulated vessel. The RMSD difference to a manual tracked centerline is used as a performance measure.

3.4 Evaluation on clinical CTA data

The second part of our evaluation consisted on an evaluation on clinical CTA data. In order to evaluate the additional value of the Bhattacharyya observation model we used a clinical problem that is difficult with standard bright-tube observation models; tracking the internal carotid arteries (ICAs) in CTA data covering the anatomical region containing the skull base. In these areas a vessel segment can be surrounded by both lower and higher intensities.

3.4.1 Data acquisition and parameter settings

Fourteen consecutive patients suffering from a transient ischemic attack underwent a CTA examination. The data, acquired on a 16-slice CT scanner (Siemens Somatom Sensation 16, Forchheim, Germany) were reconstructed using a B46f kernel, and resampled in-plane with linear interpolation to a 256×256 matrix, resulting in voxel sizes of $0.5 \times 0.5 \times 1.2$mm. Visual inspection showed that 24 out of 28 carotids contained some form of pathology; 16 carotids showed mild calcifications, stenoses or aneurysms and eight carotids contained severe calcifications and/or stenoses.

In this evaluation the following parameter settings were used: maximum number of hypotheses $N = 1000$, maximum number of hypotheses created from an hypothesis $N_{\max} = 50$, observation model parameters $c_1 = 10$, $c_2 = 30$, direction transition variance $\sigma_\varphi^2 = 2.0$, radius transition variance $\sigma_r^2 = 0.5$, intensity transition variance $\sigma_I^2 = 1000$, intensity measurement error $\sigma_{\text{im}}^2 = 1000$.

3.4.2 Carotid tracking

The method is evaluated by tracking the internal carotid arteries (ICAs) in CTA data covering the anatomical region containing the skull base. A possible approach is to acquire an additional unenhanced CT scan to mask high intensity structures, but this is associated with increased radiation dose for the patient and increased acquisition time. Recently, methods have appeared that directly track the vessels in carotid CTA, [99, 150, 154], but either lack methodological details [154] and/or an extensive validation [99, 150]. The method described in [99] seems to be the most promising to date, but evaluation was limited to visual inspection by observing whether the tracked path was fully contained within the ICAs. In this work we quantitatively compare the results of our method to manually traced paths by two observers.

The algorithm is initialized by placing a small vector in the vessel just after the carotid bifurcation. The method then iterates until a vessel length of at least 180 mm is tracked to ensure that each tracked path goes beyond the region containing the skull base and reaches the Circle of Willis. The segment length l_t was set to 4 mm, therefore the final path will contain 46 points. For comparison, two observers manually traced both ICAs in each patients (resulting in a total of 56 manually traced paths). For this manual tracing, points were first annotated on the axial slices and then visually inspected on multiplanar reformatted (MPR) images to determine whether the point was properly centralized. Points could be corrected on the MPR images if needed.

A comparison between two paths is made segment-wise, by finding the minimum distances of two neighboring points on the first path to the second path. The radii of ICAs are approximately in the range of 1.0mm-2.5mm. If both distances are found to be smaller than 2.5 mm, than this path segment is considered to be corresponding. For the corresponding path segments the average distance is computed. In this way we evaluate both the tracking success rate and tracking accuracy. Both criteria are applied in the path comparisons between the method and observers $\mathcal{M} \rightarrow \mathcal{O}_{\{1,2\}}$, and between the observers $\mathcal{O}_1 \leftrightarrow \mathcal{O}_2$.

3.4.3 Coronary tracking

The algorithm was also used to track two coronary arteries in a cardiac CTA dataset. One of the arteries had a chronic total occlusion, meaning that an artery was blocked

$\frac{MD}{\sqrt{2}}$	Manual			Automatic		
	#found	RMSD	RMSD/r	#found	RMSD	RMSD/r
0.5	14/16	3.12	0.40	5/16	2.85	0.39
1.0	16/16	1.37	0.25	16/16	1.72	0.31
1.5	16/16	1.11	0.21	16/16	1.60	0.30
2.0	16/16	1.05	0.19	16/16	1.52	0.28

Table 3.1: *Evaluation results for the evaluation with 2D synthetic data. A tube is said to be found if at least 95 percent of the length of the path was tracked at a distance to the golden standard of less than two times the radius of the tube. For tubes that were found both manually and automatically, the root mean squared difference (RMSD, in pixels) and the root of the average of the squared ratio between distance and radius RMSD/r is shown.*

and (almost) no blood was passing the point of occlusion. Such a dataset is problematic for conventional deterministic algorithms, because locally the image is lacking sufficient contrast. The algorithm was initialized by putting a seedpoint in the vessels and setting the radius to approximately twice the radius of the vessel.

3.5 Results

3.5.1 2D synthetic data

The results of the evaluation on 2D synthetic data are shown in Figure 3.2 and Table 3.1. Figure 3.2 shows the relation between the accuracy of the manually and automatically tracked paths. Table 3.1 shows the performance of the manual and automatic trackings, for the four different distinguishability measures. An example of the input data and the automatic tracking is presented in Figure 3.3. For the automatically tracked paths that had an overlap of less than 95%, the overlap was always less than 45%. Two of these tracks were also not found manually with an overlap of at least 95% (respectively 85% and 88%).

3.5.2 3D synthetic data

The algorithm tracked the full length of the path in the 3D synthetic tube. The RMSD to the golden standard was 2.3 voxels. Figure 3.4 visualizes this result.

3.5.3 Simulated CTA data

In the simulated CTA dataset the track was found automatically with an RMSD of 1.7 voxels to the manually annotated reference standard. Figure 3.5 shows the simulated CTA dataset, found path and reference standard.

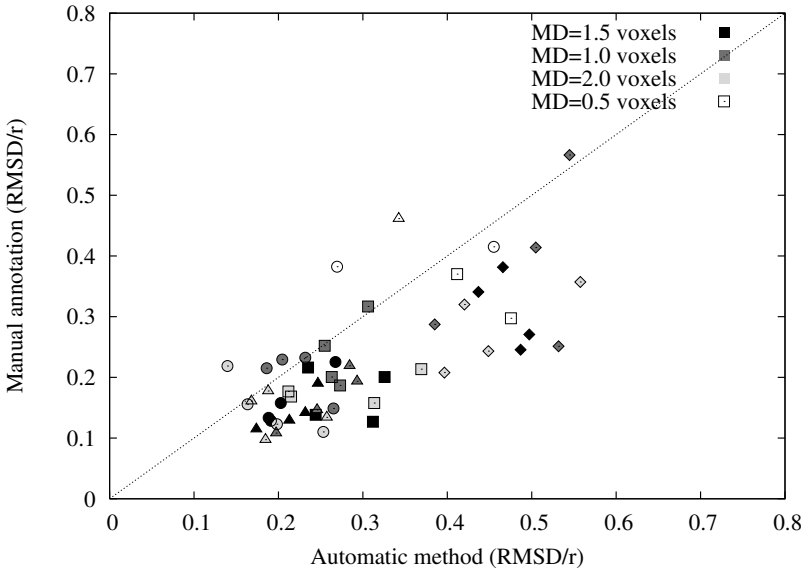

Figure 3.2 (see page 145 for a color-version): *Scatter plot demonstrating the relation between the accuracy, measured in RMSD/r (in pixels), of the manually and automatically tracked tube centerlines. The different symbols denote different radii (◇ = 2.5, □ = 5.0, ◯ = 7.5, △ = 10.0, all in voxels), and different colors denote different noise levels.*

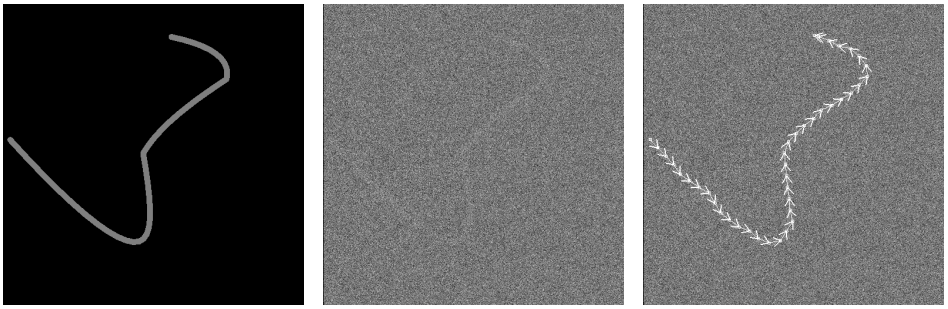

Figure 3.3: *An example of the data and results of the 2D synthetic evaluation. Figure (a) is a randomly created tube with radius of 5 voxels and 200 intensity units contrast. Figure (b) is the image after Gaussian noise with a standard deviation of 894 was added, resulting in CNR* $= 0.224$ *and* $M_d = 0.5\sqrt{2}$*. Figure (c) shows the tracking result of the algorithm.*

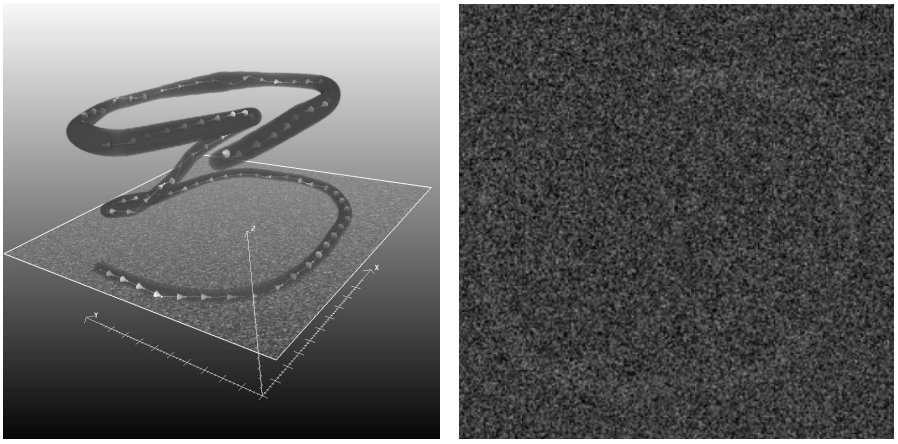

(a) (b)

Figure 3.4: *A demonstration of tracking with the developed algorithm of a 3D synthetic tube with radius varying between 5 and 10 voxels and CNR of 0.25. In (a) the tracking result in 3D is shown with the tube without noise and a slice of the volume with noise. (b) shows the noisy slice in 2D.*

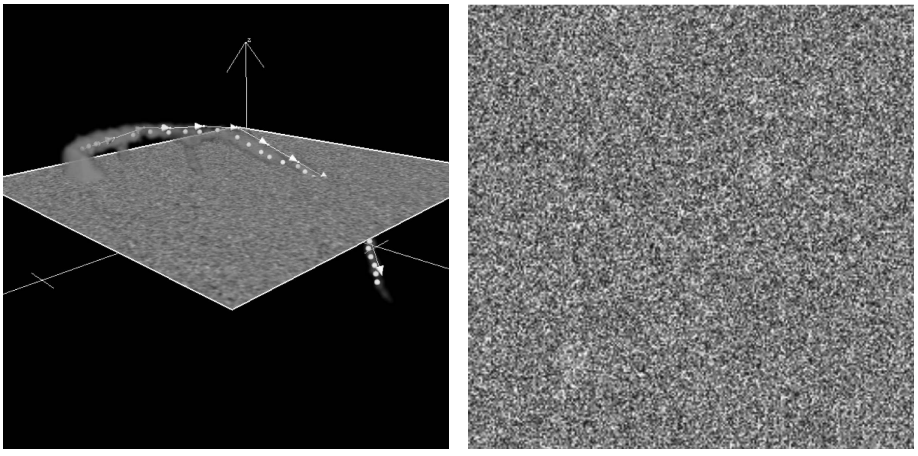

(a) (b)

Figure 3.5: *(a) An automatically found track (shown with arrows) in a simulated CTA dataset displayed together with the simulated dataset without noise, manual reference standard (white dots), and a slice from the dataset with noise. (b) shows the slice with noise in 2D.*

	$\mathcal{M} \rightarrow \mathcal{O}_1$	$\mathcal{M} \rightarrow \mathcal{O}_2$	$\mathcal{O}_1 \rightarrow \mathcal{O}_2$	$\mathcal{O}_2 \rightarrow \mathcal{O}_1$
Overlap	0.95%	0.95%	0.99%	0.99%
Avg. Dist.	0.69	0.67	0.46	0.47

Table 3.2: *The results of the comparison between observer 1 (O_1), observer 1 (O_1), and the method (M). 'Overlap' shows the percentage of the total length of all 14 tracks found within 2.5mm. 'Avg. Dist.' shows for all corresponding segments the average distance in millimeters to the reference track, weighted by segment length.*

	$\mathcal{M} \rightarrow \mathcal{O}_1$	$\mathcal{M} \rightarrow \mathcal{O}_2$
100%	23	22
90-100%	1	2
80-90%	2	3
$<$80%	3	2

Table 3.3: *The automatic tracks (M) categorized on the percentage found within 2.5mm of the reference standard, respectively observer 1 (O_1) and 2 (O_2).*

3.5.4 Carotid artery tracking

The results of tracking the carotid artery through the skull base are summarized in Table 3.2 and 3.3. An example of the results is shown in Figure 3.6. Overall, 95% of the tracked segments was found within 2.5 mm of the annotated centerlines and for these segments the average distance was approximately 0.7 mm, which only is 0.2 mm larger than the interobserver variability. Furthermore, the average distance of the corresponding segments per track was always smaller than 1.0 mm.

Seven tracks had an overlap of less than 100% with the track of one of the observers. Visual inspection showed that an annotation error was made by observer two, causing one of these cases. The other six tracks with less than 100% correspondance can be divided in a group of three tracks ($>$ 80%) where the track has some minor localization errors in the distal parts of the ICA and a group of three ($<$ 80%) where the method failed to track from the seedpoint to the Circle of Willis. For these vessels, respectively 57, 72 and 78% was tracked.

3.5.5 Coronary artery tracking

The two clinical examples are shown in Figure 3.7. Both trackings were visually evaluated and found to be well within the boundary of the vessel.

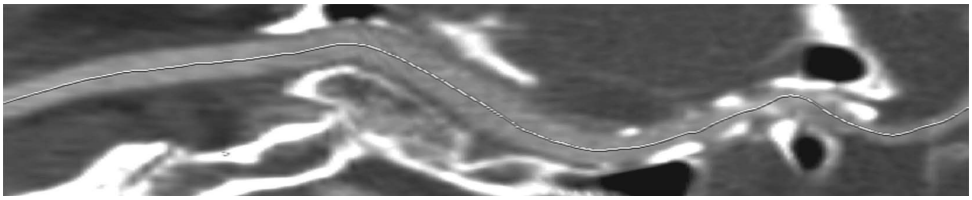

Figure 3.6: *An example of an automatically generated curved planar reformatted image. The image shows calcifications in the internal carotid artery in the region after the skull base, depicted on the right side of the image.*

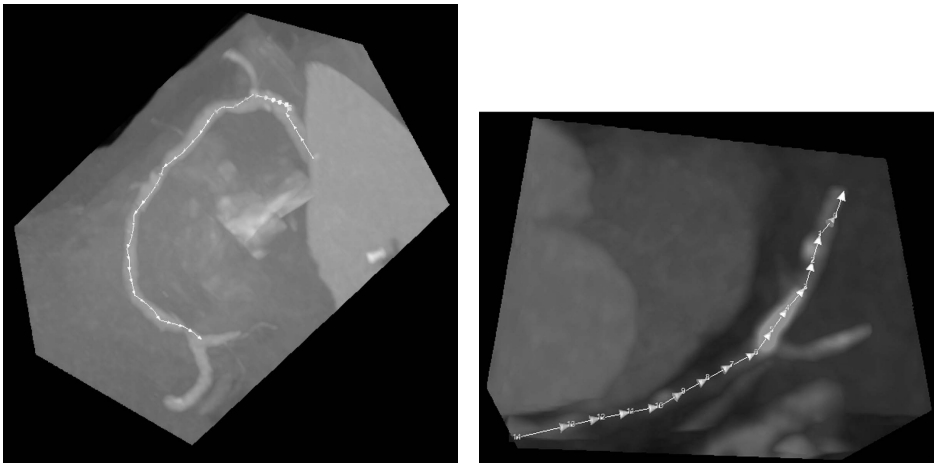

(a) (b)

Figure 3.7: *Examples of a clinical application of the probabilistic tracking algorithm. (a) shows a tracked right coronary artery in a CTA dataset. In (b) it is shown that an artery with a chronic total occlusion (CTO) in a CTA dataset is successfully tracked.*

3.6 Discussion and Conclusion

In this chapter we have presented a tubular tracking approach within a probabilistic framework. The method is flexible, since prior information can easily be incorporated. In our case, important prior information is incorporated through a novel observation model. Second, its non-deterministic character increases robustness of the method. Finally, because probabilistic tube tracking algorithms are computationally complex, several computational improvements are suggested in this chapter. We presented the use of the Kalman filter for the intensity component of the

state vector. This continuous representation of the intensity reduces the amount of necessary hypotheses significantly in comparison to a discrete representation. The second computational contribution is to adapt the amount of hypotheses on the conspicuity of the image information according to the variance of the importance weights.

The algorithm was evaluated on 2D synthetic data. Synthetic tubes in images with high additive Gaussian noise ($\mathrm{CNR}\sqrt{\mathrm{radius}} = 1.0$) were successfully tracked with accuracies that were in the same order of magnitude as the tracking results of the human observers. The evaluation on 2D synthetic data showed that the algorithm was capable of tracking tubes with a distinguishability measure MD of 1.4 with accuracies that were in the same order of magnitude as the tracking results of the human observers. As a reference for values of MD we would like to refer to [8, 97] for respectively tube enhancement and tracking. In [97] it is reported that straight tubes were successfully enhanced at $\mathrm{MD} > 5.3$. The values presented in [8] are more difficult to compare, because varying tube radii were used, but indicative values can be obtained from this article. The algorithm was also applied to 3D synthetic data (with a CNR of 0.25) and simulated high noise CTA data. Both tracks were found automatically with a RMSD of approximately 2 voxels to the reference standard.

The method has also been applied to the tracking of carotid and coronaries in CTA data. 28 ICAs were successfully tracked through the difficult region of the skull base. The resulting paths were found to be close (< 1.0mm) to manually traced paths by two observers.

The development of automatic stopping criteria, bifurcation detection, and a comparison with existing deterministic and probabilistic algorithms are subject to future work. The results are already promising for a wide variety of applications, and it should be noted that when using the presented algorithm for a specific application more a priori knowledge can be incorporated, such as expected intensity distributions and more specific observation models.

Chapter 4

Graph Cuts and Robust Kernel Regression for Coronary Segmentation

Based on:

M. Schaap, L. Neefjes, C.T. Metz, A.G. van der Giessen, A. Weustink, N. Mollet, J. Wentzel, T. van Walsum, W.J. Niessen. *Coronary Lumen Segmentation using Graph Cuts and Robust Kernel Regression, Proceedings of Information Processing In Medical Imaging, 2009.*

This chapter presents a novel method for segmenting the coronary lumen in CTA data. The method is based on graph cuts, with edge-weights depending on the intensity of the centerline, and robust kernel regression. A quantitative evaluation in 28 coronary arteries from 12 patients is performed by comparing the semi-automatic segmentations to manual annotations. This evaluation showed that the method was able to segment the coronary arteries with high accuracy, compared to manually annotated segmentations, which is reflected in a Dice coefficient of 0.85 and average symmetric surface distance of 0.22 mm.

4.1 Introduction

Coronary artery disease (CAD) is one of the leading causes of death worldwide [134]. One of the imaging methods for diagnosing CAD is Computed Tomography Angiography (CTA) (see Figure 4.1(a) for a volume rendering of a CTA dataset), a non-invasive technique that allows the assessment of the coronary lumen and the evaluation of the presence, extent, and type (non-calcified or calcified) of coronary plaque [82]. Cardiac CTA therefore has large potential to improve risk stratification of CAD, requiring methods for objective and accurate quantification of coronary lumen and plaque parameters.

Since manual annotation of the lumen, calcium and soft plaque is very labor intensive, (semi-)automatic techniques are needed to efficiently quantify these parameters in cardiac CTA data. In this chapter we focus on semi-automatic coronary lumen segmentation.

Coronary lumen segmentation is a challenging task owing to the small size of the coronary arteries (their size ranges from approximately 5 mm to less than 1 mm in diameter), the limited spatial resolution of CT (approximately 0.7 mm to 1.4 mm [132]), motion induced blurring, high intensity calcium close to the coronary lumen, and the presence of severe stenoses.

Existing coronary segmentation methods can roughly be divided into two categories: methods that segment the coronaries in one pass and methods that first find the vessel centerline and then segment the vessel. The methods that segment the vessels in one pass can further be divided into methods that use region-growing or a combination of different morphology operators [21, 22, 96], methods that track the centerline and the radius of the vessel [84, 88, 177], and methods that evolve implicit surfaces [116, 186].

The second group of methods first finds the centerline and then segments the vessel. A number of these methods uses the extracted centerline to segment the vessel with thresholding based on the image intensities on the centerline [81, 129] or by finding multiple minimal cost paths along the vessel boundary in curved planes constructed with the centerline [101, 152].

Most of the published coronary segmentation methods have been evaluated visually. Although a large body of centerline extraction methods have been quantitatively evaluated (see Chapter 2), to the best of our knowledge only Li et al. [88], Yang et al. [186], and Wesarg et al. [75, 178] have evaluated their segmentation method quantitatively. The quantitative evaluation in these papers is done with the Dice coefficient [88], the average and maximum contour distance [186], and by assessing the performance of the method for calcium and stenosis detection [75, 178].

In this chapter we present a new semi-automatic coronary CTA lumen segmentation method. The method is based on graph cuts, with edge-weights depending on the intensity of the centerline, and robust kernel regression. A vessel centerline is used for initialization of the method. From recent work it has become clear that

automatic coronary centerline extraction can be achieved with high precision and robustness [105].

A second major contribution of this chapter is the quantitative evaluation of the method on 28 manually annotated coronary artery lumen boundaries from 12 patients. In this chapter we quantitatively evaluate our method with the Dice coefficient and the average and maximum contour distance, the measures used by Li et al. [88] and Yang et al. [186].

4.2 Problem formulation

Large CT intensity gradients can be observed on the boundary of the coronary lumen in CTA, while the CT intensity within the lumen varies smoothly. Therefore the problem of coronary lumen segmentation is similar to many image segmentation problems: find the strongest edge surrounding an area with relatively similar intensities. Formalizing such a problem quickly leads to balancing a gradient and an intensity term, while often the intensity term should only be used to prevent the segmentation of structures with very different intensities.

Because we segment the lumen given a centerline, we can tailor this approach to our task: find the strongest edge surrounding areas with intensities locally similar to the centerline intensity, while not segmenting areas with intensities dissimilar to the centerline intensity. The intensity information should only be used to steer the segmentation towards the regions with appropriate intensity values; the gradient information should be used to accurately detect the border.

An additional application specific constraint that we incorporate is that we aim to segment the vessel that contains the centerline; side branches of this vessel should not be segmented. This is specifically important for subsequent quantification of the degree of stenosis in a coronary artery. The surface should interpolate the boundary of the vessel of interest and not take into account the image information arising from the side-branch.

4.3 Method

In view of the above, we propose a two step approach for segmenting the coronary lumen given a centerline:

1. **Segment the lumen** using the strong edge and similar intensity prior. This is done by solving a Markov Random Field with image terms locally depending on the intensity of the centerline.

2. **Remove falsely segmented regions not belonging to the vessel of interest** using the fact that the segmented lumen should not contain any holes, the surface should be smooth, and side-branches should not be segmented. This

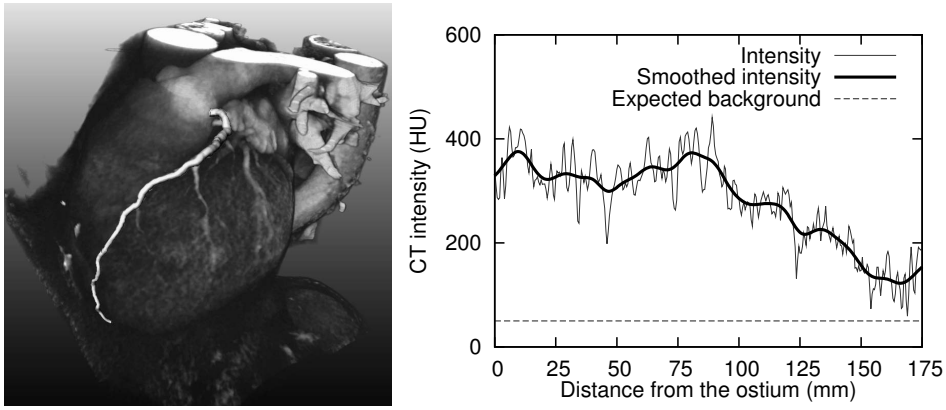

(a) (b)

Figure 4.1 (see page 147 for a color-version): *(a) A 3D rendering of a cardiac CTA dataset with in yellow a manually annotated Left Anterior Descending (LAD) Coronary Artery. (b) A graph of the CT intensities I_x along the centerline of the LAD, a graph of the intensities after Gaussian kernel regression $\hat{I}_x$ and the expected background intensity I_{bg} (see section 4.3.1).*

is done by robust kernel regression on a cylindrical parameterization of the lumen boundary.

4.3.1 Step 1: Segmenting the lumen with a Markov Random Field

In this first step we aim to find an optimal binary voxel labeling of the lumen and background. We do this by formalizing a binary Markov Random Field (MRF), which is solved using graph cuts [23, 25, 78].

A labeling $f = \{f_x | x \in \mathcal{X}\}$, with $f_x = \{0, 1\}$, is determined that has the maximum a posteriori probability given the CTA image $I = \{I_x | x \in \mathcal{X}\}$, with $\mathcal{X}$ being the set of voxels in the image. A labeling $f_x = 1$ corresponds to a voxel being lumen and $f_x = 0$ corresponds to a voxel being background. Each voxel x is associated with a set of neighborhood voxels $\mathcal{N} = \{\mathcal{N}_x | x \in \mathcal{X}\}$. The MRF is solved by factorizing the posterior $Pr(f|I)$ as follows:

$$\begin{aligned} Pr(f|I) &\propto \quad Pr(I|f)Pr(f), \\ Pr(f|I) &\propto \left(\textstyle\prod_x Pr(I_x|f_x)\right) Pr(f), \end{aligned} \tag{4.1}$$

with (see e.g. [25]):

$$Pr(f) = \exp\left(-\sum_x \sum_{y \in f_{\mathcal{N}_x}} \omega_{x,y}(1 - \delta(f_x - f_y))\right) \tag{4.2}$$

and subsequently rewriting it to the following energy functional that needs to be minimized:

$$E(f) = \sum_x -\log(Pr(I_x|f_x)) + \sum_x \sum_{y \in f_{\mathcal{N}_x}} \omega_{x,y}(1 - \delta(f_x - f_y)), \tag{4.3}$$

with $Pr(I_x|f_x = 1)$ and $\omega_{x,y}$ defined for our application below.
The minimization of the energy functional can be done with graph cuts [25]. In this approach a graph is constructed where each node corresponds to a voxel x. Each voxel is connected (with *t-links*) to two additional nodes denoted respectively as 'source' and 'sink'. A weight of $\omega_s = -\log(1 - Pr(I_x|f_x = 1))$ is assigned to the source connection and a weight of $\omega_t = -\log(Pr(I_x|f_x = 1))$ is assigned to the sink connection.

Each voxel is also connected (with *n-links*) to 34 neighboring voxels $y \in \mathcal{N}_x$ and weights of $\omega_{x,y}$ are assigned to these connections (see section 4.3.3 for a description of the 34-connected neighborhood model $\mathcal{N}_x$). We subsequently find a cut in the graph with minimal summed weight that separates the source and the sink. This cut corresponds to the global minimum of $E(f)$ [23, 25, 78]. Voxels still connected to the source are labeled as lumen.

4.3.1.1 Image dependent voxel likelihood

We let the likelihood of a voxel being lumen $Pr(I_x|f_x = 1)$ depend on the difference between the voxel intensity and a local estimate of the lumen intensity, and a local estimate of the intensity difference between the lumen and the surrounding tissue. For notational purposes we ignore the dependency on these local estimates in $Pr(I_x|f_x = 1)$.

The local lumen intensity is estimated with a Nadaraya-Watson estimator [115]: image intensities $I_{x=c(s)}$ are sampled along the centerline, with $x = c(s)$ a position on the centerline and s the geodesic length from the start of the centerline c. This 1D function is smoothed with a Gaussian function with standard deviation σ_c to obtain a local estimate $\hat{I}_x$ of the lumen intensity.

Background tissue is modeled with a fixed intensity of I_{bg}, resulting in an estimated difference between the lumen and the background of $\hat{D}_{x'} = \hat{I}_{x'} - I_{\text{bg}}$ (see Figure 4.1(b)). Here x' denotes the position on the centerline closest to x, $D_x = |I_x - \hat{I}_{x'}|$ is the absolute difference between the intensity of a voxel I_x and the local intensity estimate, and $\hat{D}_x = \hat{D}_{x'}$ describes the estimated local contrast in the image.

Using these local estimates we formalize the likelihood of a lumen voxel given its intensity (and the intensities on the centerline) with a smooth step function (see also Figure 4.2(a)):

$$Pr(I_x|f_x = 1) = -0.5\left(0.75 - 0.25\text{erf}\left(\frac{D_x - T_{\text{in}}}{\sigma_i}\right)\right)\left(\text{erf}\left(\frac{D_x - T_{\text{out}}}{\sigma_i}\right) - 1\right),$$

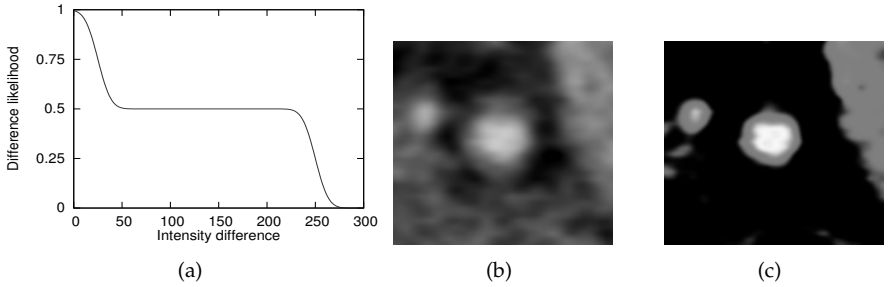

(a) (b) (c)

Figure 4.2: *(a)* $Pr(I_x|f_x = 1)$ *with* $T_{in} = 25$, $T_{out} = 250$, *and* $\sigma_i = 15$. *(b) A randomly selected cross-sectional image. (c) The application of* $Pr(I_x|f_x = 1)$.

with $T_{\text{out}} = \lambda \hat{D}_x$. It can be appreciated that this function has two soft thresholds; differences D_x smaller than T_{in} correspond to a high lumen likelihood and differences higher than T_{out} correspond to a low lumen likelihood. The parameter T_{in} is user-defined and T_{out} depends locally on the contrast of the vessel with its background tissue.

By setting T_{in} relatively low and λ relatively high we make sure that the voxel term is only used to steer the segmentation towards the regions with appropriate intensity values; the edge term is used to accurately find the border in this region. In Figure 4.2 we show an example of $Pr(I_x|f_x = 1)$ applied to a randomly selected cross-sectional image.

4.3.1.2 Edge term

For the edge term we use a Gaussian function of the squared gradient magnitude on the boundary between voxel y and x $|\nabla I|(y,x)$. A high gradient magnitude corresponds to a high probability of a label switch between lumen and background:

$$Pr(f_x \neq f_y) \propto 1 - \exp\left(\frac{-|\nabla I|^2(x,y)}{2\sigma_g^2}\right). \tag{4.4}$$

Therefore we assign the following weight to a label switch between voxel x and y:

$$\omega_{x,y} = -\log\left(1 - \exp\left(\frac{-|\nabla I|^2(x,y)}{2\sigma_g^2}\right)\right). \tag{4.5}$$

4.3.1.3 Segmentation after the first step

In this first step of the algorithm a binary segmentation of the lumen is obtained. This segmentation is close to the optimal solution but it contains several false pos-

itives, corresponding to side-branches of the vessel of interest and other contrast-filled regions, regions with a similar intensity as lumen while being (blurred) calcium, and some mis-segmentations caused by image artifacts. The result of step 1 can be seen in Figure 4.3(b).

4.3.2 Step 2: Removing outliers from the segmentation.

In the second step of the algorithm we detect and remove regions not belonging to the vessel of interest with an iterative weighted kernel regression approach.

The segmented lumen is parameterized with cylindrical coordinates $r(\phi, s)$ by calculating the intersection of the boundary of the segmented lumen with a series of radial lines perpendicular to the centerline. We do this by extracting values from the masks obtained in step 1 with normalized Gaussian interpolation [77] and calculating the intersection point by linear interpolation (see Figure 4.3(c)).

Outliers are then removed from this parameterization by a simplified version of the reweighted kernel regression approach reviewed by Debruyne et al. in [36]. To each point $r(\phi, s)$ in the parameterization we iteratively assign a weight $w(\phi, s)$ describing the belief in this point. We use a Gaussian loss function, resulting in:

$$w_(\phi, s)^t = exp\left(\frac{-(r(\phi, s)^t - r(\phi, s)^{t=0})^2}{2\sigma_r^2}\right) \tag{4.6}$$

These weights are then used to improve the estimation of $r(.)$:

$$r_(\phi, s)^{t+1} = \frac{\sum_{\phi', s'} G_{\sigma_\phi, \sigma_s}(\phi' - \phi, s' - s) w(\phi', s')^t r(\phi', s')^{t=0}}{\sum_{\phi', s'} G_{\sigma_\phi, \sigma_s}(\phi' - \phi, s' - s) w(\phi', s')^t} \tag{4.7}$$

with $G_{\sigma_\phi, \sigma_s}(.)$ a 2D Gaussian kernel with standard deviations in the angular and longitudinal direction of respectively σ_ϕ and σ_s. This process is repeated until convergence ($t = T$). See Figures 4.3(d) and 4.3(e) for an example of $r(\phi, s)^T$.

4.3.3 Implementation

All segmentations are carried out in a region of 7.5 mm (approximately 50% larger than the maximum radius of a coronary arteries) around the centerlines to reduce computation time and memory requirements. Cross-sectional images of 128 $\times$ 128 pixels are created every 0.5 mm along the centerline (resulting in a voxelsize of 0.1 $\times$ 0.1 $\times$ 0.5 mm^3).

We use a 34 connected neighborhood region $\mathcal{N}_x$, with 26 connections corresponding to all the neighborhood connections in a 3 $\times$ 3 $\times$ 3 region and 8 connections corresponding to the 8 possible knight-moves in the cross-sectional plane. Using these knight-moves significantly improved the smoothness of the resulting segmentation (see also [24]).

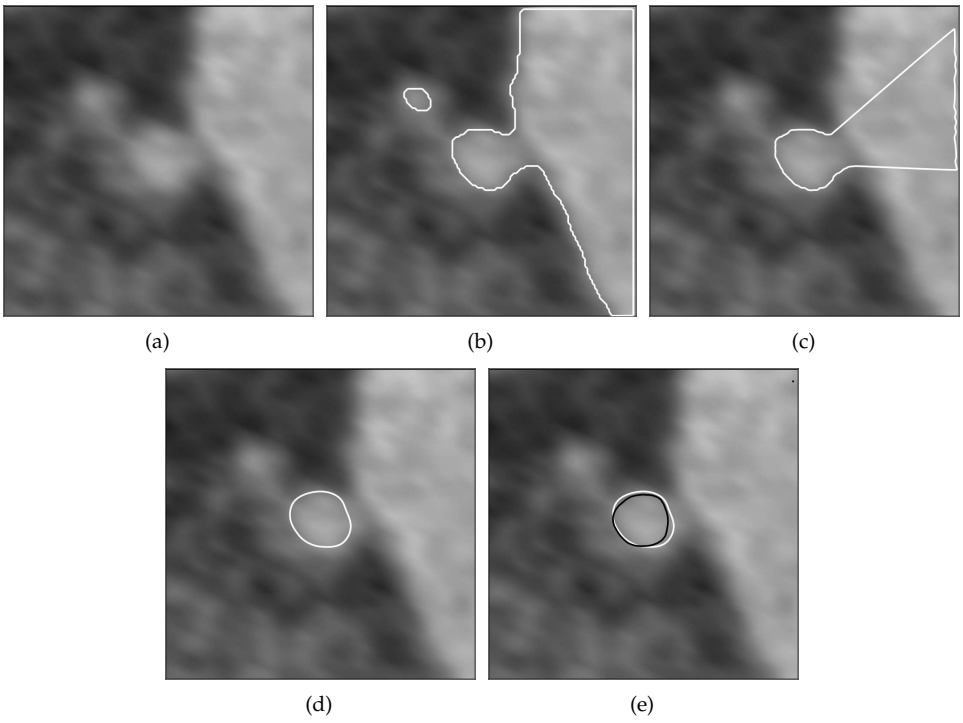

Figure 4.3: *A representative example (Dice=0.85, ASSD=0.18mm, AMCD=0.42mm, see section 4.4.3). (a) shows a cross-sectional slice of the input image, (b) shows the result after step 1, (c) shows the initialization of step 2, (d) shows the final segmentation, and (e) shows the automatic segmentation (white) together with the reference standard (black).*

4.3.4 Parameters

The parameters were empirically chosen by the authors; no extensive parameter-tuning was performed, with the exception of σ_ϕ and σ_s. These two parameters were tuned on one of the 28 vessels.

The following parameter settings were used for all the experiments in this chapter: σ_c =2 mm, σ_i =15 HU, T_{in} = 25 HU, T_{bg} = 50 HU, λ = 0.75, σ_g =15 HU, σ_r =0.1 mm, σ_ϕ = 0.2 rad, and σ_s = 1 mm.

4.4 Quantitative evaluation

The method is quantitatively evaluated by comparing the segmentations with manually annotated lumen surfaces of 28 coronary arteries. The coronary arteries were

segmented using manually annotated centerlines. These centerlines were also used to manually annotate the lumen boundary (see section 4.4.2).

4.4.1 Data

The cardiac CTA data of twelve patients was used for this study. Two main coronary arteries (RCA, LAD or LCX) were annotated in each dataset and an additional side-branch was annotated in four of the datasets. The observer annotated in total 8 RCAs, 8 LADs, 8 LCXs, and 4 side-branches.

The twelve CTA datasets were acquired in the Erasmus MC, University Medical Center Rotterdam, The Netherlands. The datasets were randomly selected from a series of patients who underwent a cardiac CTA examination between June 2005 and June 2006. The datasets were acquired with a 64-slice CT scanner and a dual-source CT scanner (Sensation 64 and Somatom Definition, Siemens Medical Solutions, Forchheim, Germany). The datasets were reconstructed using a sharp (B46f) kernel or a medium-to-smooth (B30f) kernel.

4.4.2 Manual annotation

One observer annotated the coronary arteries from the coronary ostium (i.e. the point where the coronary artery originates from the aorta), until the most distal point where the artery is still distinguishable from the background. On average the 28 coronary arteries were 147 mm long.

A tool was specifically designed for the manual annotation. The tool was developed in the free software package MeVisLab (`http://www.mevislab.de`) and has a workflow similar to the automatic approach used by Marquering et al. in [101]. After annotating a centerline the user annotates the lumen outlines longitudinally with B-splines in curved planar reformatted images created at three different angles. These curves are then intersected with planes perpendicular to the centerline spaced regularly with a distance of 1 mm. In the second annotation step the points resulting from the intersections are connected with closed B-splines to form initial contours in the cross-sectional planes. These contours can then be modified by the observer, resulting in the final annotation.

4.4.3 Evaluation measures

The Dice measure, the average symmetric surface distance (ASSD), and the average maximum contour distance (AMCD) are used to quantify the difference between the manual annotations and the automatically extracted lumen surface (see [88] and [186] for the application of these measures on coronary artery segmentation evaluation).

Table 4.1: *An overview of the quantitative results. The Dice measure, average symmetric surface distance (ASSD), and the average maximum contour distance (AMCD) are reported for the complete vessel and for the first 90 mm of the vessel.*

	Complete vessel			First 90 mm		
	Dice	**ASSD**	**AMCD**	**Dice**	**ASSD**	**AMCD**
LAD (8)	0.852	0.188	0.378	0.860	0.212	0.445
LCX (8)	0.831	0.250	0.514	0.862	0.229	0.495
RCA (8)	0.854	0.221	0.466	0.860	0.241	0.542
Side branch (8)	0.846	0.222	0.479	0.858	0.205	0.456
All (28)	0.846	0.220	0.456	0.859	0.226	0.491

The Dice measure represents the fraction of the volume of the overlap of the two segmentations and the average volume of the two segmentations:

$$\text{Dice} = \frac{2 \times \text{TP}}{2 \times \text{TP} + \text{FP} + \text{FN}} \tag{4.8}$$

The ASSD measure is determined by calculating for each point on both segmentations the distance to the closest point on the other segmentation and averaging these distances. The AMCD distance is calculated by averaging the maximum of all these distances per cross-sectional contour.

4.4.4 First 90 mm of the vessel

Correctly segmenting the complete coronary shows the capability of the method to segment very small vessels, but segmenting the distal part of the vessel is not always needed in clinical practice, because disease occurs for more than 95% in the first 90 mm of the coronary arteries [65]. Therefore we also evaluate the capability of the method to segment the first 90 mm of the vessel.

4.5 Results

Table 4.1 shows the quantitative results. Figure 4.4 shows a series of cross-sectional images with the manual annotation and an intersection of the automatic segmentation. Figure 4.5 shows two segmentations in 3D, color-coded with the distances to the reference standard.

4.6 Discussion

We have presented a new CTA coronary lumen segmentation method, which uses a vessel centerline for initialization. The method accurately aligns the boundary of

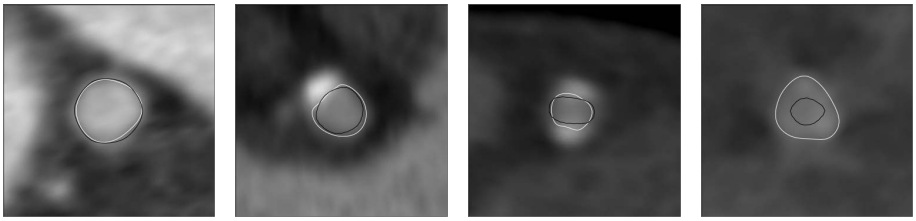

(a) [0.95,0.11,0.27] (b) [0.89,0.19,0.38] (c) [0.82,0.18,0.41] (d) [0.36,0.75,2.04]

Figure 4.4: *Cross-sectional segmentation examples (15$\times$ 15 mm^2) (in white) with corresponding reference standard (in black) and measures [Dice, ASSD in mm, AMCD in mm]. The error in (d) was caused by the false segmentation of a stent.*

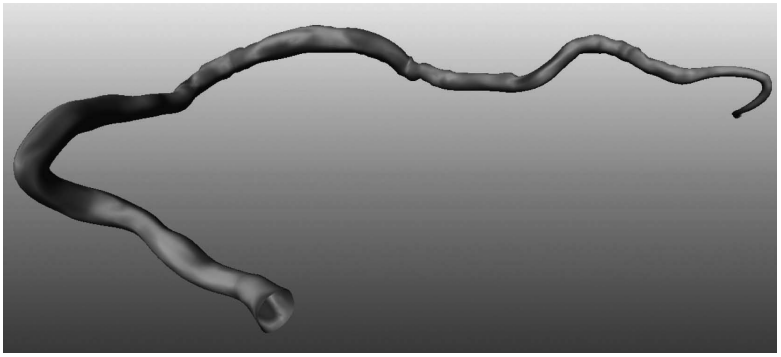

Figure 4.5 (see page 147 for a color-version): *A 3D example of a coronary segmentation color-coded with the distance to the reference standard. Red corresponds to the segmentation being locally 0.5 mm larger than the reference standard, green corresponds to a perfect fit, and blue corresponds to a 0.5 mm under-segmentation.*

the segmentation with the strongest edge surrounding areas with intensities that are locally similar to the centerline intensity, while not segmenting intensities that are dissimilar to the centerline intensities. A successive robust regression step is used to remove outliers from the segmentation.

The method is quantitatively evaluated on 28 vessels in 12 cardiac CTA datasets. The average symmetric surface distance between the method and the manual reference is 0.22 mm, and the average maximum contour distance is 0.46 mm, with a mean voxel size of $0.32 \times 0.32 \times 0.40$ mm^3. Furthermore, the method obtains an average Dice coefficient of 0.85. As a rough reference one could compare these numbers to the quantitative results obtained by Li et al. and Yang et al. Li et al. obtain a Dice coefficient of 0.58 [88] and Yang et al. [186] obtain an ASSD of 0.37 mm and an AMCD of 1.36 mm. However, it should be noted that these two methods are

evaluated in different patients and the datasets are most probably acquired with a different type of CT scanner. Furthermore, our method is initialized with a centerline in contrary to the methods of Li et al. and Yang et al. To objectively compare different coronary artery segmentation methods a standardized evaluation framework for coronary artery segmentation methods should be developed, e.g. using a similar approach as the coronary artery tracking evaluation framework [105].

A limitation of this study is that our algorithm uses the same centerlines as were used by the observers for annotating the coronary lumen, which may bias the results. In the future we will investigate the effect of perturbations of the centerline on the segmentation results.

Using the presented method for lumen segmentation can already reduce the amount of user-interaction with a factor of 10 (the manual annotation of a centerline takes approximately 5 minutes, while annotating the coronary lumen outline can take up to 50 minutes). A further reduction seems feasible, as results from the coronary artery tracking evaluation framework [105] show that semi-automated and fully automatic coronary artery centerline tracking methods that achieve high robustness and accuracy (comparable to inter-observer variability) are available. In the future we will also evaluate our method with (semi-)automatically extracted centerlines.

Another limitation is that at this moment we only have coronary lumen annotations of one observer. Manual annotations by multiple clinical experts are currently planned, and in future work we will therefore also relate the performance of the method to the inter-observer variability. Finally, we will investigate the possibility to quantify clinically relevant measures (such as the degree of stenosis) with the proposed method.

4.7 Conclusion

A high-precision coronary lumen segmentation method is presented. The method is based on graph cuts and robust kernel regression and segments the coronary lumen given a centerline. The method has been successfully applied for the segmentation of 28 coronary arteries. A quantitative evaluation showed that the method was able to segment the coronary arteries with high accuracy, compared to manually annotated segmentations.

Chapter 5

Supervised Shape-based Coronary Segmentation

Based on:

M. Schaap, T. van Walsum, M. de Bruijne, L. Neefjes, C. Metz, E. Capuano, W.J. Niessen. *Robust Shape Regression for Supervised Vessel Segmentation and its application to Coronary Segmentation in CTA, Submitted, 2010.*

This chapter presents a vessel segmentation method which learns the geometry and appearance of vessels in medical images from annotated data and uses this knowledge to segment vessels in unseen images.

Vessels are segmented in a coarse-to-fine fashion. First the vessel boundaries are estimated with multi-variate linear regression using image intensities sampled in a region of interest around an initialization curve. Subsequently the position of the vessel boundary is refined with a robust non-linear regression technique using intensity profiles sampled across the boundary of the rough segmentation and information about plausible cross-sectional vessel shapes.

The method was evaluated by quantitatively comparing segmentation results to manual annotations of 229 coronary arteries. On average the difference between the automatically obtained segmentations and manual contours was smaller than the inter-observer variability, which is a strong indicator that the method outperforms manual annotation. The method was also evaluated by using it for centerline refinement on 24 publicly available datasets of the Rotterdam Coronary Artery Evaluation Framework. Centerlines are extracted with an existing method and refined with the proposed method. This combination is currently ranked 2nd out of 17 evaluated methods. An additional qualitative expert evaluation in which 250 automatic segmentations were compared to manual segmentations showed that the automatically obtained contours were rated on average better than manual contours.

5.1 Introduction

Accurately quantifying the vessel lumen in medical images is beneficial for the diagnosis and treatment planning of a wide range of cardiovascular diseases. This chapter presents a segmentation method which learns the geometry and appearance of the vessel lumen in medical images from annotated data and uses this knowledge for the segmentation of an unseen vessel.

We demonstrate the applicability of the proposed method to a challenging vessel segmentation task: the segmentation of coronary arteries in Computed Tomography Angiography (CTA) data. Furthermore, we demonstrate that the proposed method can also be used for the refinement of inaccurate central lumen lines. Coronary artery lumen segmentation and central lumen line extraction in CTA are both challenging because of the small size of the vessels compared to the image resolution, the presence of (motion) artifacts in the data, the presence of neighboring structures with similar intensities as the vessel, and pathologies such as severe stenoses and calcifications. Figure 5.1 shows a cardiac CTA image and examples of coronary lumen segmentations.

Coronary artery lumen segmentation methods can be divided into two categories: methods that first determine a central lumen line (manually or automatically) and then segment the vessel of interest, and methods that segment the vessel in one pass. Coronary segmentation methods that use a centerline for initialization incorporate thresholding techniques [81, 129], minimal cost path approaches in longitudinal reformatted images [101], or minimal surface techniques [143, 152]. Methods that segment the vessels in one pass can be divided into methods that use region-growing or a combination of different morphological operators [21, 22, 96, 104], methods that use vessel tracking in 4D to directly find the centerline and the radius of the vessel [44, 84, 88, 89, 144, 177], and methods that evolve implicit surfaces [116, 186]. The highest accuracy was reported in previous work from our group [143]. A surface distance to manual annotations of approximately half the size of a voxel was obtained with a method based on graph cuts and robust regression, but an accurate centerline was needed as initialization. Moreover, it was observed that low contrast between the inside and outside of the lumen sometimes caused errors in the segmentation of vessels with a very small radius.

All the existing coronary artery segmentation methods as well as almost all other vessel segmentation methods are unsupervised; vessels are segmented using a variety of hand-tailored assumptions. In a recent overview of Lesage et al., vessel segmentation methods are classified according to their assumptions on geometry and appearance of the vessel [85]. Examples of geometric assumptions are the vessel's tubularity (e.g. [9, 144]), and smoothness of the centerline (e.g. [95]). Examples of appearance assumptions are homogeneity of the vessel's image intensities (e.g. [21, 143, 160]), strong image gradients on the vessel boundary (e.g. [15, 95, 101]), and a relative high intensity of the vessel compared to its surround-

ing (e.g. [9, 44, 144]). Although these assumptions hold for healthy vessels, vessel disease affects the morphology and the appearance of vessels in medical images, which often breaks these assumptions. For example, calcified plaque in the vessel wall makes the vessel lumen in CTA appear darker than its neighboring tissue.

In this work we propose a supervised approach to vessel segmentation which replaces the hand-tailored assumptions with models learnt from annotated training data. In a coarse-to-fine robust shape regression approach we jointly fit a set of local shape models to image data in order to segment the vessel of interest. Both the shape and appearance information are learnt from annotated training data with respectively local point distribution models and a non-linear boundary intensity model. By incorporating shape information in the segmentation process the resulting boundary does not necessarily have to be supported by image information; sections of the image with strong image information (e.g. sharp edges) can steer the segmentation with weak image information (e.g. low contrast).

Appearance information has been used previously in a supervised fashion to detect and segment vessels in 2D and 3D vascular images [130, 153, 193]. The combination of statistical shape and appearance models of tubular anatomical structures has been applied to the segmentation of large vessels, such as the abdominal aorta in CTA [26, 35]. However, the combination of shape-based boundary delineation of small vessels with supervised appearance information in 3D images, as proposed in this chapter, has, to the best of our knowledge, not been addressed previously. Moreover, in the cases where elongated structures are segmented by fitting a series of local shape models to the image data [35, 68], these local models are generally fitted sequentially, whereas we propose to fit all local shapes simultaneously to the image data in one global optimization strategy.

The next section explains the method, followed by a description of the experiments in Section 5.3, in which the method is evaluated both qualitatively and quantitatively on a large number of datasets. In total 107 cardiac CTA datasets were used for the evaluation. The results of these experiments are discussed in Section 5.4 and the chapter is concluded by discussing possible improvements of the proposed method in Section 5.5.

5.2 Method

The proposed vessel segmentation method is initialized with an approximate centerline, which can, for example, be found with an existing centerline extraction technique or by manually clicking a few points in the vessel (see Figure 5.1). Previous work showed that coronary centerline extraction can be done robustly with a variety of existing techniques [142].

In the next two sections we define respectively the local coordinate system and region of interest in which the vessel is segmented and our vessel shape represen-

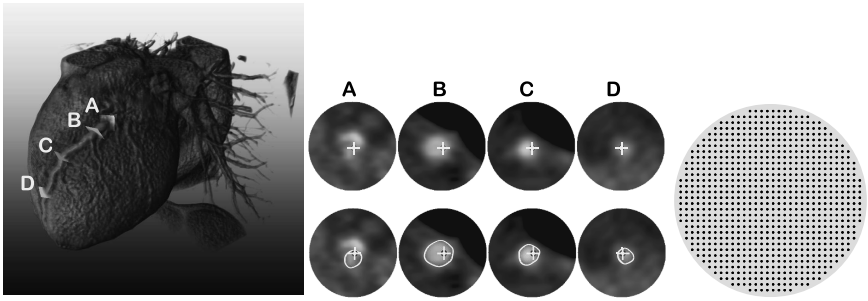

Figure 5.1 (see page 148 for a color-version): *Left: a 3-dimensional visualization of a cardiac CTA image with in red a coronary artery. Middle: four circular images sampled cross-sectionally to the coronary artery (sample radius=7.5 mm and sample distance=0.5 mm.) overlaid with the position of an initialization centerline and manually annotated coronary lumen contours. Right: The intensity sample pattern, as described in Section 5.2.1.*

tation. Subsequently we explain the segmentation process. The vessel is segmented in two steps in a coarse-to-fine fashion. First the vessel is segmented with multivariate linear regression using image intensities sampled in a region of interest around the initialization centerline (see Section 5.2.3). Subsequently the position of the vessel boundary is refined with a robust non-linear regression technique using image intensities on the boundary of the rough segmentation (see Section 5.2.4).

5.2.1 Local coordinate systems and region of interest

The segmentation process is formulated in a set of local coordinate systems that are defined according to the initialization centerline. First, the initialization centerline is sampled equidistantly, resulting in a set of L points $\{\mathbf{p}_z\} : z \in \{1, .., L\}$. These points are used to define local 2D coordinate systems at each centerline point using a rotation minimizing frame (RMF) technique [171]. For this, the normalized tangent vector $\mathbf{t}_z = (\mathbf{p}_{z+1} - \mathbf{p}_z)/\|\mathbf{p}_{z+1} - \mathbf{p}_z\|$ is calculated for each point and two normalized orthogonal vectors $\{\mathbf{v}_z, \mathbf{w}_z\}$ are determined perpendicular to these vectors. The first frame $\{\mathbf{v}_0, \mathbf{w}_0\}$ is chosen arbitrarily in the plane perpendicular to $\mathbf{t}_0$ and all other frames are calculated recursively by rotating $\mathbf{v}_z$ and $\mathbf{w}_z$ with the angle between $\mathbf{t}_z$ and $\mathbf{t}_{z-1}$:

$$\begin{aligned} \mathbf{v}_z &= \mathrm{Rot}(\mathbf{t}_{z-1} \times \mathbf{t}_z, \mathrm{acos}(\mathbf{t}_z \cdot \mathbf{t}_{z-1})) \cdot \mathbf{v}_{z-1} \\ \mathbf{w}_z &= \mathbf{v}_z \times \mathbf{t}_z, \end{aligned} \tag{5.1}$$

where $\mathrm{Rot}(\mathbf{x}, \alpha)$ is a rotation matrix describing a rotation with angle α around axis $\mathbf{x}$.

In each of the cross-sectional planes R, image intensities are sampled at regular grid positions in a circular region of interest, resulting in a set of image intensities $\mathbf{I} = \{i_{z,r}\} : z \in \{1,..,L\}, r \in \{1,..,R\}$ (see Figure 5.1). The algorithm uses these image intensities to segment the vessel.

5.2.2 Shape representation

It is challenging to model the shape and appearance of vessels with a statistical shape model because of the wide variety in shapes and appearances. Vessels can have different lengths, and vessel disease and bifurcations can be present at a variety of locations in the vasculature. Moreover, landmark correspondence between annotated shapes is required for statistical shape models, which is difficult for vessels with varying length. Even with a model representation that can represent vessels in all their variety, modeling this highly diverse information in one global model would require an impractical large amount of training data. We circumvent this shape modeling problem by learning the relevant information about the vessel's appearance and shape with a combination of local shape models, instead of one global model.

In contrast to much of the existing work on shape-based image segmentation, in this work we use the term 'shape' to refer to a set of points without invariance to the similarity transformation, i.e. we consider a shape different if it is scaled, moved or rotated. The complete shape of the vessel $\mathbf{S} = \{\mathbf{s}_{z,a}\} : z \in \{1,..,L\}, a \in \{1,..,D\}$ is represented by a set of 2-dimensional landmarks $\mathbf{s}_{z,a}$ in the cross-sectional planes described by $\{\mathbf{v}_z, \mathbf{w}_z\}$, with z the position along the centerline and a the cross-sectional landmark index. During training, landmarks are sampled on the (manually annotated) 3D surface of the lumen. Correspondence between the landmarks is obtained by sampling them according to fixed angles from the center of the cross-sectional contours. First, the center of mass of the intersections between each cross-sectional plane and the annotated 3D surface of the lumen is calculated:

$$\mathbf{c}_z = \mathbf{p}_z + \mathbf{P}_z \frac{\sum_i \|\mathbf{s}'_{z,i} - \mathbf{s}'_{z,i+1}\|(\mathbf{s}'_{z,i} + \mathbf{s}'_{z,i+1})}{2\sum_i \|\mathbf{s}'_{z,i} - \mathbf{s}'_{z,i+1}\|}, \tag{5.2}$$

with $\mathbf{s}'_{z,i}$ a point on the cross-sectional curve $\mathbf{S}'_z = \{\mathbf{s}'_{z,i}\}$, $\mathbf{P}_z = [\mathbf{v}_z\ \mathbf{w}_z]$ a matrix containing the axes of the local coordinate frame, and $\mathbf{c}_z$ the resulting center position. Subsequently landmarks $\mathbf{s}_{z,a}$ are obtained at the intersection between the annotated surface and rays $p(r)$ from the central point:

$$p(r) = \mathbf{c}_z + r\mathbf{P}_z \begin{pmatrix} \cos(\frac{2\pi d}{D}) \\ \sin(\frac{2\pi d}{D}) \end{pmatrix}, \tag{5.3}$$

with $r \geq 0$.

During the segmentation the vessel is described with L overlapping local shapes, each describing the shape of $2W + 1$ adjacent cross-sectional shapes. Each local shape $\mathbf{S}_x$, with $x \in \{1, .., L\}$, is composed of the landmarks $\{\mathbf{s}_{z,a}\}_{|z-x|\leq W}$ as follows:

$$\mathbf{S}_x = (\mathbf{s}_{x-W,1}^T, \mathbf{s}_{x-W,2}^T, \ \cdots, \mathbf{s}_{x+W,D-1}^T, \mathbf{s}_{x+W,D}^T)^T, \tag{5.4}$$

We use border clamping to obtain landmark coordinates or image intensity values at positions before the centerline ($z < 1$) and after the centerline ($z \geq L$). The thickness of each local shape was set to 5 ($W = 2$) to capture the first and second order variations of the cross-sectional shapes along the course of the vessel and $D = 32$ landmarks were sampled in each cross-section of the vessel (see Figure 5.2). More details about this and other parameter choices can be found in Section 5.3.3.

The dimensionality of the local shapes ($F = 2(2W + 1)D$) is reduced with a set of annotated training data (see section 5.3.3). The mean and modes of variation of this set of training shapes is calculated and used to represent the F-dimensional local shapes with an M-dimensional parameter vector β_x containing the first M principal components of the local shape:

$$\mathbf{S}_x \approx \bar{\mathbf{S}} + \Phi\beta_x, \tag{5.5}$$

with Φ an $F \times M$ dimensional matrix containing the first M eigenvectors of the covariance matrix of all annotated local shapes and $\bar{\mathbf{S}}$ their mean. The construction of Φ and $\bar{\mathbf{S}}$ is explained in detail in Section 5.3.3. During the segmentation, the parameter vector β_x, instead of the complete shape $\mathbf{S}_x$, is estimated. This forces the resulting shape to lie on the subspace spanned by the first M modes of variation.

5.2.3 Coarse shape estimation via Ridge Regression

In the first step of the vessel segmentation, the boundary of the vessel, described by the local shape parameters β_x, is estimated in the proximity of the true vessel boundary. Several regression approaches are possible in this coarse shape estimation step and we have chosen to use ridge regression (RR), a linear multivariate regression technique.

Because very bright calcium spots are sometimes present in the proximity of the lumen, the relation between the image intensities and the shape parameters β_x can be highly non-linear. In order to remove these nonlinearities we propose an application-specific pre-processing of the images. In this step we aim to make all the bright calcium spots neighboring the coronary lumen darker than the lumen intensity, thereby making the lumen appear as a bright spot in the image, which makes the intensity-shape relation more linear.

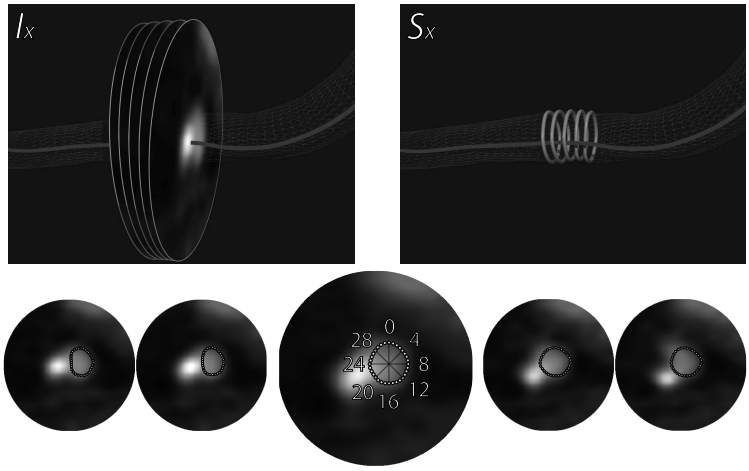

Figure 5.2 (see page 148 for a color-version): *A visualization of the intensity samples $\mathbf{I}_x$ and a local shape $\mathbf{S}_x$. Top left: a 3D surface of the lumen, the initialization centerline, and a tubular image patch $\mathbf{I}_x$ (see Section 5.2.3.2). Top right: the surface, initialization centerline, and a local shape $\mathbf{S}_x$ (see Section 5.2.2). Bottom: $\mathbf{I}_x$ overlaid with $\mathbf{S}_x$ and the rays (Equation 5.3) used for calculating the landmarks. The numbers indicate the cross-sectional landmark index d for $D = 32$. In the coarse segmentation $\mathbf{I}_x$ is used to estimate $\mathbf{S}_x$ with linear regression. During shape refinement all cross-sectional images of the vessel are used to estimate all cross-sectional shapes.*

5.2.3.1 Image pre-processing

Image pre-processing is done with an adaptive thresholding approach under the assumption that the majority of the cross-sections do not contain calcifications. In each cross-sectional image the maximum image intensity $l_z^{\max}$ is determined. This intensity is a good approximation of the lumen intensity, except at cross-sectional images with calcifications. A 1D function describing the lumen image intensity l_z along the centerline is then fitted through the measured intensities $l_z^{\max}$ with a robust re-weighted kernel regression approach [36]. The regression approach works iteratively; given a set of weights w_z^t and the measured image intensities the function is approximated with weighted Gaussian kernel regression:

$$l_z^t = \frac{\sum_{z'} \mathcal{N}(z'|z, \sigma_z) w_{z'}^{t-1} l_{z'}^{t-1}}{\sum_{z'} \mathcal{N}(z'|z, \sigma_z) w_{z'}^{t-1}} \tag{5.6}$$

and subsequently the weights are updated as follows:

$$w_z^t = \mathcal{N}(l_z^t | l_z^{\max}, \sigma_i), \tag{5.7}$$

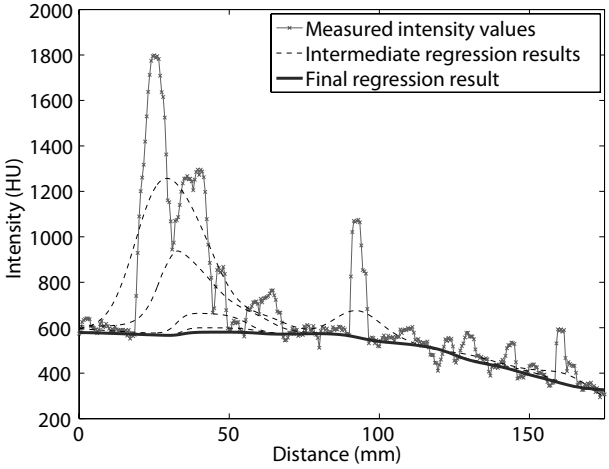

Figure 5.3: *Robust reweighted kernel regression results. The parameters described in section 5.3.2 were used to obtain these results. The red line with crosses shows the measured image intensity l_z^{max}, the black dashed lines the intermediate regression results l_z^t at $t = \{1, 5, 10, 15\}$, and the thick blue line shows the regression results at convergence ($T = 45$).*

with $\mathcal{N}(x|\mu, \Sigma)$ the value of the normal distribution with mean μ and covariance Σ at position x. A low weight implies that the point is considered to be an outlier. After convergence at $t = T$, image intensities significantly (more than c HU) above the lumen intensity are set to intensities c HU below the lumen intensity:

$$i^*_{z,q} = \begin{cases} l_z^{t=T} - c, & \text{if } i_{z,q} > l_z^{t=T} + c \\ i_{z,q}, & \text{else .} \end{cases} \tag{5.8}$$

The pre-processing step has three free parameters; σ_z, σ_i, and c, whose values are discussed in Section 5.3.2. Figure 5.3 shows an example of the estimation of the lumen intensities with the proposed robust reweighted kernel regression and Figure 5.4 shows examples of the image pre-processing results.

5.2.3.2 Ridge regression

Ridge regression is applied to learn a linear model that predicts the local shape parameters β_x from local image intensities. To describe the local appearance, we introduce $\mathbf{I}_x$, which consists of all image intensities within a $2W + 1$ slices thick region of interest, i.e. all intensities in the set $\{i_{z,q}\}_{|z-x| \leq W}$:

$$\mathbf{I}_x = (i_{x-W,0}\ , i_{x-W,1}\ , ...,\ i_{x+W,R-1}\ , i_{x+W,R})^T. \tag{5.9}$$

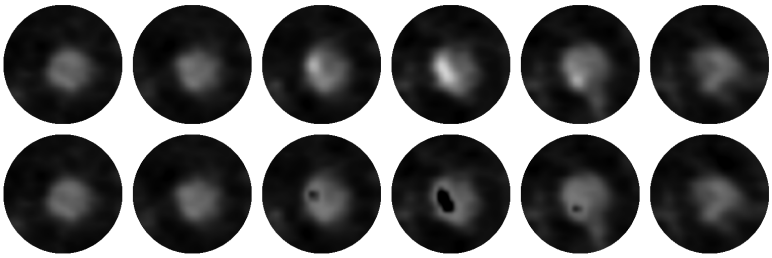

Figure 5.4: *Image pre-processing results. The top row shows the original images and the bottom row shows the images after pre-processing.*

An example of $\mathbf{I}_x$ can be found in Figure 5.2. The training data is composed of a set of N training samples: $\{\theta_i\}$, with $\theta_i = (\mathbf{I}_i, \beta_i)$ a tuple with a tubular image patch $\mathbf{I}_i \in \Re^Q$, with $Q = R(2W+1)$, and its corresponding shape parameters $\beta_i \in \Re^M$. Details about the training samples are given in Section 5.3.3. The goal is now to learn from the training data the (linear) relation between the tubular image patches and the annotated shapes.

Given an L^2 loss function, the $M \times Q$ dimensional matrix $\hat{H}$ that maps a tubular image patch to shape parameters can be found as follows:

$$\hat{H} = \underset{H}{\operatorname{argmin}} \sum_{i=1}^{N} \|\beta_i - H\tilde{\mathbf{I}}_i\|^2, \tag{5.10}$$

where $\tilde{\mathbf{I}}_i = \mathbf{I}_i - \bar{\mathbf{I}}$ denotes a centralized version of an image patch and $\bar{\mathbf{I}} = \frac{1}{N}\sum_{i=1}^{N} \mathbf{I}_i$ the mean of the training patches. This linear regression matrix $\hat{H}$ will be optimal for the training data. However, it does not necessarily have to be optimal for unseen data (i.e. overfitting can occur). Therefore, in ridge regression, the system is regularized by penalizing the sum of all squared regression coefficients:

$$\hat{H}_{\mathrm{RR}} = \underset{H}{\operatorname{argmin}} \sum_{i=1}^{N} \|\beta_i - H\tilde{\mathbf{I}}_i\|^2 + \lambda|H|^2. \tag{5.11}$$

After setting the derivative of the objective function in Equation 5.11 to zero we find:

$$\hat{H}_{\mathrm{RR}} = (\sum_{i=1}^{N} \beta_i \tilde{\mathbf{I}}_i^T)(\sum_{i=1}^{N} \tilde{\mathbf{I}}_i \tilde{\mathbf{I}}_i^T + \lambda I)^{-1}. \tag{5.12}$$

Given an unseen image $\mathbf{I}_x$ the coarsely estimated shape parameters β_x^* can now be determined as follows:

$$\beta_x^* = \hat{H}_{\mathrm{RR}}(\mathbf{I}_x - \bar{\mathbf{I}}). \tag{5.13}$$

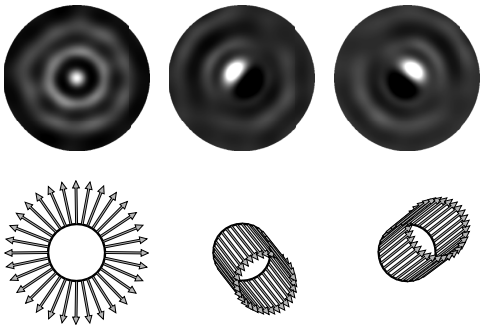

Figure 5.5: *A visualization of $\hat{H}_{RR}$ as used for the experiments described in Section 5.3. Top row: the center slice of the first three rows of $\hat{H}_{RR}$. Bottom row: the center slice of the mean shape and the corresponding modes of shape variation. The first mode of variation describes the cross-sectional size of the vessel, and the second and third describe its position relative to the initialization curve. In the regression, the contour points are moved in the direction of the modes of shape variation (the arrows in the bottom image) according to the filter response of the respective row of $\hat{H}_{RR}$ (top image). For visualization the regression coefficient images were resampled with cubic b-spline interpolation.*

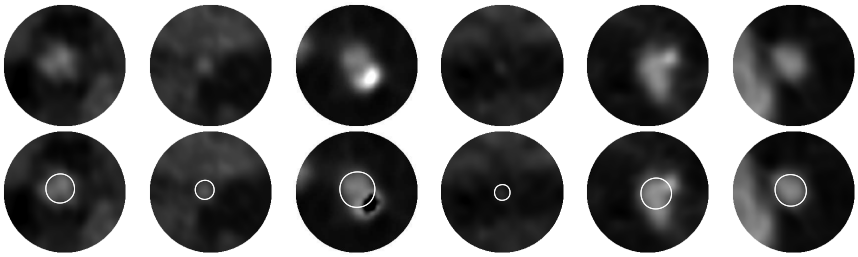

Figure 5.6: *Example results of the coarse segmentation step. Top row: input image. Bottom row: the segmentations after image pre-processing and ridge regression. These images show that the position and scale of the vessel are estimated correctly, but the linear model failed to estimate the complete shape accurately.*

And the first coarse estimate of the local shape $\mathbf{S}_x^*$ can be obtained with Equation 5.5:

$$\mathbf{S}_x^* = \bar{\mathbf{S}} + \Phi\beta_x^*. \tag{5.14}$$

Figure 5.5 visualizes the first three rows of $\hat{H}_{\mathrm{RR}}$ and Figure 5.6 shows several examples of the results of the coarse vessel segmentation step.

5.2.4 Robust supervised shape refinement

In the second and final step of the segmentation the landmark coordinates are optimized by shifting them to their most likely position given the image data across the boundary of the shapes estimated in the previous step, and the available training data. In this step the orginal image data is used without the pre-processing. The refinement step is further decomposed into two sub-steps: the independent estimation of the landmark positions based on boundary profiles and the statistical inference of a lumen segmentation from these landmark positions with information about plausible local vessel shapes.

5.2.4.1 Landmark estimation with KNN principal component regression

The location of each landmark is refined with local image information surrounding the coarsely estimated landmark using a non-linear regression approach, which learns the relation between the intensity profiles across an inaccurate boundary and the shift required to bring the boundary point in the correct position. From the estimated local shape parameters β_x^* we first obtain an estimate of each landmark position $\mathbf{s}_p^*$ by averaging the $2W+1$ estimates of the landmark position from the previous step:

$$\mathbf{s}_{z,a}^* = \frac{1}{2W+1} \sum_{\mathbf{S}_x^*:|x-z|\leq W} \Psi(z-x,a)\mathbf{S}_x^*, \tag{5.15}$$

with $\Psi(d,a)$ a $2 \times F$ matrix that extracts the 2D landmark at index $\{x+d,a\}$ from the F-dimensional vector $\mathbf{S}_x^*$. The elements of this matrix are defined as follows:

$$\Psi(d,a)_{i,j} = \begin{cases} 1, & \text{if } j = 2((d+W)D+a)+i \\ 0, & \text{else .} \end{cases} \tag{5.16}$$

After determining the initial position of each landmark, intensity profiles are sampled around the landmarks. The intensity profiles consist of G equidistantly sampled intensities in a region of interest (ROI) around the landmark $\mathbf{s}_p^*$. We use a 3D rectangular profile oriented in the direction of the normal of the estimated contour from the RR step. This local intensity information is used to improve the initial landmark estimation by searching in an annotated database with training data for similar profiles with known boundary shifts.

The training data $\mathcal{T} = \{(\mathbf{b}_j, \mathbf{d}_j)\}$ consists of boundary profiles $\mathbf{b}_j \in \Re^G$ and the associated in-plane shift of the landmark to its true position $\mathbf{d}_j \in \Re^2$. This training data is created by performing for all training samples the first steps of the algorithm (sampling image intensities around an inaccurate centerline, pre-processing the data, applying RR, and averaging the landmark estimates), and subsequently sampling for each landmark $\mathbf{s}_j^*$ the boundary profile $\mathbf{b}_j$, and calculating the difference $\mathbf{d}_j$ between the landmark $\mathbf{s}_j^*$ and its closest point on the annotated reference standard. See Figure 5.7 for an example of $\mathbf{b}_j$ and $\mathbf{d}_j$.

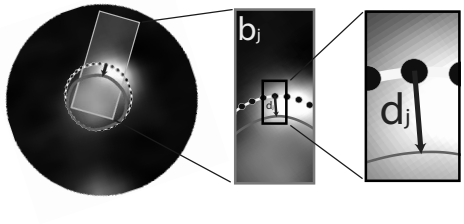

Figure 5.7 (see page 149 for a color-version): *A visualization of a boundary profile and displacement vector. Left: a cross-sectional image with the reference standard in solid orange, ridge regression result in black and white, the boundary profile* $\mathbf{b}_j$ *in transparent gray, and the displacement vector* $\mathbf{d}_j$ *in purple. Middle and right: zoomed-in versions of the left image.*

During the segmentation process we search for the K nearest neighbors (KNN) of the boundary profile $\mathbf{b}_{z,a}$, sampled at the landmark $\mathbf{s}^*_{z,a}$, in the training database by comparing the principal components of the training samples with those of the boundary profile. This is accomplished by finding the K nearest samples according to the following distance function:

$$D(\mathbf{b}_j, \mathbf{b}_{z,a}) = |\Theta(\mathbf{b}_j - \mathbf{b}_{z,a})|^2, \tag{5.17}$$

with Θ a $P \times G$ matrix containing the first P eigenvectors of the covariance matrix of the boundary profiles. With the resulting samples K refined landmark positions are created as follows:

$$\mathbf{u}^k_{z,a} = \mathbf{s}^*_{z,a} + \mathbf{d}^k_{z,a}, \tag{5.18}$$

with $\mathbf{s}^*_{z,a}$ the estimated position from the RR step and $\mathbf{d}^k_{z,a}$ the 2D shift obtained from the kth KNN sample. These K estimates of each landmark position will be combined into a segmentation of the vessel in the next step of the algorithm.

5.2.4.2 Robust Shape fitting

This section explains how the landmark estimates $\mathbf{u}^k_{z,a}$ are used to improve the estimates of the local shape parameters β^*_x. First each refined landmark position $\mathbf{u}^k_{z,a}$ is associated with a covariance matrix describing the expected error of the KNN estimation. In this work we have chosen to use a diagonal covariance matrix:

$$\Sigma^k_{z,a} = \sigma_{\text{knn}} I, \tag{5.19}$$

with σ_{knn} a free parameter of the method, describing the expected error for each KNN landmark estimate. The landmark estimates and covariance matrices together

describe $L \times D \times K$ sparse estimates (with associated uncertainty) of the lumen boundary:

$$\mathbf{U} = \{(\mathbf{u}_{z,a}^k, \Sigma_{z,a}^k)\} : z \in \{1,..,L\}, a \in \{1,..,D\}, k \in \{1,..,K\}, \tag{5.20}$$

where each tuple $(\mathbf{u}_{z,a}^k, \Sigma_{z,a}^k)$ gives a sparse estimate of the $2W+1$ shapes containing the respective landmark.

Similar to the robust regression approach in the image preprocessing step (see Section 5.2.3.1) we propose to use an iterative reweighted regression approach in this step. Here we employ this technique to statistically infer the complete shape from all the KNN landmark estimates and, at the same time, estimate a weight for each KNN landmark that describes how close to the true border the landmark was estimated. A low weight means the KNN estimate is classified as an outlier and a weight of 1 denotes that the landmark is assumed to be estimated correctly (i.e. it lies exactly on the resulting segmentation). Here we have, in contrast to the univariate case presented in Section 5.2.3.1, multiple measurements per value that needs to be estimated; each KNN estimate gives a sparse estimate of the shape parameters and $(2W+1) \times D \times K$ of these estimates are available per shape. We propose to estimate in each iteration the shape parameters with weighted linear regression from the corresponding landmark estimates. The proposed regression approach can therefore be seen as a robust iterative reweighted variant of the well-known local linear regression approach [55].

At iteration t the estimates $\mathbf{U}$ and the associated weights $w_{z,a}^{kt}$ are used to calculate the shape parameters $\hat{\beta}_x^t$ with weighted linear regression by minimizing the following weighted residual sum of squares (WRSS):

$$\text{WRSS}(\beta_x^t) = \sum_{(z,a)\in N(\beta_x)} \sum_{k=1}^{K} w_{z,a}^{kt} d(\mathbf{u}_{z,a}^k, \beta_x^t) \Sigma_{z,a}^{k^{-1}} d(\mathbf{u}_{z,a}^k, \beta_x^t), \tag{5.21}$$

with $N(\beta_x)$ the set of landmark indices that are part of the shape described by β_x and $d(\mathbf{u}_{z,a}^k, \beta_x)$ the 2D difference between a landmark estimate $\mathbf{u}_{z,a}^k$ and the position of that specific landmark described by the shape parameters β_x:

$$d(\mathbf{u}_{z,a}^k, \beta_x) = \mathbf{u}_{z,a}^k - \Psi(z-x, a)(\Phi\beta_x + \bar{\mathbf{S}}). \tag{5.22}$$

By setting the derivative of $\text{WRSS}(\beta_x^t)$ to zero we find:

$$\hat{\beta}_x^t = \left(\sum_{(z,a)\in N(\beta_x)} \sum_{k=1}^{K} w_{z,a}^{kt} \Phi^T \Psi^T \Sigma_{z,a}^{k^{-1}} \Psi\Phi \right)^{-1} \sum_{(z,a)\in N(\beta_x)} \sum_{k=1}^{K} w_{z,a}^{kt} \Phi^T \Psi^T \Sigma_{z,a}^{k^{-1}} (\mathbf{u}_{z,a}^k - \Psi\bar{\mathbf{S}}) \tag{5.23}$$

with $\Psi \triangleq \Psi(z-x,a)$ for notational purposes. After this weighted least-squares estimation new weights are calculated by assessing the quality of each landmark $\mathbf{u}_{z,a}^{k}$ in relation to the current boundary estimate:

$$w_{z,a}^{k\,t+1} = \frac{1}{2W+1} \sum_{\beta_x \in N(z)} \exp(-\frac{1}{2} d(\mathbf{u}_{z,a}^{k}, \beta_x^t) \Sigma_{z,a}^{k^{-1}} d(\mathbf{u}_{z,a}^{k}, \beta_x^t)), \tag{5.24}$$

with $N(z)$ the set of shape parameters that describe shapes containing landmarks in slice z.

The regression starts with $\beta_x^{t=0} = \beta_x^*$ from the RR step and all the shape parameters and weights are updated simultaneously. Convergence at $t = T$ is declared when $|\beta_x^t - \beta_x^{t-1}| < \epsilon$ for all β_x. In the experiments convergence was declared on average at approximately $t = 7$. After convergence the position of each landmark on the vessel boundary $\mathbf{s}_{z,a}$ can be determined with Equation 5.5 and 5.15 using the refined shape parameters $\beta_z^{t=T}$. Figure 5.8 shows several examples of the complete segmentation procedure.

5.3 Experiments

The proposed method was evaluated both quantitatively and qualitatively. The quantitative evaluation consisted of the assessment of the segmentation performance, by comparing segmentations of 229 coronary arteries in cardiac CTA data from 83 patients with manual annotations, and by the assessment of the ability to refine inaccurate centerlines. The method was qualitatively evaluated by a clinical expert who rated the segmentation quality of 250 automatic cross-sectional segmentations and compared the quality of these segmentations with manual annotations.

5.3.1 Data and manual annotation

We collected 83 cardiac CTA datasets acquired in the Erasmus MC, University Medical Center Rotterdam, The Netherlands. These datasets were randomly selected from a series of patients who underwent a cardiac CTA examination between May 2005 and December 2008. Forty-eight datasets were acquired with a 64-slice CT scanner and 35 datasets with a dual-source CT scanner (Sensation 64 and Somatom Definition, Siemens Medical Solutions). The average voxelsize was $0.37 \times 0.38 \times 0.40\text{mm}^3$. In 72 of these datasets in total 195 coronary arteries were annotated by one expert observer. In the remaining 11 datasets 34 coronary arteries were annotated by two expert observers. On average the annotated coronaries were 12.6 cm long.

The annotation was done in three steps; first a centerline was annotated, then in six different views curved planar reformatted images were generated in which the

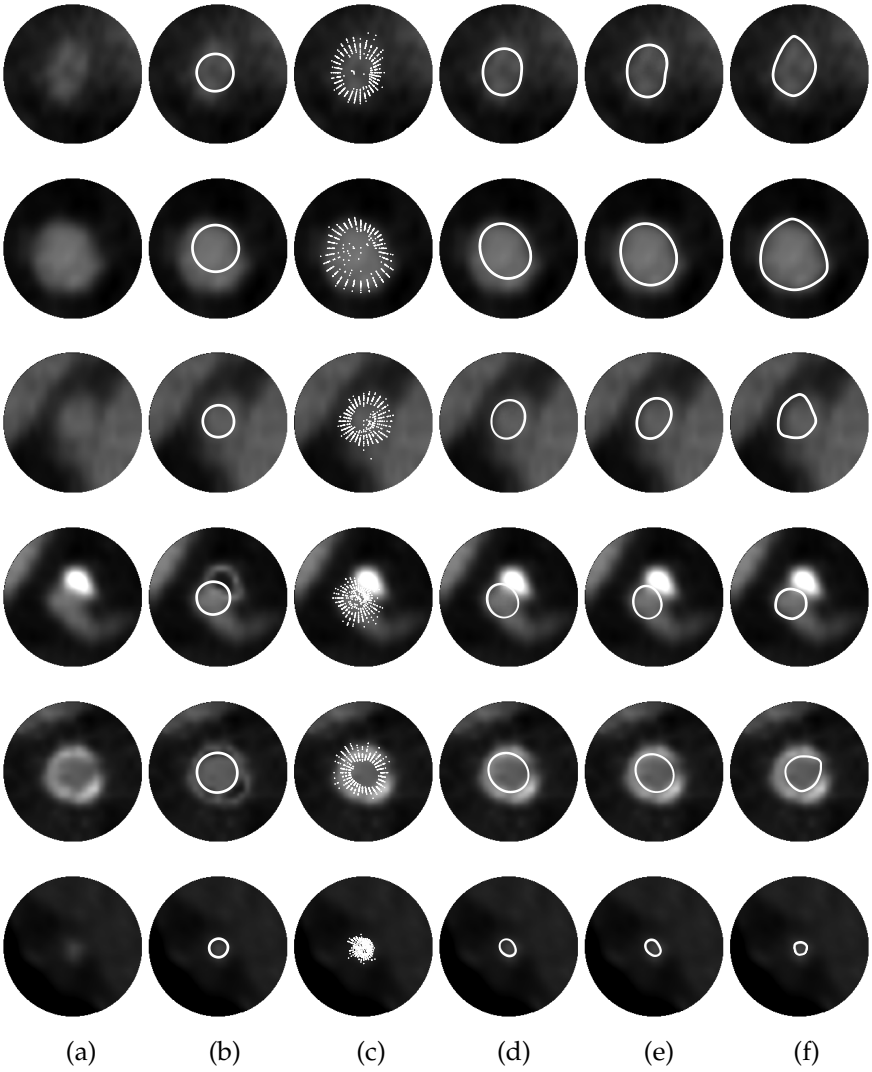

Figure 5.8: *Robust shape fitting results. a) Input image. b) Ridge Regression result. c) KNN result. d) Shape fitting at $t = 1$. e) Segmentation result (approximately $t = 7$). f) Manual annotation.*

vessel boundary was annotated longitudinally. In the last step cross-sectional images were generated every 1.0 mm containing the intersection of the longitudinally annotated mesh. These contours were inspected and modified if necessary.

5.3.2 Parameter settings

The radius of the region of interest was fixed to 7.5 mm, to accommodate for the largest coronary diameters of approximately 5.0 mm and an inaccurate centerline with an error of approximately 1.0 mm. In view of the image resolution (approximately 0.5 mm to 1.0 mm [132]), intensities were sampled every 0.5mm, both longitudinally along the centerline and in-plane.

The remaining parameters were optimized in pilot experiments using CTA datasets of 10 patients that were not included in any of the experiments. The image pre-processing parameter $\sigma_z = 15\,\text{mm}$, $\sigma_i = 250\,\text{HU}$, and $c = 200\,\text{HU}$ were empirically chosen by visually inspecting the calcium removal results. The boundary profile size (7.5 mm $\times$ 3.0 mm), the parameter that describes the expected error of the KNN estimate $\sigma_{\text{knn}} = 0.15\,\text{mm}$, the number of shape parameters $M = 12$, the number of principal components for the KNN search $G = 32$, and the number of neighbors $K = 16$ in the KNN search were optimized with a parameter optimization using a series of line searches.

5.3.3 Training procedure

The regression models and samples were obtained from the available training data in three steps. First four random centerlines were created for each of the manually annotated centerlines by adding uniform spatial noise with a magnitude of 0.5 mm to the annotated centerline. This magnitude of noise is approximately twice the accuracy attainable with existing coronary artery centerline extraction methods [142]. In this way the method is trained on inaccurate centerlines and is expected to be sufficiently robust to segment vessels with existing centerline extraction techniques. Subsequently, local coordinate frames were created with the steps described in Section 5.2.1 with a random rotation of the first local coordinate.

In the second step the image intensities were extracted in the ROI around the centerline, the image pre-processing was applied (see Section 5.2.3.1), and the intersection of the cross-sectional planes with the manual annotations were calculated. The cross-sectional images and corresponding annotations were then assembled in local shapes of 5 cross-sections, and the regression matrix $\hat{H}_{\text{RR}}$ was calculated using this set of training data according to the steps described in Section 5.2.3.2.

After estimating the regression matrix $\hat{H}_{\text{RR}}$ coarse shape estimates of the shape parameters β_x^* were determined. Given these estimates and the reference standard we calculated the boundary profiles and associated shifts $\{(\mathbf{b}_j, \mathbf{d}_j)\}$ for all the landmarks. Because this resulted in too many training samples for efficient computa-

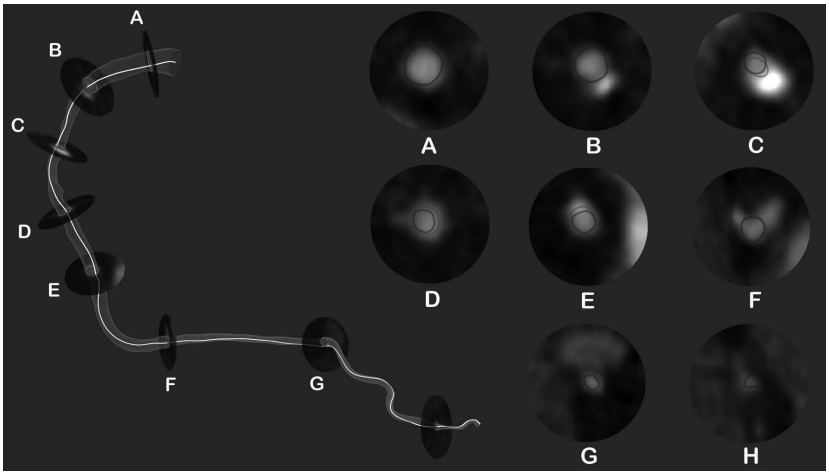

Figure 5.9 (see page 149 for a color-version): *A segmentation of the lumen of a right coronary artery showing a 3D mesh of the automatic segmentation, the initialization curve in white, and 8 selected cross-sections with results of automatic segmentation in red and reference standard in blue.*

tion, 2 million boundary profiles were randomly selected and used as KNN samples in the experiments.

5.3.4 Evaluation of segmentation accuracy and robustness

The segmentation accuracy was evaluated with a leave-one-patient-out study. For each of the 83 patients we built a training database with the 82 other patients using the steps presented in Section 5.3.3. This database was subsequently used to segment all the annotated vessels with the 4 random initialization curves described in Section 5.3.3. This resulted in $263 \times 4 = 1052$ coronary segmentations with each on average 252 cross-sectional planes, sampled every 0.5 mm.

The segmentation accuracy was assessed with the root mean squared landmark-to-surface distance (RMSD) for all the estimated landmark positions. Moreover, the cross-sectional lumen area was estimated using the resulting segmentation and this area was compared with the area manually quantified. For the 34 coronaries with two annotations we also determined the inter-observer variability and we quantified the difference of the automatic segmentations with each of the two manual annotations.

In this quantitative evaluation we also evaluated the robustness of the method for different initialization centerlines. For each coronary we determined the RMSD between the 3D surfaces obtained with the 4 different random initialization center-

lines and the average of these 4 surfaces.

5.3.5 Evaluation of centerline accuracy

A second quantitative evaluation was performed by assessing the resulting centerlines after segmentation and subsequent centerline extraction. For this purpose we used the Rotterdam Coronary Artery Evaluation Framework [142]. This framework allows easy evaluation of centerline extraction methods with standardized evaluation measures. We tested our method on the 24 testing datasets provided by the framework. In each of these datasets four coronary arteries were extracted and evaluated. We used the start- and end-point of these 96 vessels, as provided by the organizers of the framework. These points were connected with the minimal cost path technique and parameters described in [107]. Subsequently the extracted centerlines were visually inspected and in the case of complete extraction failures additional points were annotated. Given these centerlines the coronary lumen was segmented with the proposed technique, and centerlines were extracted from the resulting segmentations by calculating the center of the contours with Equation 5.2. The resulting centerlines were submitted to the online evaluation framework[†]. The method was evaluated and compared to 15 existing methods according to the four evaluation measures provided by the framework, which measure both tracking success and accuracy [142].

5.3.6 Qualitative evaluation

For the qualitative evaluation of the proposed method we randomly selected 250 cross-sections from the segmentations obtained in the first quantitative evaluation (Section 5.3.4). These cross-sections and the manual and automatic contour were shown, together with four neighboring cross-sectional images to provide longitudinal context, to one of the two observers responsible for the manual annotation. The manual contour and automatic contour were randomly displayed above or below the cross-sectional images.

Without showing the bottom contour the top contour was scored on a 5-point scale (very good, good, moderate, poor, very poor). Subsequently the bottom contour was also shown and the top contour was scored relative to the bottom on a 5-point scale (much better, better, equally good, worse, much worse). In this comparative study the expert observer was asked to try and select the 'equally good' rating as little as possible. This procedure was repeated one week later after randomly re-ordering the images. Moreover, the bottom image was now set on top and vice versa. In this way the 250 manual and automatic contours were twice qualitatively compared with each other and each image was scored on the absolute scale once. Two of the images used during this evaluation are shown in Fig. 5.10.

[†]http://coronary.bigr.nl

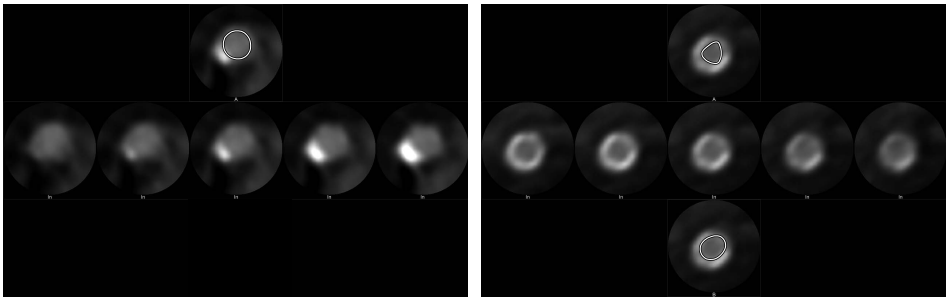

Figure 5.10: *Examples of the images used for the qualitative evaluation. Left: Image used for the absolute rating (showing a manual contour). Right: Image used for the relative rating (automatic contour is shown at the bottom). See Section 5.3.6 for details about these figures.*

5.4 Results

5.4.1 Segmentation accuracy

Figure 5.11 shows the results of the quantitative evaluation of the segmentation accuracy based on 195 annotated coronaries in 72 cardiac CTA datasets (23128 cross-sectional contours). Here, the automatic segmentations were compared to manual annotations made by one observer. The results show that the average root mean squared difference (RMSD) between automatic and manual contours is approximately 0.21 mm, and that 97% of the automatically generated landmark positions have an error smaller than the average voxelsize of the images (0.39 mm). Furthermore, 5.11 shows the results for the cross-sectional area estimation. If the coronary cross-section is circular, the difference in area has a quadratic relation with the radius of the vessel if a similar surface inaccuracy is observed. Therefore before comparison of the manually and automatically estimated area we first approximated the lumen radius from these numbers (radius $= \sqrt{\frac{\text{area}}{\pi}}$). These approximated radii were compared using Bland-Altman plots [19]. It can be seen that in 95% of the cases the manually quantified radius can be approximated with an error of maximally 0.44 mm (approximately the size of a voxel).

Figure 5.12 shows the results for the robustness experiment. For the 195 coronaries the average surface difference between the four segmentations with different initialization and the average segmentation was 0.07 mm. The 95% confidence interval for radius estimation with different initialization centerlines was also 0.07 mm.

Figure 5.13 and Fig. 5.14 show the results of the comparison of the segmentation results with the inter-observer variability. This comparison is performed with manual annotations of the lumen of 34 coronary arteries in 11 datasets drawn by 2 observers. The figures show that the surface difference between manual annotations

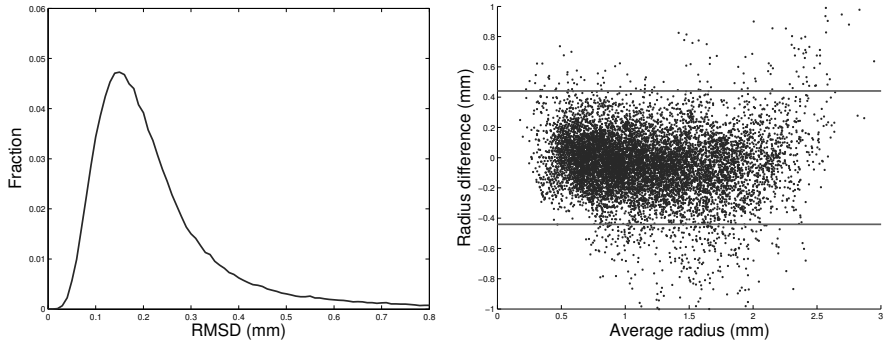

(a) Histogram of cross-sectional surface accuracy (RMSD).

(b) Bland-Altman plot for the cross-sectional area estimation.

Figure 5.11: *The accuracy of the segmentation method, quantified by comparing segmentations of 195 coronaries with manual annotations made by 1 observer.*

and automatic segmentations is on average slightly better than the inter-observer variability (0.22 mm and 0.21 mm vs. 0.24 mm) and a similar effect can be seen for the cross-sectional area estimation.

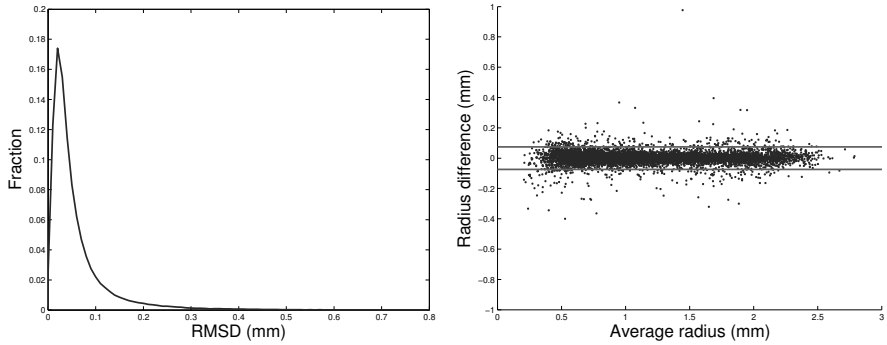

(a) Histogram of cross-sectional surface accuracy (RMSD).

(b) Bland-Altman plot for the cross-sectional area estimation.

Figure 5.12: *The robustness of the segmentation method, quantified by comparing segmentations of 195 coronaries with 4 random initialization centerlines to the average of these segmentations.*

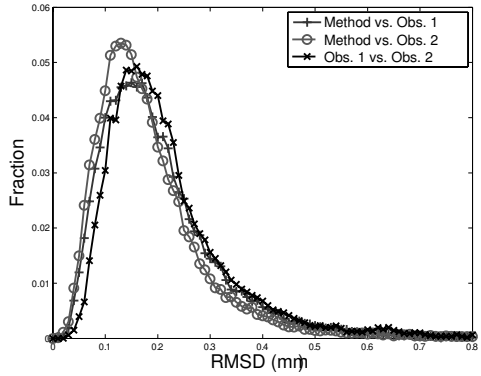

Figure 5.13: *The accuracy of the segmentation method and the inter-observer variability. The dashed green and red line show the accuracy of the method compared to manual annotations and the solid blue line shows the inter-observer variability. Quantified with segmentations of 34 coronaries made by 2 observers.*

5.4.2 Centerline extraction

Table 5.1 shows the results on the Rotterdam Coronary Artery Evaluation Framework [142] when deriving the centerlines from the coronary segmentations with the proposed method. The table shows results for four evaluation measures for the 24 testing datasets of the evaluation framework. Three of these measures describe the

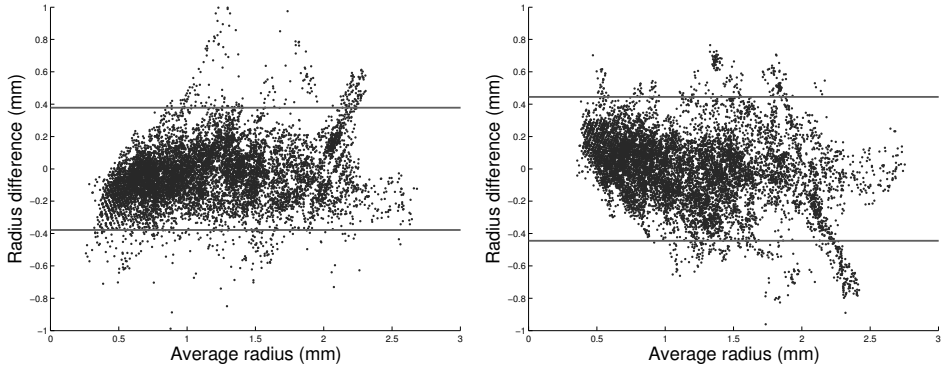

(a) Method vs. Observer 1 (b) Observer 1 vs. Observer 2

Figure 5.14: *Bland-Altman plots for the accuracy of the cross-sectional area estimation and the inter-observer variability. Quantified with segmentations of 34 coronaries made by 2 observers.*

Table 5.1: *Summary of the results obtained on the 24 testing datasets of the Rotterdam Coronary Artery Evaluation Framework. See http://coronary.bigr.nl/preview/346785 for more details.*

Measure	% / mm			score			rank		
	min.	max.	avg.	min.	max.	avg.	min.	max.	avg.
OV	77.7%	100.0%	96.9%	42.3	100.0	79.2	1	12	3.35
OF	4.4%	100.0%	72.5%	2.2	100.0	66.3	1	17	4.40
OT	77.9%	100.0%	97.1%	41.8	100.0	79.2	1	12	3.40
AI	0.09 mm	0.53 mm	0.23 mm	33.8	63.2	49.6	1	7	1.97
Total							**1**	**17**	**2.84**

Table 5.2: *Results of the comparison of the proposed method to 15 existing coronary centerline extraction methods. The best 5 competing methods, the results for the semi-automatic initialization curves, and the results of the proposed method are shown. See http://coronary.bigr.nl/preview/346785 for results on all 15 other methods and detailed results for all vessels.*

Method	Avg overlap rank	Avg accuracy rank	Avg. rank	Computation time	User-interaction
Friman et al.[49]	2.87	2.42	2.65	6 minutes	Inter. - 2.6pts.
Proposed method	**3.70**	**1.97**	**2.84**	**12 + 10 minutes**	**2.2pts**
Wang and Smedby[168]	4.69	4.83	4.76	2 minutes	Inter. - 3pts.
Szymczak[155]	6.01	3.96	4.99	30 minutes	Inter. - 2pts.
Lesage et al.[86]	3.28	7.46	5.37	4 minutes	Inter. - 2pts.
Zambal et al.[188]	8.78	5.01	6.90	4-8 minutes	Auto.
Initialization curves [107]	**6.93**	**8.07**	**7.50**	**12 minutes**	**2.2pts.**
..	..	..	..	..	..

overlap or extraction capability (i.e. how much of the centerline can be found?) in a similar fashion as the well-known Dice similarity metric. They represent the ability to find the complete centerline (OV), the ability to find the centerline without making errors (OF), and the ability to find the clinically relevant part of the vessel (OT). The accuracy is assessed with one measure describing the extraction accuracy if the centerline is located within the vessel (AI). We would like to refer to [142] for further details about these measures and other values in Table 5.1. Table 5.2 compares the overall results to the currently 5 highest ranked methods, the centerlines before refinement (which we used as initialization), and the centerlines after refinement. A comparison to the 10 other methods currently evaluated with the framework and detailed results of all vessels can be found at: http://coronary.bigr.nl/.

5.4.3 Qualitative evaluation

Table 5.3(a) and 5.3(b) show the results for the qualitative evaluation. Table 5.3(a) shows that 151 of the 250 automatically obtained segmentations were twice judged to be better than manual annotations and 42 manual contours were twice scored better than automatic segmentations. The remaining 57 segmentations were inconsistently scored. Figure 5.15 shows several example of the images used for the quali-

Table 5.3: *Results of the qualitative evaluation. (a) Comparison study. Horizontal: First comparison. Vertical: Second comparison. Each cell shows the amount of couples in the comparitive study that were scored in the two runs in a specific combination (+ + = method much better than the manual contours, + = method better, . = equally good, - = method worse, - - = method much worse. (b) The independent qualitative rating showing the percentage per rating type.*

(a) Comparison study

	+ +	+	.	-	- -
+ +	6	16	0	0	0
+	12	117	0	25	1
.	0	1	2	2	1
-	0	23	2	29	4
- -	0	0	0	7	2

(b) Independent rating

	Manual	**Method**
Very good	1.2%	10.8%
Good	31.2%	52.4%
Moderate	34.0%	19.6%
Poor	26.8%	12.4%
Very poor	6.8%	4.8%

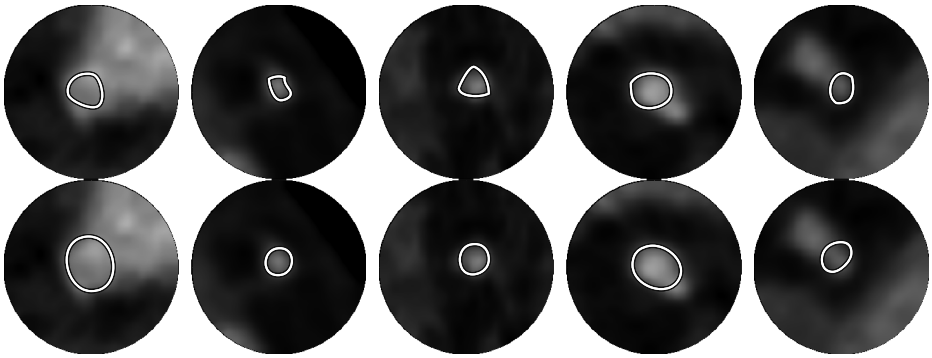

(a) $++\&++$ (b) $++\&+$ (c) $+\&$ equal (d) $--\&-$ (e) $--\&--$

Figure 5.15: *Examples of images used for the qualitative evaluation with corresponding score for the first and second rating (see caption of Table 5.3(a)). Top row: manual annotation. Bottom row: automatic segmentation.*

tative evaluation with the corresponding rating. The absolute scoring in Table 5.3(b) confirms that the automatic segmentations are on average better rated than manual annotations.

5.5 Discussion and Conclusion

We have proposed a supervised vessel segmentation approach which learns the relation between image intensities and vessel shape from an annotated dataset. The method has been successfully applied to the segmentation of coronary arteries in CTA data. Our quantitative experiments showed that the method is highly accu-

rate. On average the difference between the automatically obtained segmentations and manual contours is smaller than the inter-observer variability, which is a strong indicator that the method outperforms manual annotations. Moreover, in a qualitative expert evaluation the automatically obtained contours were rated on average better than manual contours. Very good results were also obtained on the data of the Rotterdam Coronary Artery Evaluation Framework [142] when resulting segmentations were used for centerline refinement. The proposed method is currently ranked 2nd out of 17 evaluated methods. Because the method learns the geometry and appearance of the vessel lumen in 3D images from annotated data we believe that the method, if provided with enough representative training data, can readily be applied to vessels with a wide variety of shapes and appearances. The method is currently evaluated with segmentations of complete vessels. In future work we would like to investigate the relation between pathology and segmentation accuracy.

Although the presented method showed high accuracy and robustness, we believe that several aspects of the method can still be improved. The method can handle centerlines with an inaccuracy up-to approximately 0.5 mm. During the centerline refinement experiment we observed that in very few cases the semi-automatically obtained initialization curves were not satisfactory, resulting in a locally failed segmentation. This problem can be solved by using a better centerline extraction method or by improvements of the coarse segmentation step. This can, for example, be achieved with a similar technique as presented in the shape regression work of Zhou et al. [191]. Before the ridge regression step we would add a position regression step that classifies a series of translated cross-sectional tubular image patches based on its presence of a vessel in the center of the image. The translated image patches with the highest response are then used as input for the ridge regression step and all the resulting shapes can then used for the shape refinement step.

During shape refinement, each KNN estimate is now fixed to a specific landmark. If the difference between the center of the coarse segmentation and the final segmentation is large the KNN estimate can correspond to a significantly different angle and thereby a different landmark. Although we only noticed very few cases where this could have caused problems, we believe that the shape-refinement step can possibly be improved with a correspondence-free approach, such as presented in the work by Saragih et al. [140].

In conclusion, a novel supervised vessel segmentation approach that fits a set of local shape models to image data in a coarse-to-fine approach has been presented. The applicability of the method to a challenging clinical problem, the segmentation of the coronary lumen in CTA data, was demonstrated with excellent results on both quantitative and qualitative experiments.

Chapter 6

Noise Reduction in CTA for Improved 3D Visualization

Based on:

M. Schaap, A.M.R. Schilham, K.J. Zuiderveld, E.P.A. Vonken, M. Prokop and W.J. Niessen

Fast Noise Reduction in Computed Tomography for Improved 3D Visualization, IEEE Transactions on Medical Imaging, 2008

Computed tomography (CT) has a trend towards higher resolution and higher noise. This development has increased the interest in anisotropic smoothing techniques for CT, which aim to reduce noise while preserving structures of interest. However, existing smoothing techniques are computationally demanding, which makes clinical application difficult. Furthermore, the published methods have limitations with respect to preserving small details in CT data. This chapter presents a widely applicable speed optimized framework for anisotropic smoothing techniques. A second contribution of this chapter is an extension to an existing smoothing technique aimed at better preserving small structures of interest in CT data. Based on second order image structure, the method first determines an importance map, which indicates potentially relevant structures that should be preserved. Subsequently an anisotropic diffusion process is started. The diffused data is used in most parts of the images, while structures with significant second order information are preserved. The method is qualitatively evaluated against an anisotropic diffusion method without structure preservation in an observer study to assess the improvement of three-dimensional visualizations of CT series and quantitatively by determining the reduction of the difference between low and high dose CT scans of in vitro carotid plaques.

6.1 Introduction

Computed X-ray tomography (CT) is a widely used imaging modality for, among others, the diagnosis of cancers, infectious diseases, trauma, and cardiovascular diseases. A major disadvantage of CT is the use of ionizing radiation, which may induce cancer in the exposed individual after a latent period of up to a few decades [163].

Cancer induction by ionizing radiation is a probabilistic process. Reduction of radiation dose used in CT will therefore lead to a reduction in the number of induced cancer cases. Dose reduction is therefore an important topic, especially for low-risk patients with a long life expectancy. Reduction of the radiation dose, however, will increase the amount of photon noise in CT images, which will degrade the image quality. While there are several other effects that influence CT image quality, such as background noise and subject movement, photon noise contributes the most to CT image quality degradation [16, 62, 64]. Extensive research has been performed on the relation between image quality and diagnostic value [53, 54, 70, 113]. The results of these studies show that there is a clear relation between an increase in noise and a reduction in diagnostic accuracy and vice versa, i.e. a lower amount of noise will improve the diagnostic quality. Furthermore, a higher intrinsic resolution will lead to improved tissue differentiation and thereby improved diagnostic quality. However, increasing the spatial resolution at a constant radiation dose will also increase the amount of photon noise per voxel.

Owing to the constant demand for lower dose and increasing resolution in CT, several methods have been explored to lower the relative amount of noise per voxel at constant resolution and dose. A major concern when developing an image enhancement method is preservation of structures of interest. This is one of the key aspects of the method that is presented in this chapter. A well-known group of image enhancement methods is based on the principle of anisotropic diffusion [173]. Anisotropic diffusion has the disadvantage of being computationally complex and thereby relatively slow, which hampers its introduction into clinical practice. In order to overcome this problem we present a new speed optimized implementation for anisotropic diffusion. This optimized anisotropic diffusion implementation is used in a novel image enhancement method that improves the diagnostic quality of reconstructed CT image data with constant spatial resolution and radiation dose. The method combines the well-known anisotropic diffusion technique edge enhancing diffusion (EED) [173] with structure preservation. The achieved speed optimization is determined by comparing the computation time of the new optimized method with a standard implementation. The newly developed image enhancement method is qualitatively and quantitatively compared to EED without structure preservation.

6.2 Previous work

Photon noise reduction methods for CT can be divided into two groups. Noise reduction in the measured data (the projection data) and noise reduction in the 3D attenuation map that is determined from these measurements (the reconstructed image).

Noise reduction methods in the projection data can have the advantage that a priori information about the projection data, such as the Poisson distribution of the measurement noise, can easily be incorporated in the noise reduction process. However, filtering in projection data has the disadvantage that it is very difficult to use reconstructed shape information in the noise reduction process. Furthermore, projection data is stored differently for each CT vendor, while the reconstructed image is stored according to strict guidelines. Therefore most methods, including the one presented here, focus on noise reduction in the reconstructed CT image.

The first articles about noise reduction in CT were published shortly after the invention of CT. The methods were based on low-pass filtering [34, 137]. Low-pass filtering removes high frequencies from the image, thereby reducing noise and improving the detectability of large objects in the image. However, there are problems associated with these methods. Low-pass filtering reduces the intrinsic resolution of the image; it smoothes edges and decreases the detectability of small structures.

A method which included the prior knowledge that large differences between neighboring voxels are unlikely to be caused by noise, by weighing neighboring voxels during smoothing, depending on the intensity differences, is presented by Okada [120]. Less important structures are smoothed with this method; the CT images maintain their spatial resolution better.

Lee et al. presented the sigma filter, this filter assumes a Gaussian distribution of CT noise [63, 83]. The method estimates the standard deviation of noise throughout the image and averages voxels in local neighborhoods that are within an intensity range depending on this global intensity variation.

This noise dependent averaging achieves a higher spatial resolution with an equal level of smoothing. A recent extension to this method is the NOVA method presented by Schilham et al. [147]. The NOVA filter estimates iteratively the standard deviation of noise on each position in a CT volume and uses this local measure to perform a weighted average between neighboring voxels.

Noise reduction methods that preserve specific structures of interest give even better results [71, 72, 74]. These methods take the magnitude and the direction of the difference between neighboring voxels into account and outperform undirected methods because of this anisotropic character.

A popular framework that enables inhomogeneous and anisotropic blurring is called anisotropic diffusion. A large body of literature has been devoted to anisotropic diffusion (e.g. [174]). Anisotropic diffusion is based on the diffusion equation in which the diffusion coefficient is made image dependent. Specific image struc-

tures can be directionally smoothed with anisotropic diffusion. An example of a CT noise reduction method that uses anisotropic diffusion is described by Frangakis et al. [46]. The method filters CT data with a specifically tuned algorithm that reduces noise, enhances edges, and improves elongated (vessel like) structures. Recently two 3D anisotropic CT image enhancement methods specifically designed for enhancement of elongated structures have been proposed [98, 117]. However, because these methods only enhance vessel like structures in the image they are not suited for generic image enhancement.

6.3 Method description

The proposed image enhancement method is specifically aimed at improving the diagnostic quality of CT images within the framework of anisotropic diffusion. Numerical implementations of anisotropic diffusion are computationally demanding and hence a speed optimized implementation will be introduced in section 6.4.

In this section we will first outline the requirements that we impose on our image enhancement method. Whereas image quality is a subjective measure, the two most important aspects of diagnostic image quality are considered to be low noise and high resolution [41]. Image quality is expected to improve if noise is reduced on and around important structures and if the resolution of these structures is maintained or improved. We formulate our prior knowledge on the important structures in CT images as follows: they consist both of large structures such as bones, the aorta lumen, and liver and small structures, such as small vessels, calcifications and bronchi. It is our believe that preservation of resolution is especially important for the smallest structures. Intensity variation reduction is important in the large structures, because these structures are disturbed by noise structures that can not be distinguished from the real intensity variations in these structures.

The requirements of the new method are therefore formalized as:

- Reducing the intensity variation within the large scale structures.
- Improvement of detectability of the edges between large scale structures.
- Preservation of the important small scale structures.

6.3.1 Edge enhancing diffusion

The first two requirements can be fulfilled by the anisotropic diffusion technique edge enhancing diffusion (EED) [173]. Anisotropic diffusion can be described as an extension of the linear diffusion equation, by introducing a direction dependent conduction coefficient $\mathbf{D}(\mathbf{u})$:

$$\frac{\partial u}{\partial t} = \nabla \cdot (\mathbf{D}(\mathbf{u}) \nabla u). \tag{6.1}$$

The interpretation of this equation is that the image u is diffused over time t (in voxels^2), where the diffusion is steered by the diffusion tensor $\mathbf{D}(\mathbf{u})$ that is based on local image geometry. Equation (6.1) states that the intensity change per voxel is equal to the divergence of the intensity gradient after a diffusion tensor transformation. EED expresses the diffusion tensor in terms of its eigenvectors and eigenvalues; the first eigenvector of $\mathbf{D}(\mathbf{u})$ V_1 is parallel to the gradient, calculated at scale σ:

$$V_1 \quad \| \nabla u_\sigma. \tag{6.2}$$

The eigenvalues μ_i are defined as follows:

$$\begin{aligned} \mu_1 &= \begin{cases} 1 & (|\nabla u| = 0) \\ 1 - \exp\left(\frac{-C_m}{(|\nabla u|/\lambda_{\text{eed}})^m}\right) & (|\nabla u| > 0) \end{cases} \\ \mu_{[2..n]} &= 1 \end{aligned} \tag{6.3}$$

where n is the dimension of the image. The constants $m = 4$ and $C_4 = 3.31488$ are predefined in [173]. The parameter λ_{eed} is a contrast parameter related to the expected gradient magnitude at diagnostically important edges. When $|\nabla u| >> \lambda_{\text{eed}}$ the diffusion will be anisotropic, while isotropic blurring occurs at $|\nabla u| < \lambda_{\text{eed}}$ [166].

6.3.2 Structure controlled edge enhancing diffusion (SCEED)

Our extension to EED fulfills the third stated requirement, the preservation of small scale structures. This implies that noise on the small structures is not reduced with the proposed method. This approach is chosen for two reasons: an attempt to reduce noise in the small structures is expected to reduce conspicuousness of these details, as it is too difficult to differentiate between noise and the important details on this small scale. Secondly, enhancement of the small structures can only be done if the undisturbed shape is known and this a priori information is not available.

The new method interpolates between the original image and its EED filtered version. It uses the original image data at locations with small important structures, edge enhancing diffusion if there are no small important structures, and a linear transition between these two in the intermediate regions. The weight factors of the interpolation are based on an importance map that is derived using second order information.

Construction of the importance map The measure of total curvature is used to indicate diagnostically important structural information [47]. Because important structures can exist on multiple scales, a multi-scale approach is used to construct this importance map.

The importance S_α of each voxel at scale α is calculated as α^2 times the total curvature at scale alpha [91]:

$$S_\alpha = \alpha^2 \sqrt{\mu_1^2 + \mu_2^2 + \mu_3^2}, \tag{6.4}$$

where μ_i refers to the ith eigenvalue of the Hessian matrix calculated at scale α.

The maximum response of the importance measures S_α is selected by

$$S = \max_{\alpha_{\min} \leq \alpha \leq \alpha_{\max}} S_\alpha, \tag{6.5}$$

where $\alpha_{\min}$ and $\alpha_{\max}$ correspond to the minimum and maximum scale of the Hessian calculation. The range $\alpha_{\min}$ to $\alpha_{\max}$ is based on the expected scale of the small structures in the image and it is divided exponentially in discrete steps.

Based on the calculated value for S, the importance measure T is calculated with the following exponential function

$$T(S) = 1 - \exp\left(-\left(\frac{S}{\gamma}\right)^\beta\right). \tag{6.6}$$

The threshold-like function $T(S)$ has a range from 0 to 1, tending to 1 when S is large. The position and the steepness of the transition between important and unimportant structures is controlled by β and γ.

Blending There are several ways to incorporate an importance map in a traditional anisotropic diffusion scheme. We opt for linearly blending the original and the diffused image. In order to have spatially smooth transitions between the original image and the edge enhanced image, we blend using $T_\rho(S)$, a Gaussian smoothed version of the importance measure $T(S)$. Using this approach, the final voxel intensity is calculated as follows:

$$\begin{aligned} I_{\text{out}} &= T_\rho(S) I_{\text{in}} + (1 - T_\rho(S)) I_{\text{eed}}, \\ T_\rho(S) &= T(S) * K(\rho). \end{aligned} \tag{6.7}$$

I_{in}, I_{out} and I_{eed} are respectively the input image, output image and the EED diffused image. The scale of the Gaussian kernel K is ρ.

6.4 Implementation

The conventional EED implementation is a very time consuming algorithm because of its iterative nature. Furthermore, the additions of SCEED increase the computational complexity of the image enhancement method. We therefore introduce a novel approach which considerably speeds up anisotropic smoothing techniques.

6.4.1 Standard implementation of anisotropic diffusion

A discretization of equation (6.1) is required in order to apply anisotropic diffusion on discrete images. Let U denote the discretized version of the continuous image u in equation (6.1). The scale-space axis is divided in steps of Δt with discrete scale parameter k, where $t = k\Delta t$ and U_0 denotes the original image. Iterative methods can be used to calculate an image at a specific scale t in the scale-space by choosing a specific Δt and by performing $t/\Delta t$ iterations; during each iteration, U_{k+1} is calculated using the previous iteration U_k as input. Generally, the discrete anisotropic diffusion equation is solved in the following explicit way:

$$\frac{U_{k+1} - U_k}{\Delta t} = A(U_k)U_k, \tag{6.8}$$

where $A(U_k)U_k$ is a linear discretization of the operation $\nabla \cdot (D(U_k)\nabla U_k)$.

The discrete equation can also be solved as follows

$$\frac{U_{k+1} - U_k}{\Delta t} = A(U_k)U_{k+1}. \tag{6.9}$$

Since this requires solving a linear system of equations and $A(U_k)$ is a function of U_k and not U_{k+1}, this scheme is called semi-implicit. Explicit and semi-implicit discretization of the anisotropic diffusion equation differ on various points. With explicit schemes the step in scale-space Δt has the restriction of $\Delta t < \frac{1}{2n}$, where n is the image dimension; the solution becomes numerically unstable above this step size [173]. A widely-used step size in the explicit 3D implementation ($n = 3$) is $\Delta t = 0.15$ which provides numerical stability, speed and reasonable accuracy [173].

With semi-implicit schemes it is not necessary to take small steps in scale-space and Δt can be set equal to the desired scale t. This means that the semi-implicit method can be made non-iterative and the diffusion tensors have to be calculated only once [173, 175]. A semi-implicit solution results in a linear solution of the anisotropic diffusion equation, which is often satisfactory in practice. The non-iterative character of the semi-implicit method is especially advantageous for anisotropic diffusion schemes that require complex calculations to obtain the diffusion tensors.

However, semi-implicit schemes are assumed to be too time-consuming, because of the extremely large system of linear equations that needs to be solved. Therefore, most existing anisotropic diffusion methods use the explicit scheme with $\Delta t = 0.15$ to calculate discrete solutions to the anisotropic diffusion equation.

The values for $A(U_k)$ can be calculated, for both approaches, with a finite differences scheme described in [173]. This scheme uses forward-, backward-, and central differences for calculating the different discrete derivatives. The result of this discretization of the 3D diffusion equation is that matrix $A(U_k)$ has 19 coefficients on each row that can be nonzero. Those coefficients are on fixed distances from the diagonal. The required 19 voxels span a $3 \times 3 \times 3$ neighborhood in 3D excluding

the 8 corner voxels of this cube. These coefficients are based on the values of the 7 diffusion tensors at the position of the current voxel and its direct neighbors in the x,y and z directions.

6.4.2 Existing optimization methods

Several authors have described methods for fast nonlinear isotropic diffusion. Weickert's additive operator splitting (AOS) and Acton's multigrid method [4, 175] yield a significant performance increase. However, these methods can not be applied to anisotropic diffusion. Although Acton named the multigrid paper 'multigrid anisotropic diffusion', this paper describes an optimization to nonlinear isotropic diffusion, not anisotropic diffusion.

Mrázek presented a method for anisotropic AOS, but this method has limitations to the degree of anisotropy that can be applied [111]. Standard methods for solving large systems of linear equations like the Gauss-Seidel method have also been used for solving discrete semi-implicit diffusion equations [11, 175], but none of these methods have been tailored to anisotropic diffusion.

Besides these generically applicable speed optimizations, a number of authors have published methods to quickly solve partial differential equations on graphical processing units (GPUs). These methods are used to efficiently perform anisotropic diffusion with modern GPUs. An overview of these methods can be found in [136] and [123]. Because of the limited availability of these modern graphical cards, we focused on generically applicable speed optimizations.

To our knowledge, there are no generically applicable methods that outperform the standard explicit iterative anisotropic diffusion implementation.

6.4.3 New optimization of the anisotropic diffusion implementation

The newly developed optimization of anisotropic diffusion equations contains two parts. The first part is an optimization for the discretization of the diffusion process, which can also be used for an explicit implementation, and the second is an optimized semi-implicit implementation.

6.4.3.1 Optimized anisotropic diffusion discretization

To update voxel intensities during the explicit or semi-implicit diffusion process requires for each voxel, in the standard linear discretization method, the diffusion tensors in its seven connected neighborhood [173]. If these 7 diffusion tensors are all equal to the identity matrix the diffusion changes locally from anisotropic to isotropic. In this case it is not necessary to access the values of the diffusion tensors; the discretized anisotropic diffusion calculation simplifies to a simple weighted average of the input voxel value and the values of its 6 neighbors. Owing to this potential speed up, a mask is created during the calculation of the diffusion tensors

indicating all 'identity' diffusion tensors. A diffusion tensor is marked as being identity if the difference of each element with an identity matrix is smaller then $\epsilon_{\text{identity}}$.

This mask is eroded with a 7 points kernel, equal to the described neighborhood. The eroded mask is used during the diffusion process to switch between the straightforward isotropic diffusion equation and the standard discretization of the diffusion equation.

6.4.3.2 Optimized semi-implicit diffusion

The second part of the optimization is a method which is specifically developed to quickly solve the linear system of equations for the semi-implicit anisotropic diffusion process. The method is inspired by the successive overrelaxation (SOR) method [11].

The discretized semi-implicit diffusion equation from equation (6.9) can be written as

$$U_{k+1} = (I - \Delta t A(U_k))^{-1} U_k. \tag{6.10}$$

A solution to this equation can be obtained with several iterations of the SOR method. For that purpose equation (6.10) is rewritten as

$$A'X = B \tag{6.11}$$

where A' is equal to $(I - \Delta t A(U_k))$, B is equal to U_k and X is the solution U_{k+1}. The SOR method iteratively calculates the solution X in the following way

$$x_i^k = \omega \frac{b_i - \sum_{j<i} a'_{ij} x_j^k - \sum_{j>i} a'_{ij} x_j^{k-1}}{a'_{ii}} + (1-\omega) x_i^{k-1} \tag{6.12}$$

where $\omega < 2$ describes the amount of extrapolation. The coefficients a'_{ij} are calculated as follows

$$a'_{ii} = 1 - \Delta t A_{ii} \tag{6.13}$$

$$a'_{ij} = -\Delta t A_{ij} \tag{6.14}$$

where A_{ij} is computed using isotropic diffusion (averaging) if it is part of the eroded identity mask, and computed using the discretized diffusion tensors otherwise. Because the matrix A' is diagonal dominant [173], the SOR method will converge to a solution [10].

For the 3D anisotropic diffusion equation, there are 19 nonzero values of a'_{ij}. From equation (6.12) it follows that it is not necessary to evaluate the coefficients that are zero. Also, if all 19 x_j^k terms associated with the nonzero elements of a'_{ij} are equal to x_j^{k-1}, the value of x_i will not change between iterations.

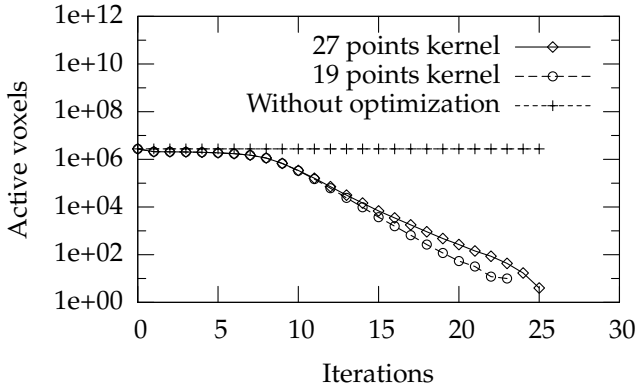

Figure 6.1: *The number of active voxels per iteration while solving the semi-implicit EED diffusion equation of a characteristic* 140^3 *voxels CT image. The graph shows the number of active voxels for the 19-points dilation method and for the method with the 27-points dilation. The parameters were as described in section 6.5.1.4 with* $t = 2$

.

If a voxel changes during an iteration, the voxel is marked in a so-called active voxel mask. A voxel is marked as changed if the absolute difference with the previous value is above a certain threshold ϵ_{sor}. This threshold is set to a very small value, in order to ensure that the speed optimization does not result in a less accurate solution. Convergence is achieved if all voxels change less than ϵ_{sor}. The active voxel map is dilated with a kernel of the described 19 points neighborhood, resulting in a map that indicates the voxels that can change in the next SOR iteration. Because dilation with a $3 \times 3 \times 3$ kernel can be computed fast with a separable filter, the dilation is done with a 27 points kernel, instead of the described 19 points kernel. Fig. 6.1 shows how the number of active voxels for a $140 \times 140 \times 140$ subvolume of a typical CT dataset decreases with a 19 and 27 points kernel. After five iterations, the number of active voxels decreases approximately exponentially. A 19 point dilation is comparable to a 27 points dilation on a densely populated mask, therefore the results of the first few iterations are comparable. The 27 points dilation results in a slight increase in the number of iterations, but because of the increased efficiency the overall processing time is lower.

It should be noted that the resulting optimization method with active voxel masking and dilation does not have exactly the same convergence as an unmodified successive overrelaxation method because not all of the elements of x^{k-1} are evaluated every iteration. With the successive overrelaxation method, the new value x_i^k of a voxel x_i^{k-1} influences the calculation of the elements x_j^k where $j > i$ and $A_{ij} \neq 0$.

The SOR method can be seen as an extended version of the Jacobi method [11]. With the Jacobi method, the new value of x_i^k has only influence on the calculation of the values of x^{k+1}. It is possible with the presented method that x_i is active and that one or more elements x_j are inactive. If all the elements x_j are inactive the newly developed method resembles the Jacobi method locally. If all the x_j elements are active, the new method resembles the successive overrelaxation method locally. Therefore, the presented method can be interpreted as a hybrid of the Jacobi method and the successive overrelaxation.

6.4.4 Implementation details

The Hessian was calculated using Gaussian derivative operators [47]. Five discrete steps were used to divide the range between α_{min} and α_{max} exponentially. An experiment with 32 different CT images showed that a value for the SOR parameter ω of 1.33 gave the highest convergence speed. This parameter setting was used in all experiments. The threshold ϵ_{sor} in the optimized SOR method was set to 0.001. The threshold $\epsilon_{\text{identity}}$ in the creation of the identity mask was set to 0.0001.

With these values the resulting images have a negligible small difference with the result of an un-optimized implementation; when rounded to the nearest integer, 0.04% of the voxels have a difference of 1HU and no voxels have a difference of more than 1HU, for an intensity range of -1000HU to 4000HU.

6.5 Evaluation

SCEED was evaluated qualitatively with an observer study and quantitatively by assessment of the reduction of the root mean squared difference (RMSD) between the intensity values of the low and high dose data after low dose data processing with SCEED.

The performance of the optimized semi-implicit method was evaluated by determining the reduced computation time of the new optimized method in comparison with a standard explicit EED implementation.

6.5.1 Qualitative evaluation

In thoracic and abdominal CTA, electrocardiogram (ECG) gating can be used to reduce motion artifacts caused by heart motion or to detect and quantify motion due to heart or vessel pulsation [56, 112, 119, 146]. With this technique, CT projection data and ECG-recordings are combined to reconstruct images at different phases of the heart cycle. ECG-gated reconstruction is always necessary in anatomical regions with severe movement, like the area of the coronary arteries. The term ‘gated scan’ will be used to refer to one reconstructed phase of such a retrospectively gated CT scan.

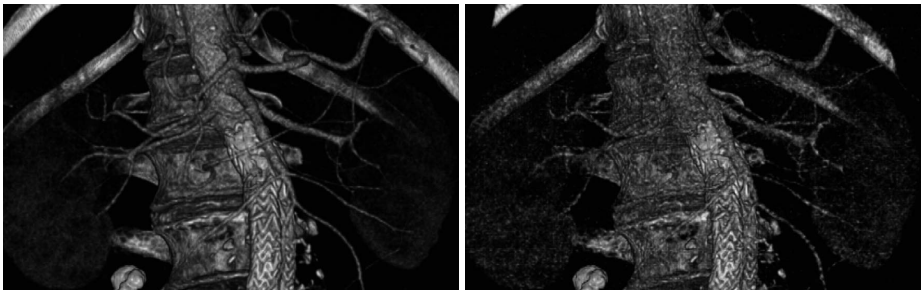

(a) (b)

Figure 6.2 (see page 150 for a color-version): *A section of a volume rendering of (a) gated and (b) ungated AAA CTA images. Note the difference in quality between, for example, the visualization of small vessels.*

The separate phases of an ECG gated scan contain only approximately 10% to 20% of the projection data, therefore they contain more noise than an ungated CT scan. As a result, 3D visualizations of gated CT studies are of a lesser quality than 3D visualizations of ungated CT studies in anatomical areas with limited movement (Fig. 6.2).

One of the application of CTA is the diagnosis of an abdominal aortic aneurysm (AAA). An AAA is a focal dilation of the abdominal aorta usually below the level of the renal arteries. The diagnosis of AAAs requires a clear insight in the size and orientation of the aneurysm and the aorta. Volume rendering has become widely used for this purpose [20]. Because the area of the abdominal aorta undergoes limited motion, ungated images are most often used for this purpose.

In this study we evaluated whether SCEED noise reduction improves 3D volume renderings of gated CTA AAA images. As 3D volume rendering of ungated reconstructions is currently the prevalent method, we use these ungated scans as a reference standard. 3D visualization of noise reduced gated images were compared with those of unfiltered gated images and ungated images. Noise reduction was performed with SCEED and EED.

6.5.1.1 Scan protocol

Ten CTA data sets were randomly selected from patients who were referred for pre- or post-operative CTA AAA examinations. Retrospectively ECG gated CTA scans were acquired in the University Medical Center Utrecht with a 64 slice CT scanner (Brilliance CT, Philips Medical Systems, Best, The Netherlands) using a dose of 300 mAs and tube voltage of 120kV. Images were reconstructed with a slice thickness of 1.4 mm, slice spacing of 0.7 mm, a field of view of 300 mm and a resolution of 512×512; resulting in an in-plane resolution of approximately 0.6 mm. Both stan-

dard ungated and eight-phase gated reconstructions were made. The gated reconstructions were reconstructed at 70% between two consecutive R-peaks of the ECG [108].

6.5.1.2 Volume renderings

Three-dimensional visualizations of the CTA AAA datasets were created with Vital Images' Vitrea 2 workstation using the default 3D vessel preset; except for the window width, all settings were kept at their default values.

Volume renderings from two different viewing directions were made for each of the four different images (gated, ungated, EED, and SCEED); because of the excellent reproducibility, a standard anteroposterior and an orthogonal lateral view were used. To increase the visibility of the aorta, a part of the hip bone was trimmed from each lateral image. This resulted in a total of $10 \times 2 \times 4 = 80$ different 3D renderings.

The optimal window level for 3D visualization depends on the goal of the visualization and the noise characteristics of the image. Appropriate window width values were calculated using grey level intensities from areas in the background tissue and their standard deviation. Visual inspection by a team of experts suggested that a lower bound of the intensity window of two times the standard deviation of the liver above the average liver intensity yields a good compromise between the visibility of details of interest and low disturbance by unconnected noise voxels, for 3D visualizations of AAAs. The window level was set to its default value of 270HU. Therefore the window width value (in Hounsfield units) was selected with:

$$W = 2(270 - (HU_{\text{liver}} + 2SD_{\text{liver}})), \tag{6.15}$$

where HU_{liver} is the average intensity of the liver and SD_{liver} the standard deviation of the liver, both measured in Hounsfield units.

6.5.1.3 Observation study

For both visualization directions, the four different pictures of each patient were compared side-by-side. All the six combinations of the four pictures were scored relatively on a 3 points scale (left better, equal, or right better).

Each of the pairs of visualizations were scored using four aspects which experienced radiologists judged to be important for diagnostic value; the four aspects were:

- Detail visibility: The expected amount of correctly visualized small details and the visibility of these small details; the preservation of structures known to be present by anatomical knowledge.

- Structural sharpness: The sharpness, low blurriness, and expected accuracy of the edges between the important structures, like calcifications and stents, and their surroundings.

- Smoothness of surfaces: The low graininess on the hip bone and aorta.

- Overall impression: The preferability for a normal clinical work flow.

Three observers (two radiologists and a fourth year radiology resident) graded the 3D visualizations. Each observer scored a total of 10 (patients) $\times 2$ (directions) $\times 4$ (aspects) $\times 6$ (combinations) $= 480$ pairs of pictures. All of the 480 pairs were ordered using patient, direction, and aspect. The six combinations and the position of the images within each pair were randomly ordered and the resulting combinations were shown sequentially. The observers were blinded for the method that was used to generate the volume renderings.

6.5.1.4 Parameters

The noise filtering parameters for the qualitative evaluation were found empirically; an expert panel determined the parameter settings for the two algorithms by tuning them for the specific task of three-dimensional visualization of gated AAA CT images. Three CT AAA datasets were used for this purpose. These datasets were not used in the qualitative evaluation. The parameters were determined as follows.

EED Scale parameter $t = 0.8$, gradient scale $\sigma = 4.0$ mm, preservation function parameter $\lambda_{\text{eed}} = 6.0$.

SCEED Scale parameter $t = 3.0$, Hessian calculation from $\alpha_{\text{min}} = 0.9$ mm to $\alpha_{\text{max}} = 3.0$ mm, importance function parameters $\beta = 5$ and $\gamma = 26.5$, scale of importance map blurring: $\rho = 0.5$ mm. $\sigma = 4.0$ mm, and $\lambda_{\text{eed}} = 6.0$.

Explicit and optimized semi-implicit EED diffusion solutions were visually compared by the expert panel on the reduction of intensity variations in the large scale structures and the improvement of the detectability of edges between the large structures. These aspects were judged to be similar. Therefore, only the semi-implicit solutions were used in the rest of the study in order to reduce the necessary computation time and workload of the observers.

6.5.2 Quantitative evaluation

To quantitatively compare SCEED and EED, their ability to reduce the root mean square difference (RMSD) of low dose and high dose CT imaging data was assessed.

6.5.2.1 Data

Five ex-vivo carotid plaque samples were placed in a 16 cm water-filled neck phantom to achieve a realistic noise level. Ten scans of the phantom were acquired on a Siemens Sensation 64 with 120 kVp, slice thickness and spacing of 0.6 mm, and the following exposures: 25, 60, 90, 120, 150, 180, 210, 240, 270, and 427mAs. Exposures of 25 and 427mAs were respectively the minimum and maximum possible mAs settings with the specified slice thickness and spacing. The acquired data was reconstructed to datasets with cubic voxels of $(0.6\text{mm})^3$. The 'golden' standard for the quantitative evaluation was created by averaging the five highest dose scans (Fig. 6.6(b)). Only the five highest dose scans were used because the lower dose series contained noticeable streaking artifacts.

6.5.2.2 Implementation and parameters

To achieve optimal parameter settings for minimization of the RMSD between the low dose scans and the golden standard, filter parameter optimization was carried out with the Powell optimization technique [126]. This was done for both SCEED and EED in a region of interest of about 1200 mm^3 around two carotid plaques. The following parameters were optimized for EED: t, σ, and λ. SCEED was optimized by varying the following parameters: t, σ, λ, α_{min}, α_{max}, β, γ, and ρ. The optimal parameter settings were subsequently used to measure the RMSD around the three remaining carotid plaques in the image.

6.5.3 Speed-optimization evaluation

To evaluate the time required for EED filtering using both a normal explicit and the optimized semi-implicit method, both methods were implemented in C++ and the code was compiled with the Intel® compiler 8.1 on the highest optimization level. The resulting programs were executed on a workstation with a dual Xeon 2.8 gigahertz and 4 gigabytes of internal memory using 512×512×256 CTA images as input data.

Two different scale-space steps were used for the explicit method. $\Delta t = 0.01$ provided a solution with a high accuracy while $\Delta t = 0.15$ provided a more practical balance of accuracy, stability, and speed. To demonstrate the relation between scale and computation time, the comparison was done for multiple evolution times t, varying from 0.15 to 4.0. Ten different 512×512×256 CTA images were diffused with all these evolution times.

The overall processing time of the SCEED method was also determined; the method contains several additional steps compared to the EED method. It calculates the Hessian matrix for every voxel, constructs an interpolation factor from the calculated matrix, and blends the EED image with the original image. These extra steps do not depend on the scale parameter, therefore only one setting of the scale

Table 6.1: *Rankings for the four qualitative aspects for the lateral and anteroposterior direction. A lower ranking number means a better ranking. The rankings were identical for the two volume rendering directions and the analysis also produced the same results if the observations of the two viewing directions were combined. TODO: lijn*

	Rank			
Aspect	**Ungated**	**Gated**	**EED**	**SCEED**
Detail visibility	1	4	3	2
Smoothness of surfaces	1	4	3	2
Structural sharpness	1	2	3	3
Overall impression	1	4	3	2

parameter was necessary for this measurement. An evolution time of $t = 3$ was chosen, i.e. the same evolution time as used for the qualitative evaluation.

6.6 Results

6.6.1 Qualitative results

In total 1440 observations were made by the three observers. These observations were used to create rankings per aspect and viewing direction.

The scores for each filter combination and aspect were combined into 48 lists for each of the six combination, four aspects and two directions. Each list contained three (observers) $\times$ ten (patients) = 30 measurements. A sign-test, with a significance threshold of 0.01, determined whether one of the two combinations scored significantly better [1].

These statistics were used to create rankings for each aspect and direction. A filter is ranked higher than another filter if it is significantly better according to the sign-test. The resulting rankings are listed in table 6.1.

The rankings show that both evaluated noise reduction methods were considered beneficial for improved detail visibility, surface smoothness, and overall impression of gated scans. Furthermore, it shows that the observers ranked the visualizations of the ungated scans the highest.

Fig. 6.3 contains examples of volume renderings used for the qualitative evaluation. All three observers ranked these images for the aspect overall impression in the increasing order of gated (worst), EED, and SCEED (best). In order to demonstrate the preservation of real structures and variation reduction in axial slices, figure 6.4 shows the effect of SCEED on a low dose ungated scan.

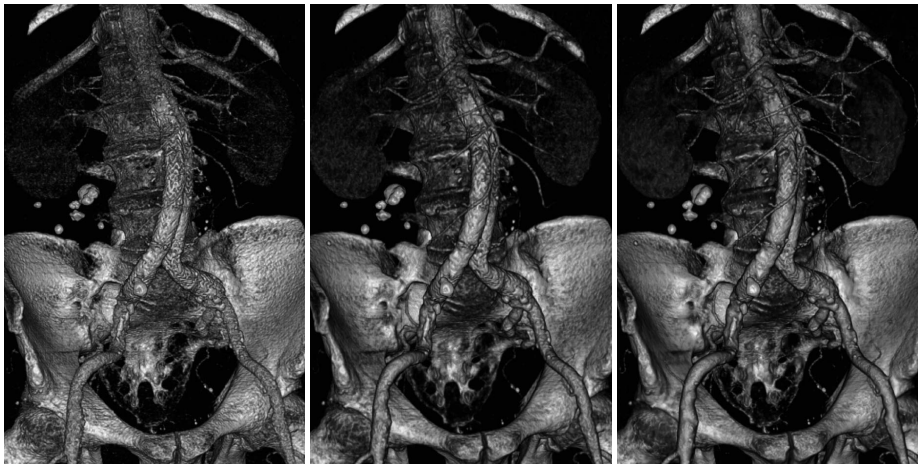

(a) (b) (c)

Figure 6.3 (see page 150 for a color-version): *(a) 3D visualization of a gated dataset. (b) Visualization after EED filtering. (c) Visualization after SCEED filtering. Note the improved visibility of the small vessels in the SCEED visualization in comparison with the EED visualization. The unimproved structural sharpness in the EED and SCEED scan can be seen in the region of the stent, located in the bifurcation of the abdominal aortic artery in the center of the image.*

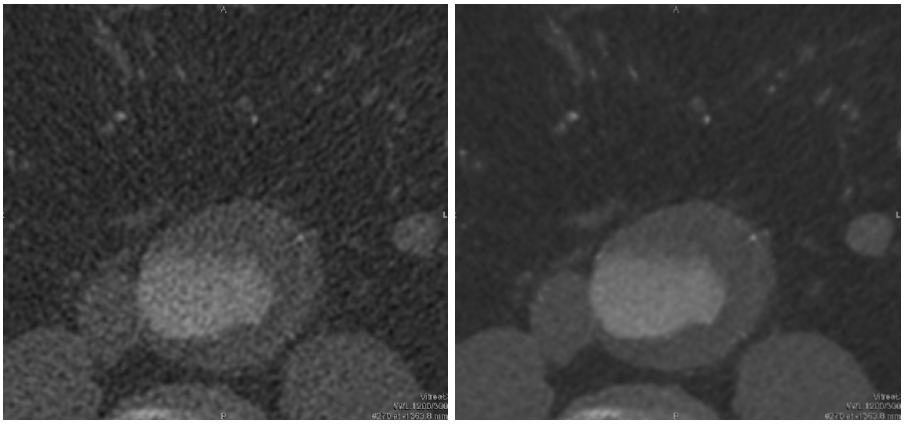

Figure 6.4: *Left: A section of an unfiltered axial CTA slice. Right: The same section filtered with SCEED. Note the improved detectability of the small vessels anterior to the aorta and calcifications in the aorta.*

Table 6.2: *Inter-observer agreement. (a) The three observers scored a pair the same. (b) Two observers rated a combination the same and the third was one category off. (c) Two observers scored a pair the same and the third observer was two categories off. (d) The observers had all different scores.*

Aspect	(a)	(b)	(c)	(d)
Detail visibility	85%	10%	3%	3%
Surface smoothness	89%	5%	5%	1%
Structural sharpness	68%	14%	13%	5%
Overall impression	82%	9%	8%	2%
Total	81%	10%	7%	3%

6.6.1.1 Inter-observer variability

Table 6.2 shows the relation between the scores from the different observers; the table shows that image pairs are scored identically in more than 80% of the cases. The table also shows that observers agreed most on the surface smoothness aspect and least on the structural sharpness aspect.

6.6.1.2 Intra-observer statistics

Table 6.3 shows statistics on the ability of observers to make consistent rankings and the differences between their lateral and anteroposterior observations of the same pair. These numbers give additional information about the reliability of the observer study. Ranking consistency determines the percentage of the observations where an observer made a consistent scoring. An example of an inconsistent observation is that an observer scored for a specific patient, direction and aspect EED better than ungated, ungated better than SCEED, and SCEED better than EED. The directional consistency indicates the percentage of the observations in which the observer scored a pair identical in the lateral and anteroposterior direction. The table shows that observers made a consistent ranking 90% of the cases; it also indicates that the structure sharpness aspect has the highest variability. The same can be concluded for the directional consistency aspect, with even higher percentages.

6.6.2 Quantitative results

The RMSD around the three carotid plaques decreased considerably with both EED and SCEED in all the low dose scans. Fig. 6.5 shows the unprocessed RMSD and the RMSD after EED and SCEED noise reduction. In all scans the lowest RMSD was achieved with SCEED, which was statistically significant ($p < 0.0001$), assessed with a one-tailed paired Student's t-test. Fig. 6.6 shows an example of the data that was used for quantitative evaluation.

Table 6.3: *Average consistency of the three observers. The percentage of the observations when the observer made a consistent ranking and the percentage of the cases when an observer scored a pair the same for the lateral and anteroposterior direction.*

Aspect	Ranking consistency	Directional consistency
Detail visibility	91%	98%
Surface smoothness	94%	100%
Structural sharpness	86%	93%
Overall impression	88%	98%
Total	90%	98%

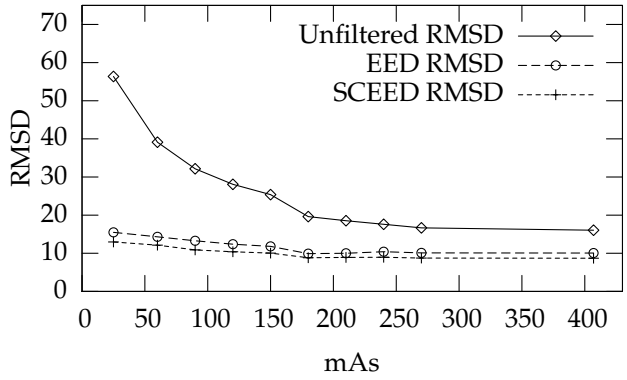

Figure 6.5: *Quantitative results. The RMSD in areas around the carotid plaques before and after noise reduction with EED and SCEED.*

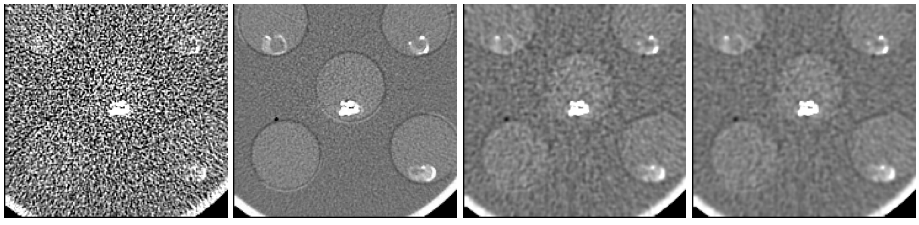

Figure 6.6: *Examples of the quantitative evaluation data. (a) A section of the 25mAs scan. (b) A section of the 'golden' standard. (c) A section of the low dose image after SCEED filtering. (d) A section of the low dose image after EED filtering.*

6.6.3 Speed-optimization results

6.6.3.1 EED computation time

Linear regression was used to model the relation between the scale t and the computation time in seconds s for the explicit, semi-implicit and optimized semi-implicit method. The computation time of the explicit method with $\Delta t = 0.15$ can be approximated with $s = 2.2 \cdot 10^2 t - 2.0$. The explicit method with $\Delta t = 0.01$ can be approximated with $s = 3.9 \cdot 10^3 t + 52$ and the fully optimized semi-implicit method with $s = 31t + 27$. The Pearson's coefficient of regression R^2 was larger than 0.99 for all three linear approximations.

For scales around $t = 1$, the optimized semi-implicit method is approximately 52 times faster than the explicit method with $\Delta t = 0.01$ and 3.2 times as fast as the explicit method with $\Delta t = 0.15$. The optimized method is approximately 5.6 times faster than the explicit method with $\Delta t = 0.15$ for the scale ($t = 3$) which was used in the qualitative evaluation of the SCEED method.

6.6.3.2 SCEED computation time

The SCEED filter operation takes on average 257.6 seconds, this means that, for a $512 \times 512 \times 256$ image, the total computation time of SCEED is still more than a factor of two less than EED with the standard explicit implementation and without structure preservation.

6.7 Discussion and conclusion

An image enhancement method has been developed, with the aim to improve CT images in clinically more acceptable computation times. The two main novelties of the approach are a new implementation for anisotropic diffusion which is optimized for speed and an extension to an existing anisotropic image diffusion approach, namely edge enhancing diffusion (EED), to better preserve structures of interest. The new implementation achieves a considerable speed up and the new image enhancement method SCEED outperforms EED, both qualitatively and quantitatively.

Both SCEED and EED improve the quality of 3D visualizations of gated CTA images for three of the four evaluated aspects, namely detail visibility, surface smoothness and overall impression. SCEED outperforms EED on all these aspects. The visualizations of ungated images were qualitatively judged better compared to visualizations of the noise reduced gated images. However, gated images are preferred when anatomical regions with severe movement, like the area of the coronary arteries, are imaged. SCEED is expected to improve diagnosis if these images are filtered prior to evaluation.

There was a high inter-observer agreement and intra-observer consistency for the three improved qualitative aspects. This, together with the low value for the significance threshold of the sign-test, suggests that the results are very reproducable.

In contrast, the structural sharpness aspect had a lower inter- and intra-observer agreement. This may be explained by the fact that images with a high amount of noise can have good structural sharpness and yield a bad overall impression. It is conceivable that observers were sometimes biased towards the image with highest overall impression, which would explain the lower inter- and intra-observer variability of the structural sharpness aspect.

SCEED and EED both reduce the difference between a high dose scan and low dose scans of ex vivo carotid plaque samples considerably and SCEED outperforms EED significantly on this quantitative aspect, which further demonstrates the superiority of SCEED over EED.

Chapter 7

Summary and Discussion

In this thesis we proposed and evaluated algorithms towards quantitative image analysis of CTA data, namely automated techniques for centerline extraction and coronary lumen segmentation. In addition, a noise reduction method for improving the diagnostic quality of 3D visualization of CTA images has been presented and evaluated. In this chapter we summarize the main contributions of the work presented and discuss future research directions.

7.1 Summary

7.1.1 Evaluation of coronary centerline extraction techniques

Efficiently obtaining a reliable coronary artery centerline from computed tomography angiography data is relevant in clinical practice. Although a large number of authors have proposed methods for automated coronary centerline extraction a method for the objective evaluation of different algorithms has been lacking.

In Chapter 2 we have described the development and employment of a standardized evaluation framework for the quantitative evaluation of coronary centerline extraction techniques. We presented a method to create a reliable consensus coronary centerline with multiple observers, introduced well-defined measures for the evaluation of tracking accuracy and success, and we publically released and disseminated an annotated image database. This validation database comprises thirty-two cardiac annotated CTA datasets with corresponding reference standard of which 8 are made available for algorithm training purposes, and 24 reference standards are kept blind for objective quantitative evaluation. Thirteen coronary artery centerline extraction algorithms, implemented by different research groups, were quantitatively evaluated and compared. The presented evaluation framework is made available to the medical imaging community for benchmarking existing or newly developed coronary centerline extraction algorithms.

7.1.2 Probabilistic centerline extraction

In Chapter 3 we proposed a probabilistic centerline extraction technique. Compared to other state-of-the art centerline extraction methods, often based on dynamic programming, probabilistic methods have the advantage that more sophisticated priors on the appearance and geometry of the vessel centerlines can be incorporated. Moreover, because of the inherent Monte Carlo based approach, these methods are capable of finding global optima of objective functions defined by complex appearance and geometric priors. For the method described in Chapter 3 we have incorporated prior information on the path curvature in the geometric model and priors on the intensity homogeneity along the centerline in the appearance model, respectively. Both of these priors are very hard to incorporate in dynamic programming approaches that underly minimum cost path techniques,

because of the very large increase in computational resources associated with the integration of these priors.

Probabilistic methods can also be computationally very demanding, because often a large number of hypotheses needs to be taken into account for correct path extraction. Moreover, if too few hypotheses are kept in the optimization the high-dimensional space representing all the possible paths can be sampled too sparsely, resulting in incorrect path extractions. We therefore also implemented optimizations that made these approaches computationally less intensive, while maintaining the robustness of probabilistic tracking.

We demonstrated that the presented approach can extract elongated structures in synthetic images with very high noise levels and that it can find central lumen lines of the coronary arteries and the internal carotid artery passing through the skull base in CTA data. This latter clinical problem was chosen to demonstrate the potential of the novel observation model. In this application the lumen is sometimes brighter than its surrounding tissue and sometimes it is in close proximity to bone, which has higher intensity values or overlaps in intensity values with the lumen voxels. A quantitative clinical evaluation, which was performed by comparing extracted centerline of 28 carotid arteries to manual tracings of two observers, demonstrated good results.

7.1.3 Unsupervised and supervised coronary lumen segmentation

Coronary lumen assessment is important for quantifying coronary artery disease (e.g. stenosis quantification) and as a pre-processing step for coronary plaque analysis. Two coronary lumen segmentation techniques have been presented in Chapter 4 and 5 of this thesis. The first method, proposed in Chapter 4, is based on graph-cuts and robust kernel regression. It is computationally efficient and an evaluation with 28 manually annotated coronaries from 12 patients showed that the method can be applied with good results in situations where accurate centerlines are available.

Chapter 5 presents the second coronary lumen segmentation algorithm. Whereas the method in Chapter 4 is hand-tuned and specifically designed by us, the method in Chapter 5 is developed using a generic machine learning method that learns the geometry and appearance of plausible vessels from annotated data and subsequently applies this knowledge to unseen data. Therefore this method, if provided with enough representative training data, can readily be applied to vessels with a wide variety of shapes and appearances. The latter method achieved better results than the unsupervised method and was less sensitive to inaccurate centerlines, which makes the method also applicable to situations where less accurate, for example automatically extracted, centerlines are available. The method was evaluated by quantitatively comparing segmentation results to manual annotations of 229 coronary arteries. On average the difference between the automatically obtained

segmentations and manual contours was smaller than the inter-observer variability, which is a strong indicator that the method outperforms manual annotation. The method was also evaluated with the evaluation framework presented in Chapter 2 by using it for centerline refinement, resulting in a 2nd place ranking out of all evaluated methods. An additional qualitative expert evaluation in which 250 automatic segmentations were compared to maual segmentations showed that the automatically obtained contours were rated better than manual contours. A drawback of this method is that it is computationally more demanding than the first method, and that it requires a substantial amount of representative training data.

7.1.4 Noise reduction for improved diagnostic quality

A major disadvantage of CT is the use of ionizing radiation, which may induce cancer in an exposed individual after a latent period of up to a few decades [163]. Therefore, there is a constant demand for lowering the radiation dose while preserving the diagnostic quality of CT. This has resulted in the development of CT noise reduction methods that allow to lower the radiation dose by reducing the image noise while maintaining the resolution and image quality. A major concern when developing a noise reduction method is the preservation of small structures in the image. In Chapter 6 we presented a method for noise reduction in CTA images that preserves small structures of interest. The developed method can improve the diagnostic quality of 3D CTA visualizations while preserving the required amount of radiation dose. The method has been evaluated qualitatively with 10 in-vivo CTA datasets of abdominal aortic aneurysms and quantitatively with CTA data of five excised carotid artery plaques. In the qualitative evaluation a panel of radiologists judged the 3D visualizations of CTA abdominal aortic aneurysms before and after noise reduction with our approach and a standard edge enhancing diffusion (EED) technique; the results demonstrated that the noise-reduced images were diagnostically of statistically significant better quality than 3D visualizations of original images and EED-filtered images. Moreover, the quantitative evaluation showed that the proposed method reduced the difference between a high dose scan and low dose scans of ex-vivo carotid plaque samples considerably better than EED.

7.2 Discussion

In this thesis, we have not only developed image processing techniques for cardiovascular CTA, but we have also focused on extensive validation of such techniques. Thorough evaluation of medical image analysis techniques was not addressed systematically in the field of medical imaging until the first MICCAI 'Grand Challenge' Workshop in 2007, where a ground truth database and evaluation metrics

were presented for three specific medical imaging segmentation problems. We contributed to this important trend with our coronary centerline extraction framework (see Chapter 2), which was presented at the 2nd Grand Challenge workshop at MICCAI 2008. The fully automated web-based evaluation framework allowes an on-going comparison of novel techniques, fulfilling an apparent need, which is witnessed by the numerous times it has been used by the medical imaging community. At the time of writing this thesis 88 research groups have registered for the framework and downloaded the evaluation data, and 21 groups uploaded results to the online framework in order to validate their method. Several publications have appeared with results from the framework (see e.g. [169] , [86]) and we believe that this evaluation framework has set a new standard for coronary centerline extraction methods. At this point the community has accepted the framework as the standard for the evaluation of coronary centerline extraction techniques and it is almost not possible anymore (it should not be!) to publish a new method without evaluating the method with our evaluation framework.

A disadvantage of the current evaluation framework is that not all false-positive extractions are taken into account during the evaluation of fully automatic methods. Therefore, it does not accurately measure to what extent complete vascular trees can be extracted automatically. In future work we potentially would like to improve this aspect of the evaluation framework. Moreover, we aim to extend the framework to evaluation of coronary lumen segmentation methods using the large body of manual lumen segmentation acquired during the development of the method in Chapter 5, we aim to include multi-vendor data, and we would like to create evaluation possibilities for the extraction of quantitative imaging biomarkers, such as the degree of stenosis.

The algorithms proposed in this thesis are a substantial step towards improved diagnosis and automated analysis of coronary arteries in CTA images. The proposed noise reduction method has shown to be capable of improving the diagnostic quality of 3D CTA visualizations and we have demonstrated how probabilistic and machine-learning techniques can be used to develop robust algorithms for respectively tracking and segmentation of vascular structures, such as the coronary arteries.

The next step towards automation would be a fully automated centerline tracking algorithm, which is able to find the centerlines of all major coronary arteries from a standard clinical CTA dataset. Good results for fully automatic centerline extraction have been demonstrated with the evaluation framework presented in Chapter 2. However, at the moment, several semi-automatic methods achieve significantly better results and it is not completely clear to what extent complete vascular trees can be extracted automatically. Based on our experiences with machine learning methods in Chapter 5, we are confident that machine learning approaches can also be applied to improve fully automated centerline extraction. Whereas state-of-the-art centerline approaches are almost all unsupervised and uti-

lize explicit appearance and geometric models, we believe that spatial priors and geometric and appearance models learned from annotated datasets are important components of robust algorithms for centerline extraction in clinically obtained cardiac CTA datasets with a variety of pathologies.

The combination of a robust centerline extraction technique and the coronary lumen segmentation technique presented in Chapter 5 would yield a fully automatic supervised coronary lumen segmentation technique. Because accurate training data is very important for this approach, an active learning approach [149] may be required to boost the available training data and thereby further increase the performance of the method.

Fully automatic coronary artery segmentation of data acquired in clinical routine, which we think can be realized in the near future, will provide an ideal environment for addressing many clinically relevant questions. First of all, automated extraction of coronaries allows quantification of lumen morphology, such as stenoses, vessels lengths and lumen diameters, which are relevant for diagnosis and therapy planning. Furthermore, techniques to detect and analyze calcified and non-calcified plaques can be built on the results of the coronary extraction.

By applying coronary lumen segmentation, quantification and plaque detection methods on large volumes of routinely acquired clinical data, or in the context of large scale studies, insights can be gained in the appearance and variation of coronary artery disease on CTA. Combined with medical patient records, these techniques can be used to relate image-derived parameters such as coronary topology, morphology and plaque parameters to gender, age, body-mass-index, and other risk factors. Moreover, these parameters can be related to clinical outcomes (e.g. mortality or left ventricle ejection fraction) in order to gain insights into the value of these CTA-based quantifications for the diagnosis and risk stratification of CAD, thereby facilitating imaging based personalized prognosis and therapy.

Chapter 8

Additional work

Other contributions in CTA image processing

This chapter presents summaries of four additional 2nd-author Computed Tomography Angiography (CTA) image processing contributions. **Section 8.1** *presents the development and evaluation of a minimal user-interaction coronary centerline extraction technique,* **Section 8.2** *the evaluation of atlas-based segmentation of the heart chambers,* **Section 8.3** *the development and evaluation of a method for the segmentation of the carotid arteries, and* **Section 8.4** *describes research into improved separation between iodine and calcium with dual-energy CTA.*

8.1 Coronary centerline extraction with minimal user-interaction

C.T. Metz, M. Schaap, A.C. Weustink, N.R.A. Mollet, T. van Walsum, W.J. Niessen.
Coronary centerline extraction from CT coronary angiography images using a minimum cost path approach, Medical Physics, 2009.

Coronary centerlines are important for the analysis of cardiac computed tomography angiography CTA datasets as they can facilitate visual inspection of coronary pathology and they can be used as initialization for automated coronary lumen segmentation and plaque quantification (see Chapter 1).

In this work we developed and evaluated a semi-automatic minimal cost path method for coronary centerline extraction from CTA data. The method is evaluated using two different cost functions. The first cost function is based on a frequently used vesselness measure and intensity information, and the second is a recently proposed cost function based on region statistics. User interaction is minimized to one or two mouse clicks distally in the coronary artery. The starting point for the minimum cost path search is automatically determined using a newly developed method that finds a point in the center of the aorta in one of the axial slices. This step ensures that all computationally expensive parts of the algorithm can be precomputed.

The performance of the aorta localization procedure was demonstrated by a success rate of 100% in 75 images. Both the extraction capability (overlap) and accuracy of centerline extraction was quantitatively evaluated on 48 coronary arteries in 12 images by comparing extracted centerlines with a manually annotated reference standard. The method was able to extract 88% and 47% of the vessel centerlines correctly using the vesselness/intensity cost function and region statistics cost function, respectively. For only the proximal part of the vessels these values were 97% and 86%, respectively (see Figures 8.1 and 8.2). Accuracy of centerline extraction, defined as the average distance from correctly automatically extracted parts of the centerline to the reference standard, was 0.64 mm for the vesselness/intensity and 0.51 mm for the region statistics cost function. The interobserver variability was 99% for the overlap measure and 0.42 mm for the accuracy measure. Qualitative evaluation using the best performing cost function resulted in successful centerline extraction for 233 out of the 252 coronaries (92%) in 63 additional CTA images.

In conclusion, the presented results, in combination with minimal user interaction and low computation time, show that minimum cost path approaches can effectively be applied as a preprocessing step for subsequent analysis of coronary artery disease.

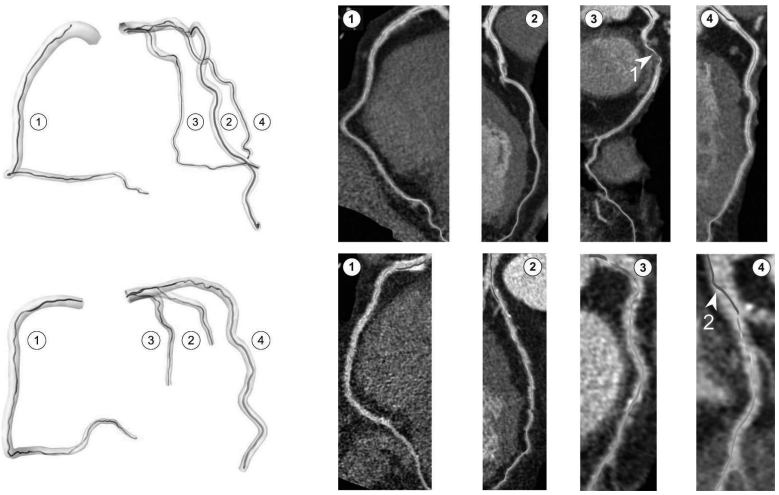

Figure 8.1 (see page 151 for a color-version): *Examples of successfully extracted coronary artery centerlines using the vesselness/intensity cost function. The first column shows the reference standard light gray tubes and the automatically extracted centerlines. The other columns show the automatically extracted centerlines projected onto CPR images based on the reference standard centerline. The greyscale coding of the extracted centerline indicates the distance from the automatically extracted centerline to the reference standard. Arrows in the CPR images indicate: 1) Successfully crossed low contrast region caused by extreme pathology, 2) decreased extraction accuracy at a location where the radius of the vessel is relatively large.*

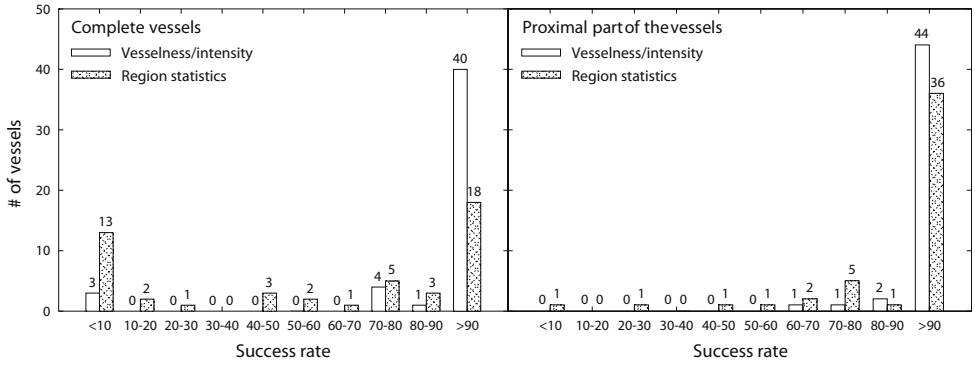

Figure 8.2: *Histogram showing the success rate distribution for centerline extraction using cost functions based on the vesselness measure presented by Frangi et al. [47] and region statistics, presented by Li and Yezzi [88], for the complete coronary arteries and the proximal part of the vessels.*

8.2 Fully automatic cardiac chamber segmentation

H.A. Kirisli, M. Schaap, S. Klein, S.L. Papadopoulou, M. Bonardi, C.H. Chen, A.C. Weustink, N.R. Mollet, E.J. Vonken, R.J. van der Geest, T. van Walsum, W.J. Niessen
Evaluation of a multi-atlas based method 1 for segmentation of cardiac CTA data: a large-scale, multicenter and multivendor study, Medical Physics, 2010.

Computed tomography angiography (CTA) is increasingly being used for the diagnosis of coronary artery disease (CAD). Whereas CTA is not commonly used for the assessment of ventricular and atrial function, this information could potentially improve its diagnostic value. Because the extraction of ventricular and atrial functional information, such as stroke volume and ejection fraction, requires accurate delineation of cardiac chambers, objective and accurate techniques are necessary for cardiac chamber segmentation.

In this work, we presented and evaluated a fully automatic method for segmenting the whole heart (i.e. the outer surface of the pericardium) and cardiac chambers from CTA datasets. The method uses a multi-atlas based approach, in which eight manually labeled atlas images are registered to a cardiac CTA scan, followed by a per voxel majority voting procedure, to obtain a cardiac segmentation. See Figure 8.3 for examples of the resulting segmentations.

We evaluated our method on a multicenter/multivendor database, consisting of: 1) a set of 1380 Siemens scans, from 795 patients, and 2) a set of 60 multivendor scans (Siemens, Philips and GE) from different patients, acquired in 6 different institutions worldwide. A leave-one-out 3D quantitative validation was carried out on the eight atlas images, resulting in a mean segmentation error of 0.94 $\pm$ 1.12 mm and an average Dice coefficient of 0.93. A 2D quantitative evaluation was performed on the 60 multivendor datasets. Here, a mean segmentation error of 1.26 $\pm$ 1.25 mm and an average Dice coefficient of 0.91 were measured. In addition to this quantitative evaluation, a large scale 2-D and 3-D qualitative evaluation was performed on 1380 images. Experts evaluated that 49% of the 1380 image were very accurately segmented (< 1 mm error) and that 29% were accurate (error between 1 and 3 mm), which confirms the robustness of the method (see Figure 8.4).

In conclusion, a fully automatic method for whole heart and cardiac chamber segmentation was evaluated, using a multicenter/multivendor CTA database. By successfully applying the method to 1420 multicenter/multivendor datasets we demonstrated that automatic segmentation method could be used to replace manual outlining by experts.

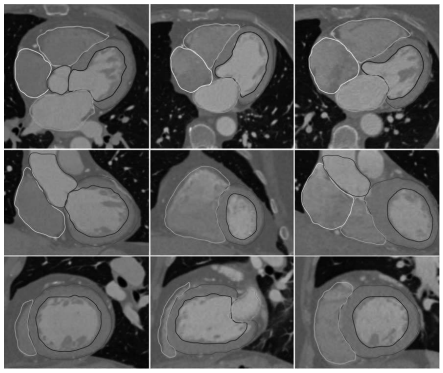

Figure 8.3 (see page 152 for a color-version): *Fully automatic segmentation results from three patients out of the 1380 CTA datasets. Axial view (top), coronal view (middle) and sagittal view (bottom). Segmentations: epicardium left ventricle (pink), endocaridum left ventricle (purple), right ventricle (blue), left atrium (green), right atrium (yellow) and aorta (red).*

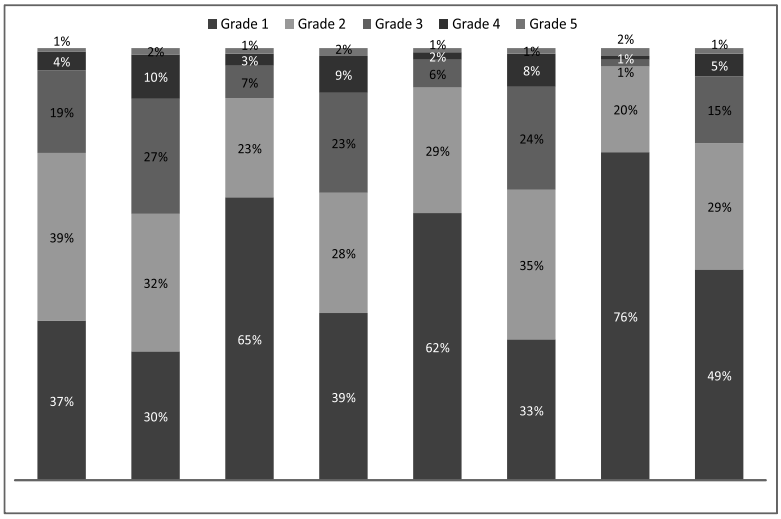

Figure 8.4: *Results of the large-scale qualitative evaluation study using 1380 CTA images. Showing side-by-side the results for a 5-point expert qualitative grade classification for 7 cardiac structures. Grade 1: Very accurate: Deviation up to 1 mm, Grade 2: Most regions accurate: 1-2 regions may deviate up to 3mm, Grade 3: Most regions accurate: 1 region may deviate up to 1cm or >2 regions may deviate up to 3mm, Grade 4: A significant region (up to 50%) has not been segmented or has been incorrectly segmented, Grade 5: Segmentation failed.*

8.3 Semi automatic carotid artery segmentation

R. Manniesing, M. Schaap, S. Rozie, K. Hameeteman, D. Vukadinovic, A. van der Lugt and W.J. Niessen

Robust CTA Lumen Segmentation of the Atherosclerotic Carotid Artery Bifurcation in a Large Patient Population, Medical Image Analysis, 2010.

The carotid arteries are responsible for blood supply to the head and neck. They bifurcate in the neck into the internal and external, which respectively supply blood to the brain and other parts of the head and neck. A widely-used modality for diagnosis of atherosclerosis in the carotid artery is computed tomography angiography (CTA). Clinical, longitudinal and epidemiological studies on the atherosclerotic bifurcation, imaged using CTA, often require a robust lumen segmentation method.

In this work we proposed and evaluated a semi-automatic method for lumen segmentation of the carotid bifurcation in CTA. First, the central vessel axis is obtained using path tracking between three user-defined points. Second, starting from this path, a segmentation is automatically obtained using a level set. The cost and speed functions for path tracking and segmentation make use of intensity and homogeneity slice-based image features. The method is validated on a large data set of 234 carotid bifurcations of 129 ischemic stroke patients with atherosclerotic disease. The results are compared to manually obtained lumen segmentations. Parameter optimization is carried out on a subset of 24 representative carotid bifurcations. With the optimized parameter settings the method successfully tracked the central vessel paths in 201 of the remaining 204 bifurcations (99%) which were not part of the training set.

Comparison with manually drawn segmentations shows that the average overlap between the semi-automatic segmentations and manual segmentation is similar to the inter-observer variability (92% vs. 87%) and the intra-observer variability (94% vs. 94%). Therefore the method has potential to replace the manual procedure of lumen segmentation of the atherosclerotic bifurcation in CTA. See Figures 8.5 and 8.6 for examples of segmentation obtained with the proposed method.

In conclusion, we have presented a semi-automatic lumen segmentation method that has potential to replace manual segmentation of the lumen of atherosclerotic carotid bifurcations in CTA.

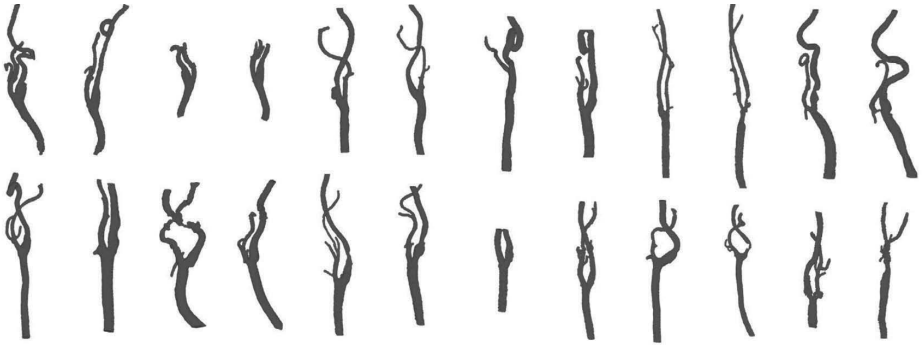

Figure 8.5: *Three-dimensional visualization of carotid lumen segmentations of the carotid bifurcation from 24 CTA images.*

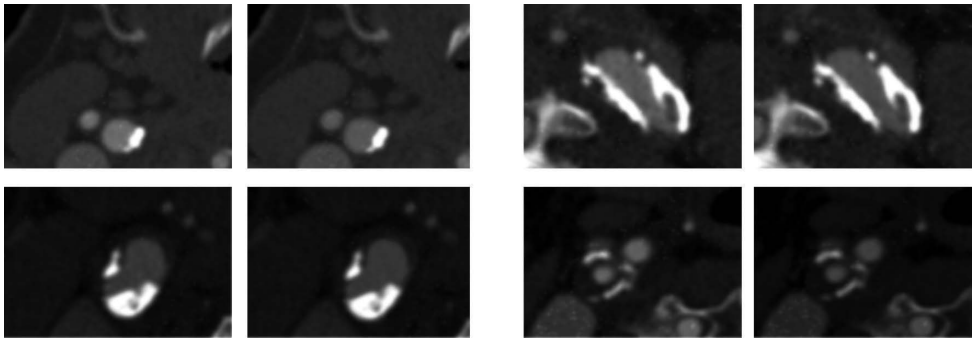

Figure 8.6 (see page 152 for a color-version): *Right: Representative examples of the segmentation results obtained by the proposed method in case of the presence of severe calcified and non-calcified (soft) plaque (rows 1 to 3) and stenosis (bottom row).*

8.4 Dual-energy bone removal

M. van Straten, M. Schaap, M.L. Dijkshoorn, M.J. Greuter, A. van der Lugt, G.P. Krestin, W.J. Niessen

Automated Bone Removal in CT Angiography: Additional Value of Tin Filtration in Dual-Energy Based Methods and Comparison to Single Energy Methods, Submitted, 2010.

Recently, dual-energy CT has been presented in clinical practice using a dual-source CT scanner. Dual-source dual-energy CT uses simultaneously acquires data with two X-ray sources operating at different x-ray tube voltages, i.e. energies. Because the amount of X-ray attenuation (and thereby the CT intensity) is a function of the density and the effective atomic number of a specific tissue, but also of the effective energy of the X-ray signal, this technique is expected to allow better differentiation between iodine and bone, or calcified plaque. Iodine is commonly used to highlight the vascular lumen. Lack of differentiation between iodinated blood and calcified plaques can cause problems in the assessment of the lumen morphology.

In this work we focused on the performance of different methods for bone removal in CTA images and we investigated the added value of tin-filtration, a recent development in dual-energy CT. In a phantom study, we investigated dual-energy based masking and subtraction methods for bone removal and compared these methods with single-energy masking and subtraction methods that use an additional, non-enhanced, CT scan (dual-scan based methods). The bone removal methods were applied to CT scans of an anthropomorphic thorax phantom containing parts which mimic vessels of various diameters in direct contact with bone (see Figure 8.7). Image quality was quantified by the contrast-to-noise ratio (CNR) normalized to the dose (CNRD). At locations where vessels touch bone, the quality of the bone removal and the vessel preservation were visually assessed as well.

All methods removed the bone successfully. Single-energy based methods had a higher CNRD value than the corresponding dual-energy based methods and CNRD values were higher for the masking based methods than for the subtraction based methods. For the subtraction based dual-energy method (see Figure 8.8), tin-filtration improved the CNRD value with approximately 50%. Compared to the dual-scan based methods, the dual-energy based methods have the advantage that all data is acquired within a single scan. Therefore there is no need for image registration, and the method might be implemented more easily in clinical practice, especially if applied to the imaging of coronary arteries.

We concluded that, in general, dual-scan based methods that use a single, optimized, tube voltage for CTA, have a higher CNR than the dual-energy based approaches at the same dose level. Tin-filtration improves the ability to differentiate between iodine and bone for the dual-energy based masking method. In clinical practice, the advantages of the dual-energy masking method might outweigh its disadvantage of a slightly higher dose penalty compared to the conventional dual-scan masking method.

Figure 8.7: *Anthropomorphic thorax phantom with homemade insert for the evaluation of bone removal techniques. The insert contains six cylinders which can be filled with iodinated contrast agents and bony structures.*

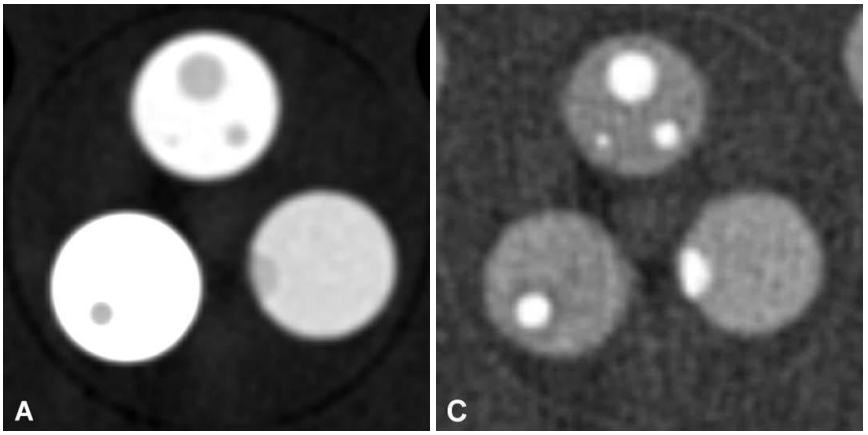

(a) (b)

Figure 8.8: *An example of enhanced contrast between iodine and calcium with dual-energy CT. (a) Image before bone removal (image intensity remapping window width/center [W/C] 510/330 HU). Image after bone removal by dual-energy subtraction (W/C 510/-930 HU).*

References

[1] H. Abdi. Binomial Distribution: Binomial and Sign Tests. In *Encyclopedia of Measurement and Statistics*, volume 1. 2007.

[2] S. Achenbach and P. Feyter. Cardiac CT and Detection of Coronary Artery Disease. *The ESC Textbook of Cardiovascular Imaging*, pages 267–286, 2010.

[3] S. Achenbach, M. Marwan, D. Ropers, T. Schepis, T. Pflederer, K. Anders, A. Kuettner, W. G. Daniel, M. Uder, and M. M. Lell. Coronary computed tomography angiography with a consistent dose below 1 mSv using prospectively electrocardiogram-triggered high-pitch spiral acquisition. *European Heart Journal*, 31(3):340–346, 2010.

[4] S. T. Acton. Multigrid anisotropic diffusion. *IEEE TIP*, 7(3):280–291, 1998.

[5] S. G. Armato, G. McLennan, M. F. McNitt-Gray, C. R. Meyer, D. Yankelevitz, D. R. Aberle, C. I. Henschke, E. A. Hoffman, E. A. Kazerooni, H. MacMahon, A. P. Reeves, B. Y. Croft, L. P. Clarke, and L. I. D. C. R. Group. Lung image database consortium: developing a resource for the medical imaging research community. *Radiology*, 232(3):739–748, 2004.

[6] E. Arnoldi, M. Gebregziabher, U. J. Schoepf, R. Goldenberg, L. Ramos-Duran, P. L. Zwerner, K. Nikolaou, M. F. Reiser, P. Costello, and C. Thilo. Automated computer-aided stenosis detection at coronary CT angiography: initial experience. *European Radiology*, 20(5):1160–1167, 2010.

[7] S. Arulampalam, S. Maskell, N. Gordon, and T. Clapp. A Tutorial on Particle Filters for On-line Non-linear/Non-Gaussian Bayesian Tracking. *IEEE Transactions on Signal Processing*, 50(2):174–188, 2002.

[8] S. Aylward and E. Bullit. Initialization, Noise, Singularities, and Scale in Height Ridge Traversal for Tubular Object Centerline Extraction. *IEEE Transactions on Medical Imaging*, 21(2):61–75, 2002.

[9] S. Aylward, S. Pizer, D. Eberly, and E. Bullitt. Intensity ridge and widths for tubular object segmentation and description. In *Proceedings of MMBIA*, page 0131, 1996.

[10] R. Bagnara. A Unified Proof for the Convergence of Jacobi and Gauss–Seidel Methods. *SIAM Review*, 37(1):93–97, 1995.

[11] R. Barrett, M. Berry, T. F. Chan, J. Demmel, J. Donato, J. Dongarra, V. Eijkhout, R. Pozo, C. Romine, and H. V. der Vorst. *Templates for the Solution of Linear Systems: Building Blocks for Iterative Methods, 2nd Edition*. 1994.

[12] D. Bartz and S. Lakare. Scaffolding-based Segmentation of Coronary Vascular Structures. In *Proceedings of International Workshop on Volume Graphics*, pages 47–54, 2005.

[13] C. Bauer and H. Bischof. A Novel Approach for Detection of Tubular Objects and its Application to Medical Image Analysis. In *Proceedings of DAGM*, pages 163–172, 2008.

[14] C. Bauer and H. Bischof. Edge Based Tube Detection for Coronary Artery Centerline Extraction. *The Midas Journal*, 2008 MICCAI Workshop - Grand Challenge Coronary Artery Tracking:<http://hdl.handle.net/10380/1403>, 2008.

[15] C. Bauer and H. Bischof. Extracting Curve Skeletons from Gray Value Images for Virtual Endoscopy. In *Proceedings of MIAR*, pages 393–402, 2008.

[16] U. Baum, K. Anders, G. Steinbichler, M. Lell, H. Greess, T. Riedel, M. Kachelriess, W. A. Kalender, and W. A. Bautz. Improvement of image quality of multislice spiral CT scans of the head and neck region using a raw data-based Multidimensional Adaptive Filtering (MAF) technique. *European Journal of Radiology*, 14(10):1873–1881, 2004.

[17] R. Bellman and S. Dreyfus. Functional approximations and dynamic programming. *Mathematical Tables and Other Aids to Computation*, 13(68):247–251, 1959.

[18] F. Benmansour and L. Cohen. Tubular Structure Segmentation Based on Minimal Path Method and Anisotropic Enhancement. *International Journal of Computer Vision*, pages 1–19, 2010.

[19] M. Bland and D. Altman. Statistical methods for assessing agreement between two methods of clinical measurement. *The Lancet*, 327(8476):307–310, 1986.

[20] D. T. Boll, J. S. Lewin, J. L. Duerk, D. Smith, K. Subramanyan, and E. M. Merkle. Assessment of automatic vessel tracking techniques in preoperative planning of transluminal aortic stent graft implantation. *Journal of Computer Assisted Tomography*, 28(2):278–285, 2004.

[21] T. Boskamp, D. Rinck, F. Link, B. Kümmerlen, G. Stamm, and P. Mildenberger. New vessel analysis tool for morphometric quantification and visualization of vessels in CT and MR imaging data sets. *Radiographics*, 24(1):287–297, 2004.

[22] B. Bouraoui, C. Ronse, J. Baruthio, N. Passat, and P. Germain. Fully automatic 3D segmentation of coronary arteries based on mathematical morphology. In *Proceedings of ISBI*, pages 1059–1062, 2008.

[23] Y. Boykov and G. Funka-Lea. Graph Cuts and Efficient N-D Image Segmentation. *International Journal of Computer Vision*, 70:109–131, 2006.

[24] Y. Boykov and V. Kolmogorov. Computing Geodesics and Minimal Surfaces via Graph Cuts. In *Proceedings of ICCV*, 2003.

[25] Y. Boykov, O. Veksler, and R. Zabih. Markov random fields with efficient approximations. In *Proceedings of CVPR*, pages 648–655, 1998.

[26] M. Bruijne, B. Ginneken, M. Viergever, and W. Niessen. Adapting active shape models for 3D segmentation of tubular structures in medical images. In *Proceedings of IPMI*, pages 136–147, 2002.

[27] T. Bülow, C. Lorenz, and S. Renisch. A General Framework for Tree Segmentation and Reconstruction from Medical Volume Data. In *Proceedings of MICCAI*, volume 3216 of *Lecture Notes in Computer Science*, pages 533–540, 2004.

[28] S. Busch, T. R. C. Johnson, K. Nikolaou, F. von Ziegler, A. Knez, M. F. Reiser, and C. R. Becker. Visual and automatic grading of coronary artery stenoses with 64-slice CT angiography in reference to invasive angiography. *European Radiology*, 17(6):1445–1451, 2007.

[29] F. Cademartiri, L. L. Grutta, A. Palumbo, P. Malagutti, F. Pugliese, W. B. Meijboom, T. Baks, N. R. Mollet, N. Bruining, R. Hamers, and P. J. de Feyter. Non-invasive visualization of coronary atherosclerosis: state-of-art. *Journal of Cardiovascular Medicine*, 8(3):129–137, 2007.

[30] M. F. D. Carli and R. Hachamovitch. New technology for noninvasive evaluation of coronary artery disease. *Circulation*, 115(11):1464–1480, 2007.

[31] J. Carrillo, M. Hoyos, E. Davila, and M. Orkisz. Recursive tracking of vascular tree axes in 3D medical images. *International Journal of Computer Assisted Radiology and Surgery*, 1(6):331–339, 2007.

[32] C. Castro, M. A. Luengo-Oroz, A. Santos, and M. J. Ledesma-Carbayo. Coronary Artery Tracking in 3D Cardiac CT Images Using Local Morphological Reconstruction Operators. *The Midas Journal*, 2008.

[33] C. Caussin, C. Larchez, S. Ghostine, D. Pesenti-Rossi, B. Daoud, M. Habis, A. Sigal-Cinqualbre, E. Perrier, C.-Y. Angel, B. Lancelin, and J.-F. Paul. Comparison of coronary minimal lumen area quantification by sixty-four-slice computed tomography versus intravascular ultrasound for intermediate stenosis. *American Journal of Cardiology*, 98(7):871–876, 2006.

[34] E. Chew, G. Weiss, R. Brooks, and G. D. Chiro. Effect of CT noise on detectability of test objects. *American Journal of Roentgenology*, 131(4):681–685, 1978.

[35] M. De Bruijne, B. Van Ginneken, M. Viergever, and W. Niessen. Interactive segmentation of abdominal aortic aneurysms in CTA images. *Medical Image Analysis*, 8(2):127–138, 2004.

[36] M. Debruyne, M. Hubert, and J. Suykens. Model selection in kernel based regression using the influence function. *Journal of Machine Learning Research*, 9:2377–2400., 2008.

[37] M. Dewey, D. Schnapauff, M. Laule, A. Lembcke, A. C. Borges, W. Rutsch, B. Hamm, and P. Rogalla. Multislice CT coronary angiography: evaluation of an automatic vessel detection tool. *Rofo*, 176(4):478–483, 2004.

[38] E. Dikici, T. O'Donnell, L. Grady, R. Setser, and R. White. Coronary Artery Centerline Tracking Using Axial Symmetries. *The Midas Journal*, 2008.

[39] A. Doucet, S. Godsill, and C. Andrieu. On Sequential Monte Carlo Sampling Methods for Bayesian Filtering. *Statistics and Computing*, page 197Ű208, 2000.

[40] A. J. Einstein, K. W. Moser, R. C. Thompson, M. D. Cerqueira, and M. J. Henzlova. Radiation dose to patients from cardiac diagnostic imaging. *Circulation*, 116(11):1290–1305, 2007.

[41] B. Escalante-Ramirez, J.-B. Martens, and H. de Ridder. Multidimensional Characterization of the Perceptual Quality of Noise-Reduced Computed Tomography Images.

Journal of Visual Communication and Image Representation, 6(4):317–334, 1995.

[42] P. J. D. Feyter, W. B. Meijboom, A. Weustink, C. V. Mieghem, N. R. A. Mollet, E. Vourvouri, K. Nieman, and F. Cademartiri. Spiral multislice computed tomography coronary angiography: a current status report. *Clinical Cardiology*, 30(9):437–442, 2007.

[43] C. Florin, R. Moreau-Gobard, and J. Williams. Automatic Heart Peripheral Vessels Segmentation Based On Normal MIP ray casting technique. In *Proceedings of MICCAI*, volume 1, pages 483–490, 2004.

[44] C. Florin, N. Paragios, and J. Williams. Particle filters, a quasi-Monte Carlo solution for segmentation of coronaries. In *Proceedings of MICCAI*, pages 246–253, 2005.

[45] C. Florin, N. Paragios, and J. Williams. Globally Optimal Active Contours, Sequential Monte Carlo and On-Line Learning for Vessel Segmentation. In *Proceedings of ECCV*, pages 476–489, 2006.

[46] A. Frangakis, A. Stoschek, and R. Hegerl. Wavelet transform filtering and nonlinear anisotropic diffusion assessed for signal reconstruction performance on multidimensional biomedical data. *IEEE Transactions on Biomedical Engineering*, 48(2):213–222, 2001.

[47] A. F. Frangi, W. J. Niessen, K. L. Vincken, and M. A. Viergever. Muliscale Vessel Enhancement Filtering. In *Proceedings of MICCAI*, pages 130–137, 1998.

[48] O. Friman, M. Hindennach, C. Kühnel, and H.-O. Peitgen. Multiple hypothesis template tracking of small 3D vessel structures. *Medical Image Analysis*, 14(2):160–171, 2010.

[49] O. Friman, C. Kühnel, and H.-O. Peitgen. Coronary Artery Centerline Extraction Using Multiple Hypothesis Tracking and Minimal Paths. *The Midas Journal*, 2008.

[50] M. L. Giger, H.-P. Chan, and J. Boone. Anniversary paper: History and status of CAD and quantitative image analysis: the role of Medical Physics and AAPM. *Medical Physics*, 35(12):5799–5820, 2008.

[51] N. J. Gordon, D. J. Salmond, and A. F. M. Smith. Novel approach to nonlinear/non-Gaussian Bayesian state estimation. In *Proceedings of Radar and Signal Processing*, volume 140, pages 107–113, 1993.

[52] M. A. Gulsun and H. Tek. Robust Tree Modeling. In *Proceedings of MICCAI*, 2008.

[53] K. M. Hanson. Detectability in computed tomographic images. *Medical Physics*, 6(5):441–451, 1979.

[54] K. M. Hanson. *Radiology of the Skull and Brain, Vol. 5: Technical Aspects of Computed Tomography*, chapter 113, Noise and computed contrast discrimination in tomography, pages 3941–3955. 1981.

[55] T. Hastie, R. Tibshirani, and J. Friedman. *The elements of statistical learning*. 2009.

[56] M. Hayakawa, K. Katada, H. Anno, S. Imizu, J. Hayashi, K. Irie, M. Negoro, Y. Kato, T. Kanno, and H. Sano. CT angiography with electrocardiographically gated reconstruction for visualizing pulsation of intracranial aneurysms: identification of aneurys-

mal protuberance presumably associated with wall thinning. *American Journal of Neuroradiology*, 26(6):1366–1369, 2005.

[57] H. S. Hecht. Applications of multislice coronary computed tomographic angiography to percutaneous coronary intervention: how did we ever do without it? *Catheterization and Cardiovascular Interventions*, 71(4):490–503, 2008.

[58] A. Hennemuth, S. Bock, T. Boskamp, D. Fritz, C. Kühnel, D. Rinck, M. Scheuering, and H. O. Peitgen. One-click coronary tree segmentation in CT angiographic images. *Computer Assisted Radiology and Surgery*, 1:317–321, 2005.

[59] T. Henzler, S. Oberoi, M. Weininger, P. Costello, C. Thilo, and U. Schoepf. Reproducibility of Non-Calcified Coronary Artery Plaque Burden Assessment Across Different Image Analysis Platforms. In *Proceedings of Annual Scientific Meeting of the Society of Cardiovascular Computed Tomography*, 2010.

[60] M. Hernández Hoyos, M. Orkisz, P. C. Douek, and I. E. Magnin. Assessment of carotid artery stenoses in 3D contrast-enhanced magnetic resonance angiography, based on improved generation of the centerline. *Machine Graphics & Vision*, 14(4):349–378, 2005.

[61] M. Hernández Hoyos, M. A. Zuluaga, M. Lozano, J. C. Prieto, P. C. Douek, I. E. Magnin, and M. Orkisz. Coronary centerline tracking in CT images with use of an elastic model and image moments. *The Midas Journal*, 2008.

[62] M. Hilts and C. Duzenli. Image filtering for improved dose resolution in CT polymer gel dosimetry. *Medical Physics*, 31(1):39–49, 2004.

[63] M. Hilts and C. Duzenli. Image noise in X-ray CT polymer gel dosimetry. *Journal of Physics: Conference Series*, 3:252–256, 2004.

[64] O. Honda et al. Comparison of Quality of Multiplanar Reconstructions and Direct Coronal Multidetector CT Scans of the Lung. *American Journal of Roentgenology*, 179(4):875–879, 2002.

[65] M.-K. Hong et al. The site of plaque rupture in native coronary arteries: a three-vessel intravascular ultrasound analysis. *Journal of the American College of Cardiology*, 46(2):261–265, 2005.

[66] A. Hoover, G. Jean-Baptiste, X. Jiang, P. Flynn, H. Bunke, D. Goldgof, K. Bowyer, D. Eggert, A. Fitzgibbon, and R. Fisher. An experimental comparison of range image segmentation algorithms. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 18(7):673–689, 1996.

[67] G. N. Hounsfield. Computerized transverse axial scanning (tomography). 1. Description of system. *The British Journal of Radiology*, 46(552):1016–1022, 1973.

[68] M. Jiang, J. Q, and M. BF. Automated extraction of fine features of kinetochore microtubules and plus-ends from electron tomography volume. *IEEE Transactions on Image Processing*, 15(7), 2006.

[69] T. R. C. Johnson, B. Krauss, M. Sedlmair, M. Grasruck, H. Bruder, D. Morhard, C. Fink, S. Weckbach, M. Lenhard, B. Schmidt, T. Flohr, M. F. Reiser, and C. R. Becker. Material differentiation by dual energy CT: initial experience. *European Radiology*, 17(6):1510–

1517, 2007.

[70] P. Judy, R. Swensson, and M. Szulc. Lesion detection and signal-to-noise ratio in CT images. *Medical Physics*, 8(1):13–23, 1981.

[71] M. K. Kalra, M. M. Maher, D. V. Sahani, M. A. Blake, P. F. Hahn, G. B. Avinash, T. L. Toth, E. Halpern, and S. Saini. Low-dose CT of the abdomen: evaluation of image improvement with use of noise reduction filters pilot study. *Radiology*, 228(1):251–256, 2003.

[72] M. K. Kalra, C. Wittram, M. M. Maher, A. Sharma, G. B. Avinash, K. Karau, T. L. Toth, E. Halpern, S. Saini, and J.-A. Shepard. Can Noise Reduction Filters Improve Low-Radiation-Dose Chest CT Images? Pilot Study. *Radiology*, 228(1):257–264, 2003.

[73] A. Kanitsar, D. Fleischmann, R. Wegenkittl, P. Felkel, and M. E. Gröller. CPR - Curved Planar Reformation. In *Proceedings of IEEE Visualization*, 2002.

[74] L. Keselbrener, Y. Shimoni, and S. Akselrod. Nonlinear filters applied on computerized axial tomography: Theory and phantom images. *Medical Physics*, 19(4):1057–1064, 1992.

[75] M. F. Khan, S. Wesarg, J. Gurung, S. Dogan, A. Maataoui, B. Brehmer, C. Herzog, H. Ackermann, B. Assmus, and T. J. Vogl. Facilitating coronary artery evaluation in MDCT using a 3D automatic vessel segmentation tool. *European Radiology*, 16(8):1789–1795, 2006.

[76] P. Kitslaar, M. Frenay, E. Oost, J. Dijkstra, B. Stoel, and J. H. Reiber. Connected Component and Morpholgy Based Extraction of Arterial Centerlines of the Heart (CocomoBeach). *The Midas Journal*, 2008.

[77] H. Knutsson and C.-F. Westin. Normalized and Differential Convolution: Methods for Interpolation and Filtering of Incomplete and Uncertain data. In *Proceedings of CVPR*, pages 515–523, 1993.

[78] P. Kohli and P. H. S. Torr. Dynamic Graph Cuts for Efficient Inference in Markov Random Fields. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 29(12):2079–2088, 2007.

[79] K. Krissian, H. Bogunovic, J. Pozo, M. Villa-Uriol, and A. Frangi. Minimally Interactive Knowledge-based Coronary Tracking in CTA using a Minimal Cost Path. *The Midas Journal*, 2008.

[80] A. Larralde, C. Boldak, M. Garreau, C. Toumoulin, D. Boulmier, and Y. Rolland. Evaluation of a 3D Segmentation Software for the Coronary Characterization in Multi-slice Computed Tomography. In *Proceedings of FIMH*, pages 1005–1005, 2003.

[81] G. Lavi, J. Lessick, P. Johnson, and D. Khullar. Single-seeded coronary artery tracking in CT angiography. In *Proceedings of NSSC*, 2004.

[82] A. W. Leber et al. Accuracy of 64-slice computed tomography to classify and quantify plaque volumes in the proximal coronary system: a comparative study using intravascular ultrasound. *Journal of the American College of Cardiology*, 47(3):672–677, 2006.

[83] J. Lee. Digital Image Smoothing and the Sigma Filter. *Computer Vision, Graphics and*

Image Processing, 24:255–269, 1983.

[84] D. Lesage, E. Angelini, I. Bloch, and G. Funka-Lea. Medial-based Bayesian tracking for vascular segmentation: Application to coronary arteries in 3D CT angiography. In *Proceedings of ISBI*, pages 268–271, 2008.

[85] D. Lesage, E. Angelini, I. Bloch, and G. Funka-Lea. A review of 3D vessel lumen segmentation techniques: Models, features and extraction schemes. *Medical Image Analysis*, 2009.

[86] D. Lesage, E. Angelini, I. Bloch, and G. Funka-Lea. Bayesian maximal paths for coronary artery segmentation from 3d ct angiograms. In *Proceedings of MICCAI*, pages 222–229, 2009.

[87] S. Leschka, H. Alkadhi, A. Plass, L. Desbiolles, J. Grünenfelder, B. Marincek, and S. Wildermuth. Accuracy of MSCT coronary angiography with 64-slice technology: first experience. *European Heart Journal*, 26(15):1482–1487, 2005.

[88] H. Li and A. Yezzi. Vessels as 4-D Curves: Global Minimal 4-D Paths to Extract 3-D Tubular Surfaces and Centerlines. *IEEE Transactions on Medical Imaging*, 26:1213–1223, 2007.

[89] H. Li, A. Yezzi, and L. Cohen. 3D Multi-branch Tubular Surface and Centerline Extraction with 4D Iterative Key Points. In *Proceedings of MICCAI*, pages 1042–1050, 2009.

[90] P. Libby. Inflammation in atherosclerosis. *Nature*, 420(6917):868–874, 2002.

[91] T. Lindeberg. Edge detection and ridge detection with automatic scale selection. *International Journal of Computer Vision*, 30(2):117–154, 1998.

[92] D. Lloyd-Jones, R. Adams, M. Carnethon, G. D. Simone, T. B. Ferguson, K. Flegal, E. Ford, K. Furie, A. Go, K. Greenlund, N. Haase, S. Hailpern, M. Ho, V. Howard, B. Kissela, S. Kittner, D. Lackland, L. Lisabeth, A. Marelli, M. McDermott, J. Meigs, D. Mozaffarian, G. Nichol, C. O'Donnell, V. Roger, W. Rosamond, R. Sacco, P. Sorlie, R. Stafford, J. Steinberger, T. Thom, S. Wasserthiel-Smoller, N. Wong, J. Wylie-Rosett, Y. Hong, A. H. A. S. Committee, and S. S. Subcommittee. Heart disease and stroke statistics–2009 update: a report from the American Heart Association Statistics Committee and Stroke Statistics Subcommittee. *Circulation*, 119(3):480–486, 2009.

[93] C. Lorenz, S. Renisch, T. Schlathoelter, and T. Buelow. Simultaneous segmentation and tree reconstruction of the coronary arteries in MSCT images. In *Proceedings of SPIE Medical Imaging*, volume 5031, pages 167–177, 2003.

[94] C. Lorenz and J. von Berg. A comprehensive shape model of the heart. *Medical Image Analysis*, 10(4):657–670, 2006.

[95] L. Lorigo, O. Faugeras, W. Grimson, R. Keriven, R. Kikinis, A. Nabavi, and C. Westin. Curves: Curve evolution for vessel segmentation. *Medical Image Analysis*, 5(3):195–206, 2001.

[96] M. A. Luengo-Oroz, M. J. Ledesma-Carbayo, J. J. Gómez-Diego, M. A. García-Fernández, M. Desco, and A. Santos. Extraction of the Coronary Artery Tree in Cardiac Computer Tomographic Images Using Morphological Operators. In *Proceedings*

of FIMH, pages 424–432, 2007.

[97] R. Manniesing and W. Niessen. Multiscale Vessel Enhancing Diffusion in CT Angiography Noise Filtering. In *Proceedings of IPMI*, volume 3565 of *Lecture Notes in Computer Science*, pages 138–149, 2005.

[98] R. Manniesing, M. Viergever, and W. Niessen. Vessel Enhancing Diffusion - A Scale Space Representation of Vessel Structures. *Medical Image Analysis*, 10(6):815–25, 2006.

[99] R. Manniesing, M. Viergever, and W. Niessen. Vessel Axis Tracking Using Topology Constrained Surface Evolution. *IEEE Transaction on Medical Imaging*, 26(3):309–316, 2007.

[100] E. Marieb and K. Hoehn. *Human anatomy & physiology*. 2007.

[101] H. A. Marquering, J. Dijkstra, P. J. H. de Koning, B. C. Stoel, and J. H. C. Reiber. Towards quantitative analysis of coronary CTA. *International Journal of Cardiovascular Imaging*, 21(1):73–84, 2005.

[102] D. Martin, C. Fowlkes, D. Tal, and J. Malik. A Database of Human Segmented Natural Images and its Application to Evaluating Segmentation Algorithms and Measuring Ecological Statistics. In *Proceedings of ICCV*, volume 2, pages 416–423, 2001.

[103] J. Melonakos, E. Pichon, S. Angenent, and A. Tannenbaum. Finsler active contours. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 30(3):412–423, 2008.

[104] C. Metz, M. Schaap, A. van der Giessen, T. van Walsum, and W. Niessen. Semi-automatic coronary artery centerline extraction in computed tomography angiography data. In *Proceedings of ISBI*, pages 856–859, 2007.

[105] C. Metz, M. Schaap, T. van Walsum, and W. Niessen. Two Point Minimum Cost Path Approach for CTA Coronary Centerline Extraction. *The Midas Journal*, 2008.

[106] C. Metz, M. Schaap, T. van Walsum, A. van der Giessen, A. Weustink, N. Mollet, G. Krestin, and W. Niessen. 3D segmentation in the clinic: A Grand Challenge II - Coronary Artery Tracking. *The Midas Journal*, 2008 MICCAI Workshop - Grand Challenge Coronary Artery Tracking:<http://hdl.handle.net/10380/1399>, 2008.

[107] C. Metz, M. Schaap, A. Weustink, N. Mollet, T. van Walsum, and W. Niessen. Coronary centerline extraction from CT coronary angiography images using a minimum cost path approach. *Medical physics*, 36:5568, 2009.

[108] N. R. Mollet, F. Cademartiri, and P. J. de Feyter. Non-invasive multislice CT coronary imaging. *Heart*, 91(3):401–407, 2005.

[109] G. Mowatt, E. Cummins, N. Waugh, S. Walker, J. Cook, X. Jia, G. S. Hillis, and C. Fraser. Systematic review of the clinical effectiveness and cost-effectiveness of 64-slice or higher computed tomography angiography as an alternative to invasive coronary angiography in the investigation of coronary artery disease. *Health Technology Assessment*, 12(17):iii–iv, ix–143, 2008.

[110] D. Mozaffarian, P. W. F. Wilson, and W. B. Kannel. Beyond established and novel risk factors: lifestyle risk factors for cardiovascular disease. *Circulation*, 117(23):3031–3038, 2008.

[111] P. Mrázek and M. Navara. Consistent Positive Directional Splitting of Anisotropic Diffusion. In *Proceedings of Computer Vision Winter Workshop*, pages 37–48, 2001.

[112] B. E. Muhs, A. Teutelink, M. Prokop, K. L. Vincken, F. L. Moll, and H. J. M. Verhagen. Endovascular aneurysm repair alters renal artery movement: a preliminary evaluation using dynamic CTA. *Journal of Endovascular Therapy*, 13(4):476–480, 2006.

[113] M. E. Mullins, M. H. Lev, P. Bove, C. E. O'Reilly, S. Saini, J. T. Rhea, J. H. Thrall, G. J. Hunter, L. M. Hamberg, and R. G. Gonzalez. Comparison of image quality between conventional and low-dose nonenhanced head CT. *American Journal of Neuroradiology*, 25(4):533–538, 2004.

[114] D. Myatt, S. Nasuto, and S. Maybank. Towards the automatic reconstruction of dendritic trees using particle filters. In *Proceedings of Linear Statistical Signal Processing Workshop*, 2006.

[115] E. A. Nadaraya. On estimating regression. *Theory of Probability and its Applications*, 10:186Ű190, 1964.

[116] D. Nain, A. Yezzi, and G. Turk. Vessel Segmentation Using a Shape Driven Flow. In *Proceedings of MICCAI*, pages 51–59, 2004.

[117] O. Nemitz, M. Rumpf, T. Tasdizen, and R. Whitaker. Anisotropic Curvature Motion for Structure Enhancing Smoothing of 3D MR Angiography Data. *Journal of Mathematical Imaging and Vision*, 27(3):217–229, 2007.

[118] K. Nieman, T. W. Galema, L. A. Neefjes, A. C. Weustink, P. Musters, A. D. Moelker, N. R. Mollet, R. de Visser, E. Boersma, and P. J. de Feijter. Comparison of the value of coronary calcium detection to computed tomographic angiography and exercise testing in patients with chest pain. *American Journal of Cardiology*, 104(11):1499–1504, 2009.

[119] B. Ohnesorge, T. Flohr, C. Becker, A. Kopp, U. Schoepf, U. Baum, A. Knez, K. Klingenbeck-Regn, and M. Reiser. Cardiac imaging by means of electrocardiographically gated multisection spiral CT: initial experience. *Radiology*, 217(2):564–571, 2000.

[120] M. Okada. Noise evaluation and filter design in CT images. *IEEE Transactions on Biomedical Engineering*, 32(9):713–719, 1985.

[121] S. Olabarriaga, M. Breeuwer, and W. Niessen. Minimum Cost Path Algorithm for Coronary Artery Central Axis Tracking in CT images. In *Proceedings of MICCAI*, 2003.

[122] M. Oudkerk, A. E. Stillman, S. S. Halliburton, W. A. Kalender, S. Möhlenkamp, C. H. McCollough, R. Vliegenthart, L. J. Shaw, W. Stanford, A. J. Taylor, P. M. A. van Ooijen, L. Wexler, P. Raggi, E. S. of Cardiac Radiology, and N. A. S. for Cardiovascular Imaging. Coronary artery calcium screening: current status and recommendations from the European Society of Cardiac Radiology and North American Society for Cardiovascular Imaging. *European Radiology*, 18(12):2785–2807, 2008.

[123] J. D. Owens, D. Luebke, N. Govindaraju, M. Harris, J. Krüger, A. E. Lefohn, and T. J. Purcell. A Survey of General-Purpose Computation on Graphics Hardware. In *Proceedings of Eurographics*, pages 21–51, 2005.

[124] I. Pantos, E. P. Efstathopoulos, and D. G. Katritsis. Two and three-dimensional quantitative coronary angiography. *Cardiology Clinics*, 27(3):491–502, 2009.

[125] H. H. Pien, A. J. Fischman, J. H. Thrall, and A. G. Sorensen. Using imaging biomarkers to accelerate drug development and clinical trials. *Drug Discovery Today*, 10(4):259–266, 2005.

[126] M. J. D. Powell. An efficient method for finding the minimum of a function of several variables without calculating derivatives. *Computer Journal*.

[127] S. Ramcharitar, F. Pugliese, C. Schultz, J. Ligthart, P. de Feyter, H. Li, N. Mollet, M. van de Ent, P. W. Serruys, and R. J. van Geuns. Integration of multislice computed tomography with magnetic navigation facilitates percutaneous coronary interventions without additional contrast agents. *Journal of the American College of Cardiology*, 53(9):741–746, 2009.

[128] J. H. Reiber, C. J. Kooijman, C. J. Slager, J. J. Gerbrands, J. C. Schuurbiers, A. D. Boer, W. Wijns, P. W. Serruys, and P. G. Hugenholtz. Coronary artery dimensions from cineangiograms methodology and validation of a computer-assisted analysis procedure. *IEEE Transactions on Medical Imaging*, 3(3):131–141, 1984.

[129] F. Renard and Y. Yang. Image analysis for detection of coronary artery soft plaques in MDCT images. In *Proceedings of ISBI*, pages 25–28, 2008.

[130] E. Ricci and R. Perfetti. Retinal blood vessel segmentation using line operators and support vector classification. *IEEE Transactions on Medical Imaging*, 26(10):1357–1365, 2007.

[131] W. Richter. Imaging biomarkers as surrogate endpoints for drug development. *European Journal of Nuclear Medicine and Molecular Imaging*, 33:6–10, 2006.

[132] E. Rollano-Hijarrubia, R. Stokking, F. van der Meer, and W. J. Niessen. Imaging of small high-density structures in CT; A phantom study. *Academic Radiology*, 13(7):893–908, 2006.

[133] D. Ropers, J. Rixe, K. Anders, A. Küttner, U. Baum, W. Bautz, W. G. Daniel, and S. Achenbach. Usefulness of multidetector row spiral computed tomography with 64- x 0.6-mm collimation and 330-ms rotation for the noninvasive detection of significant coronary artery stenoses. *American Journal of Cardiology*, 97(3):343–348, 2006.

[134] W. Rosamond et al. Heart disease and stroke statistics–2008 update: a report from the American Heart Association Statistics Committee and Stroke Statistics Subcommittee. *Circulation*, 117(4):e25–146, 2008.

[135] A. Rose. *Vision: Human and Electronic*. 1973.

[136] M. Rumpf and R. Strzodka. Graphics Processor Units: New Prospects for Parallel Computing. In *Numerical Solution of Partial Differential Equations on Parallel Computers*, volume 51 of *Lecture Notes in Computational Science and Engineering*, pages 89–134. 2005.

[137] R. Rutherford, B. Pullan, and I. Isherwoord. Measurement of effective atomic number and electron density using an EMI scanner. *American Journal of Neuroradiology*, 11:15–21, 1976.

[138] F. J. Rybicki, H. J. Otero, M. L. Steigner, G. Vorobiof, L. Nallamshetty, D. Mitsouras, H. Ersoy, R. T. Mather, P. F. Judy, T. Cai, K. Coyner, K. Schultz, A. G. Whitmore, and M. F. D. Carli. Initial evaluation of coronary images from 320-detector row computed tomography. *International Journal of Cardiovascular Imaging*, 24(5):535–546, 2008.

[139] E. Saff and A. Kuijlaars. Distributing Many Points on a Sphere. *The Mathematical Intelligencer*, 19(1):5–11, 1997.

[140] J. Saragih, S. Lucey, and J. Cohn. Face Alignment through Subspace Constrained Mean-Shifts . In *Proceedings of ICCV*, 2009.

[141] M. Schaap, R. Manniesing, I. Smal, T. Van Walsum, A. Van Der Lugt, and W. Niessen. Bayesian tracking of tubular structures and its application to carotid arteries in CTA. In *Proceedings of MICCAI*, pages 562–570, 2007.

[142] M. Schaap, C. Metz, T. van Walsum, A. van der Giessen, A. Weustink, N. Mollet, C. Bauer, H. Bogunović, C. Castro, X. Deng, et al. Standardized evaluation methodology and reference database for evaluating coronary artery centerline extraction algorithms. *Medical Image Analysis*, 13(5):701–714, 2009.

[143] M. Schaap, L. Neefjes, C. Metz, A. van der Giessen, A. Weustink, N. Mollet, J. Wentzel, T. van Walsum, and W. Niessen. Coronary Lumen Segmentation using Graph Cuts and Robust Kernel Regression. In *Proceedings of IPMI*, pages 528–539, 2009.

[144] M. Schaap, I. Smal, C. Metz, T. van Walsum, and W. Niessen. Bayesian Tracking of Elongated Structures in 3D Images. In *Proceedings of IPMI*, volume 4584, pages 74–85, 2007.

[145] D. Scharstein and R. Szeliski. A taxonomy and evaluation of dense two-frame stereo correspondence algorithms. *International Journal of Computer Vision*, 47:7–42, 2002.

[146] T. Schepis, O. Gaemperli, P. Koepfli, I. Valenta, K. Strobel, A. Brunner, S. Leschka, L. Desbiolles, L. Husmann, H. Alkadhi, and P. A. Kaufmann. Comparison of 64-slice CT with gated SPECT for evaluation of left ventricular function. *Journal of Nuclear Medicine*, 47(8):1288–1294, 2006.

[147] A. Schilham, B. van Ginneken, H. Gietema, and M. Prokop. Local Noise Weighted Filtering for Emphysema Scoring of Low-dose CT Images. *IEEE Transactions on Medical Imaging*, 25:451–463, 2006.

[148] S. Schroeder, A. F. Kopp, A. Baumbach, C. Meisner, A. Kuettner, C. Georg, B. Ohnesorge, C. Herdeg, C. D. Claussen, and K. R. Karsch. Noninvasive detection and evaluation of atherosclerotic coronary plaques with multislice computed tomography. *Journal of the American College of Cardiology*, 37(5):1430–1435, 2001.

[149] B. Settles. Active Learning Literature Survey. *Machine Learning*, 15(2):201–221, 1994.

[150] H. Shim, I. D. Yun, K. M. Lee, and S. U. Lee. Partition-Based Extraction of Cerebral Arteries from CT Angiography with Emphasis on Adaptive Tracking. In *Proceedings of IPMI*, volume 3565 of *LNCS*, pages 357–368, 2005.

[151] N. L. Smith, J. F. Felix, A. C. Morrison, S. Demissie, N. L. Glazer, L. R. Loehr, L. A. Cupples, A. Dehghan, T. Lumley, W. D. Rosamond, W. Lieb, F. Rivadeneira, J. C. Bis, A. R.

Folsom, E. Benjamin, Y. S. Aulchenko, T. Haritunians, D. Couper, J. Murabito, Y. A. Wang, B. H. Stricker, J. S. Gottdiener, P. P. Chang, T. J. Wang, K. M. Rice, A. Hofman, S. R. Heckbert, E. R. Fox, C. J. O'Donnell, A. G. Uitterlinden, J. I. Rotter, J. T. Willerson, D. Levy, C. M. van Duijn, B. M. Psaty, J. C. M. Witteman, E. Boerwinkle, and R. S. Vasan. Association of genome-wide variation with the risk of incident heart failure in adults of European and African ancestry: a prospective meta-analysis from the cohorts for heart and aging research in genomic epidemiology (CHARGE) consortium. *Circulation: Cardiovascular Genetics*, 3(3):256–266, 2010.

[152] M. Sonka, M. D. Winniford, and S. M. Collins. Robust simultaneous detection of coronary borders in complex images. *IEEE Transactions on Medical Imaging*, 14(1):151–161, 1995.

[153] J. Staal, M. Abràmoff, M. Niemeijer, M. Viergever, and B. van Ginneken. Ridge-based vessel segmentation in color images of the retina. *IEEE Transactions on Medical Imaging*, 23(4):501–509, 2004.

[154] S. Suryanarayanan, R. Mullick, Y. Mallya, V. Kamath, and N. Nagaraj. Automatic Partitioning of Head CTA for enabling Segmentation. In *Proceedings of SPIE Medical Imaging*, volume 5370, pages 410–419, 2004.

[155] A. Szymczak. Vessel Tracking by Connecting the Dots. *The Midas Journal*, 2008.

[156] A. Szymczak, A. Stillman, A. Tannenbaum, and K. Mischaikow. Coronary vessel trees from 3D imagery: a topological approach. *Medical Image Analysis*, 10(4):548–559, 2006.

[157] H. Tek, M. A. Gulsun, S. Laguitton, L. Grady, D. Lesage, and G. Funka-Lea. Automatic Coronary Tree Modeling. *The Midas Journal*, 2008.

[158] N. A. Thacker, F. J. Aherne, and P. I. Rockett. The Bhattacharyya Metric as an Absolute Similarity Measure for Frequency Coded Data. In *Proceedings of TIPR*, volume 34, pages 363–368, 1998.

[159] S. Tu, Z. Huang, G. Koning, K. Cui, and J. H. C. Reiber. A novel three-dimensional quantitative coronary angiography system: In-vivo comparison with intravascular ultrasound for assessing arterial segment length. *Catheterization and Cardiovascular Interventions*, 76(2):291–298, 2010.

[160] J. Tyrrell, E. di Tomaso, D. Fuja, R. Tong, K. Kozak, R. Jain, and B. Roysam. Robust 3-D modeling of vasculature imagery using superellipsoids. *IEEE Transactions on Medical Imaging*, 26(2):223–237, 2007.

[161] R. S. Uppoor, P. Mummaneni, E. Cooper, H. H. Pien, A. G. Sorensen, J. Collins, M. U. Mehta, and S. U. Yasuda. The use of imaging in the early development of neuropharmacological drugs: a survey of approved NDAs. *Clinical Pharmacology and Therapeutics*, 84(1):69–74, 2008.

[162] B. van Beers, C. Cuenod, L. Marti-Bonmati, C. Matos, W. Niessen, and A. Padhani. White paper on imaging biomarkers. *White paper of the European Society of Radiology (ESR)*, 2010.

[163] R. E. van Gelder, H. W. Venema, J. Florie, C. Y. Nio, I. W. O. Serlie, M. P. Schutter, J. C. van Rijn, F. M. Vos, A. S. Glas, P. M. M. Bossuyt, J. F. W. Bartelsman, J. S. Laméris, and

J. Stoker. CT colonography: feasibility of substantial dose reduction–comparison of medium to very low doses in identical patients. *Radiology*, 232(2):611–20, 2004.

[164] B. van Ginneken, T. Heimann, and M. Styner. 3D Segmentation in the Clinic: A Grand Challenge. In *Proceedings of 3D Segmentation in the Clinic: A Grand Challenge, MICCAI workshop*, pages 7–15, 2007.

[165] T. van Walsum, M. Schaap, C. Metz, A. van der Giessen, and W. Niessen. Averaging Center Lines: Mean Shift on Paths. In *Proceedings of MICCAI*, 2008.

[166] F. Voci, S. Eiho, N. Sugimote, and H. Seikguchi. Estimating the gradient threshold in the Perona-Malik equation. *IEEE Signal Processing magazine*.

[167] C. Wang and O. Smedby. Coronary Artery Segmentation and Skeletonization Based on Competing Fuzzy Connectedness Tree. In *Proceedings of MICCAI*, volume 4791 of *Lecture Notes in Computer Science*, pages 311–318, 2007.

[168] C. Wang and O. Smedby. An Automatic Seeding Method For Coronary Artery Segmentation and Skeletonization in CTA. The Midas Journal. *The Midas Journal*, 2008.

[169] C. Wang and O. Smedby. Integrating automatic and interactive methods for coronary artery segmentation: let the PACS workstation think ahead. *International Journal of Computer Assisted Radiology and Surgery*, 5(3):275–285, 2010.

[170] J. C. Wang, S.-L. T. Normand, L. Mauri, and R. E. Kuntz. Coronary artery spatial distribution of acute myocardial infarction occlusions. *Circulation*, 110(3):278–284, 2004.

[171] W. Wang, B. Jüttler, D. Zheng, and Y. Liu. Computation of rotation minimizing frames. *ACM Transactions on Graphics*, 27(1):1–18, 2008.

[172] S. Warfield, K. Zou, and W. Wells. Simultaneous truth and performance level estimation (STAPLE): an algorithm for the validation of image segmentation. *IEEE Transactions on Medical Imaging*, 23(7):903–921, 2004.

[173] J. Weickert. *Anisotropic diffusion in image processing,*. PhD thesis, Dept. of Mathematics, University of Kaiserslautern, Germany, 1996.

[174] J. Weickert. A Review of Nonlinear Diffusion Filtering. In *Proceedings of Scale-Space Theory in Computer Vision*, volume 1252, pages 3–28, 1997.

[175] J. Weickert, B. Romeny, and M. Viergever. Efficient and Reliable Schemes for Nonlinear Diffusion Filtering. *IEEE Transactions on Image Processing*, 7(3):398–410, 1998.

[176] E. W. Weisstein. Rotation Matrix. From MathWorld–A Wolfram Web Resource.

[177] S. Wesarg and E. Firle. Segmentation of Vessels: The Corkscrew Algorithm. In *Proceedings of SPIE, Medical Imaging*, volume 9, page 10, 2004.

[178] S. Wesarg, M. F. Khan, and E. A. Firle. Localizing calcifications in cardiac CT data sets using a new vessel segmentation approach. *Journal of Digital Imaging*, 19(3):249–257, 2006.

[179] J. West, J. M. Fitzpatrick, M. Y. Wang, B. M. Dawant, C. R. Maurer, R. M. Kessler, R. J. Maciunas, C. Barillot, D. Lemoine, A. Collignon, F. Maes, P. Suetens, D. Vandermeulen, P. A. van den Elsen, S. Napel, T. S. Sumanaweera, B. Harkness, P. F. Hemler, D. L.

Hill, D. J. Hawkes, C. Studholme, J. B. Maintz, M. A. Viergever, G. Malandain, and R. P. Woods. Comparison and evaluation of retrospective intermodality brain image registration techniques. *Journal of Computer Assisted Tomography*, 21(4):554–566, 1997.

[180] A. C. Weustink, N. R. Mollet, F. Pugliese, W. B. Meijboom, K. Nieman, M. H. Heijenbrok-Kal, T. G. Flohr, L. A. E. Neefjes, F. Cademartiri, P. J. de Feyter, and G. P. Krestin. Optimal electrocardiographic pulsing windows and heart rate: effect on image quality and radiation exposure at dual-source coronary CT angiography. *Radiology*, 248(3):792–798, 2008.

[181] O. Wink, W. Niessen, B. Verdonck, and M. Viergever. Vessel axis determination using wave front propagation analysis. In *Proceedings of MICCAI*, pages 845–853, 2001.

[182] World Health Organization. The world health report 2004 - changing history. Annex Table 2: Deaths by cause, sex and mortality stratum in WHO regions, estimates for 2002., 2004.

[183] World Health Organization. The top ten causes of death - Fact sheet 317, 2009.

[184] C. Xu and J. L. Prince. Snakes, Shapes, and Gradient Vector Flow. *IEEE Transactions on Image Processing*, 7(3):359–369, 1998.

[185] G. Yang, A. Bousse, C. Toumoulin, and H. Shu. A multiscale tracking algorithm for the coronary extraction in MSCT angiography. In *Proceedings of EMBS*, pages 3066–3069, 2006.

[186] Y. Yang, A. Tannenbaum, D. Giddens, and A. Stillman. Automatic Segmentation of Coronary Arteries Using Bayesian Driven Implicit Surfaces. In *Proceedings of ISBI*, pages 189–192, 2007.

[187] Y. Yang, L. Zhu, S. Haker, A. R. Tannenbaum, and D. P. Giddens. Harmonic skeleton guided evaluation of stenoses in human coronary arteries. In *Proceedings of MICCAI*, volume 8, pages 490–497, 2005.

[188] S. Zambal, J. Hladuvka, A. Kanitsar, and K. Bühler. Shape and Appearance Models for Automatic Coronary Artery Tracking. *The Midas Journal*, 2008.

[189] P. Zanzonico, L. N. Rothenberg, and H. W. Strauss. Radiation exposure of computed tomography and direct intracoronary angiography: risk has its reward. *Journal of the American College of Cardiology*, 47(9):1846–1849, 2006.

[190] Y. Zhang, K. Chen, and S. Wong. 3D Interactive Centerline Extraction. *The Midas Journal*, 2008.

[191] S. Zhou and D. Comaniciu. Shape regression machine. In *Proceedings of IPMI*, pages 13–25, 2007.

[192] H. Zhu, Z. Ding, R. N. Piana, T. R. Gehrig, and M. H. Friedman. Cataloguing the geometry of the human coronary arteries: A potential tool for predicting risk of coronary artery disease. *International Journal of Cardiology*, 2008.

[193] F. Zhuge, G. Rubin, S. Sun, and S. Napel. An abdominal aortic aneurysm segmentation method: level set with region and statistical information. *Medical physics*, 33:1440, 2006.

Color image section

Chapter 2 & 3

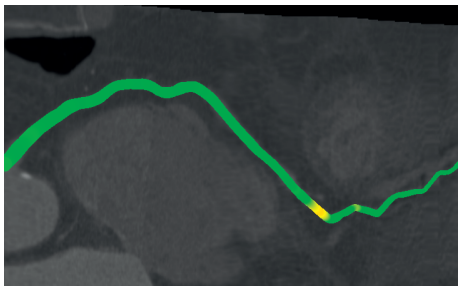

Figure 2.2 (see page 20): *An example of one of the color-coded curved planar reformatted images used to detect possible annotation errors.*

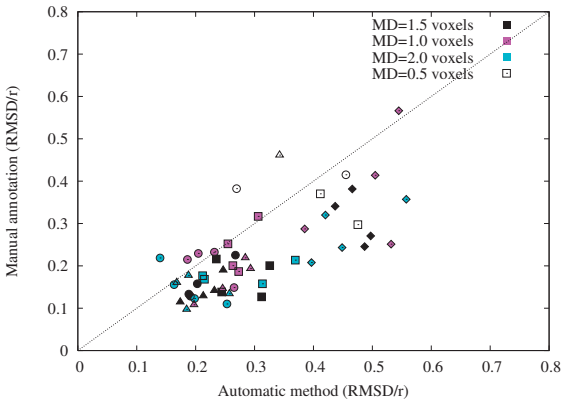

Figure 3.2 (see page 52): *Scatter plot demonstrating the relation between the accuracy, measured in RMSD/r, of the manually and automatically tracked tube centerlines. The different symbols denote different radii (◊ = 2.5, □ = 5.0, ◯ = 7.5, △ = 10.0, all in voxels), and different colors denote different noise levels.*

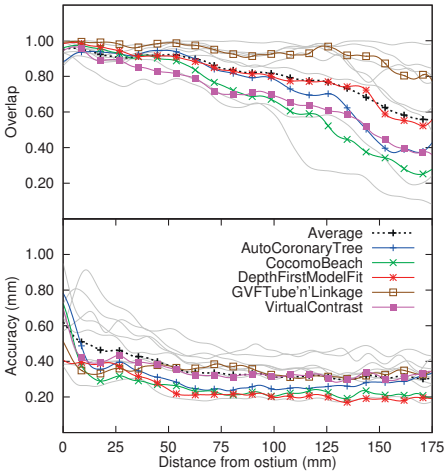

(a) Fully automatic coronary artery centerline extraction methods

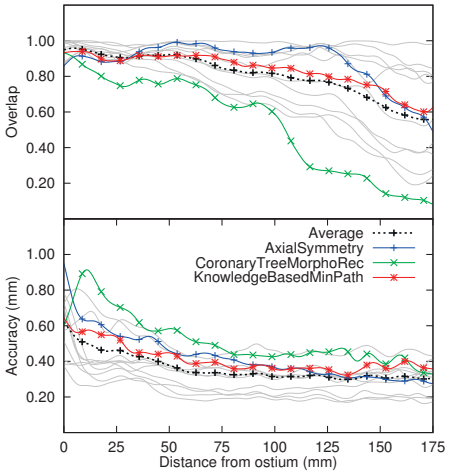

(b) Semi automatic coronary artery centerline extraction methods

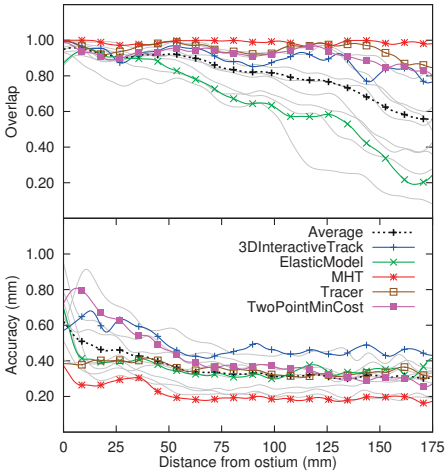

(c) Interactive coronary artery centerline extraction methods

Figure 2.7 (see page 32): *The algorithm performance of each method with respect to the distance from the ostium averaged over all 96 evaluated vessels over the first 175ṁm (only 10% of the vessels were longer than 175ṁm). Overlap: the fraction of points on the reference standard marked as true positive. Accuracy: the average distance to the centerline if the point is marked true positive. Each of the three graphs shows in light-gray the results of all the thirteen evaluated methods and in color the results of the respective algorithm category. The graphs also show in black the average accuracy and overlap for all thirteen evaluated methods.*

Chapter 4

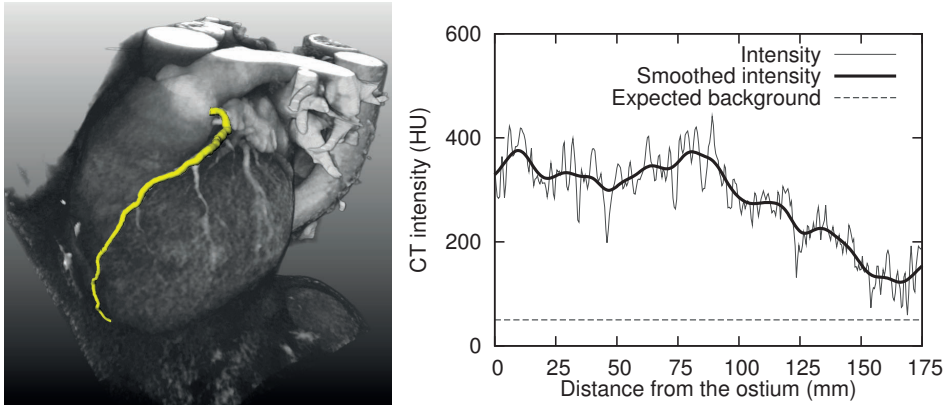

(a) (b)

Figure 4.1 (see page 60): *(a) A 3D rendering of a cardiac CTA dataset with in yellow a manually annotated Left Anterior Descending (LAD) Coronary Artery. (b) A graph of the CT intensities I_x along the centerline of the LAD, a graph of the intensities after Gaussian kernel regression $\hat{I}_x$ and the expected background intensity I_{bg} (see section 4.3.1).*

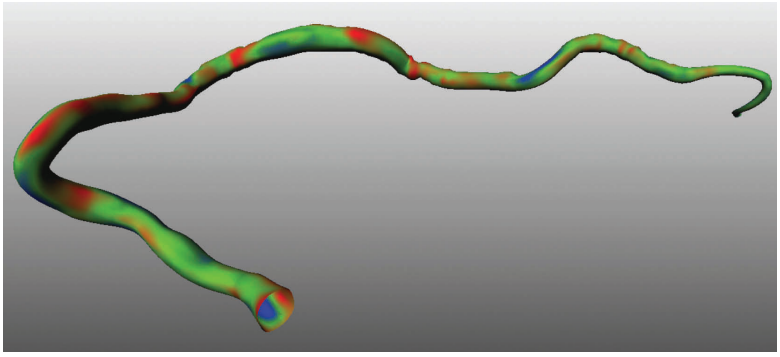

Figure 4.5 (see page 67): *A 3D example of a coronary segmentation color-coded with the distance to the reference standard. Red corresponds to the segmentation being locally 0.5 mm larger than the reference standard, green corresponds to a perfect fit, and blue corresponds to a 0.5 mm under-segmentation.*

Chapter 5

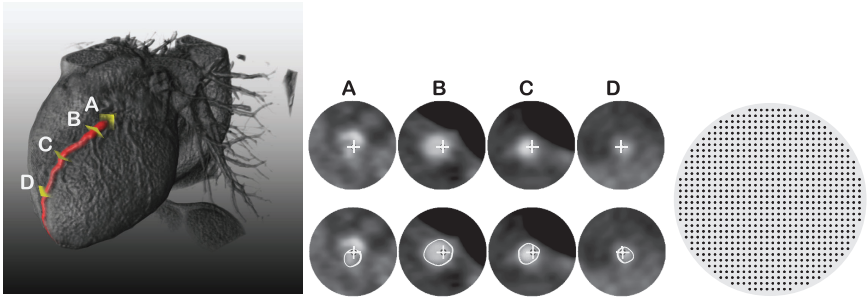

Figure 5.1 (see page 72): *Left: a 3-dimensional visualization of a cardiac CTA image with in red a coronary artery. Midle: four circular images sampled cross-sectionally to the coronary artery (sample radius=7.5 mm and sample distance=0.5 mm.) overlaid with the position of an initialization centerline and manually annotated coronary lumen contours. Right: The intensity sample pattern, as described in Section 5.2.1.*

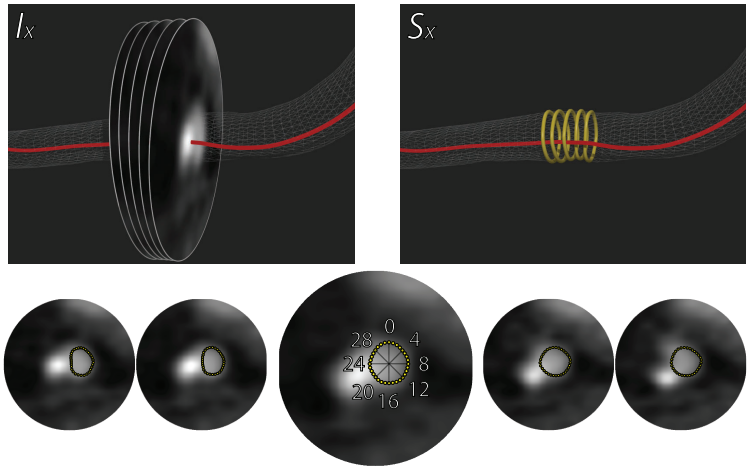

Figure 5.2 (see page 75): *A visualization of the intensity samples $\mathbf{I}_x$ and a local shape $\mathbf{S}_x$. Top left: a 3D surface of the lumen, the initialization centerline, and a tubular image patch $\mathbf{I}_x$ (see Section 5.2.3.2). Top right: the surface, initialization centerline, and a local shape $\mathbf{S}_x$ (see Section 5.2.2). Bottom: $\mathbf{I}_x$ overlaid with $\mathbf{S}_x$ and the rays (Equation 5.3) used for calculating the landmarks. The numbers indicate the cross-sectional landmark index d for $D = 32$. In the coarse segmentation $\mathbf{I}_x$ is used to estimate $\mathbf{S}_x$ with linear regression. During shape refinement all cross-sectional images of the vessel are used to estimate all cross-sectional shapes.*

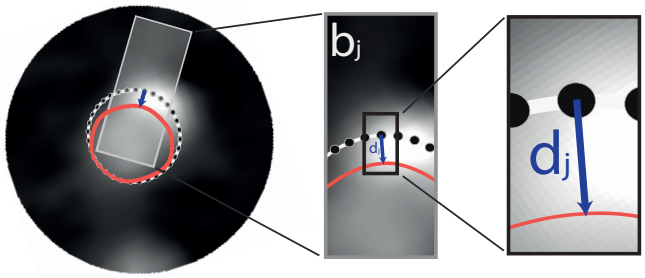

Figure 5.7 (see page 80): *A visualization of a boundary profile and displacement vector. Left: a cross-sectional image with the reference standard in solid orange, ridge regression result in black and white, the boundary profile* $\mathbf{b}_j$ *in transparent gray, and the displacement vector* $\mathbf{d}_j$ *in purple. Middle and right: zoomed-in versions of the left image.*

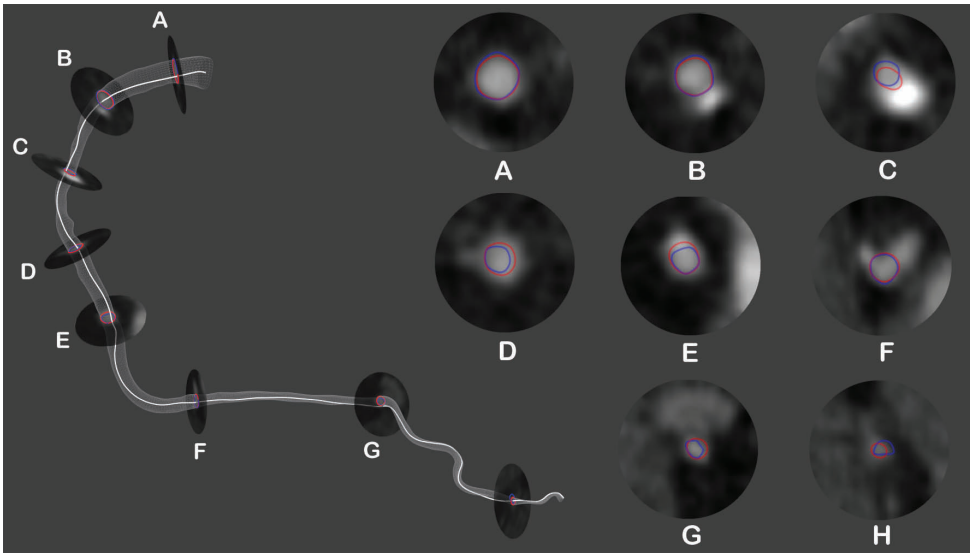

Figure 5.9 (see page 85): *A segmentation of the lumen of a right coronary artery showing a 3D mesh of the automatic segmentation, the initialization curve in white, and 8 selected cross-sections with results of automatic segmentation in red and reference standard in blue.*

Chapter 6

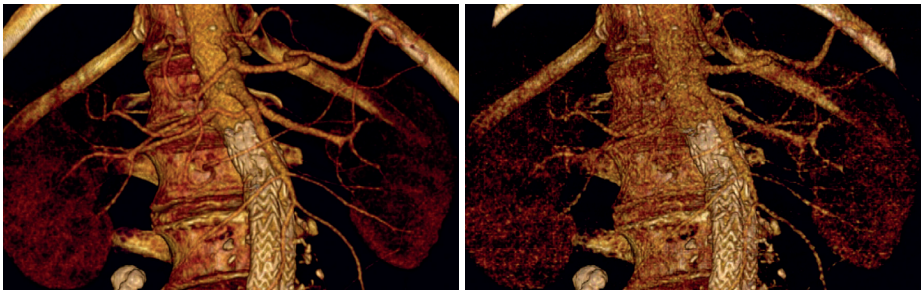

(a) (b)

Figure 6.2 (see page 104): *A section of a volume rendering of (a) gated and (b) ungated AAA CTA images. Note the difference in quality between, for example, the visualization of small vessels.*

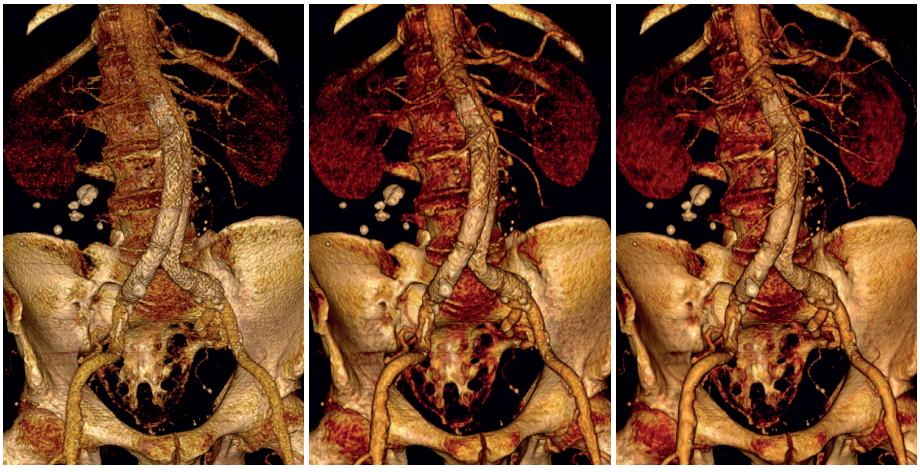

(a) (b) (c)

Figure 6.3 (see page 109): *(a) 3D visualization of a gated dataset. (b) Visualization after EED filtering. (c) Visualization after SCEED filtering. Note the improved visibility of the small vessels in the SCEED visualization in comparison with the EED visualization. The unimproved structural sharpness in the EED and SCEED scan can be seen in the region of the stent, located in the bifurcation of the abdominal aortic artery in the center of the image.*

Chapter 8

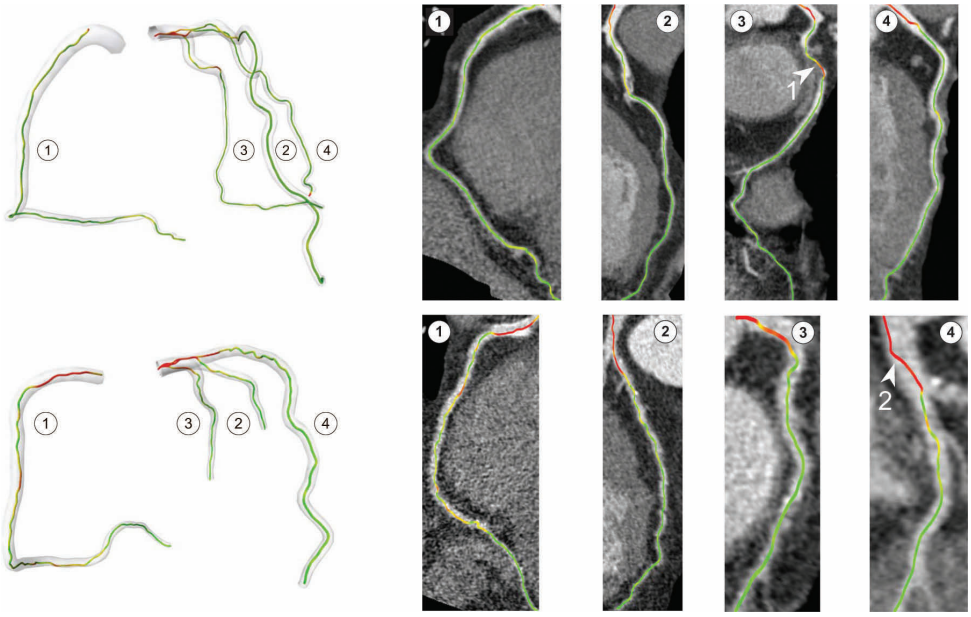

Figure 8.1 (see page 123): *Examples of successfully extracted coronary artery centerlines using the vesselness/intensity cost function. The first column shows the reference standard light gray tubes and the automatically extracted centerlines. The other columns show the automatically extracted centerlines projected onto CPR images based on the reference standard centerline. The greyscale coding of the extracted centerline indicates the distance from the automatically extracted centerline to the reference standard. Arrows in the CPR images indicate: 1) Successfully crossed low contrast region caused by extreme pathology, 2) decreased extraction accuracy at a location where the radius of the vessel is relatively large.*

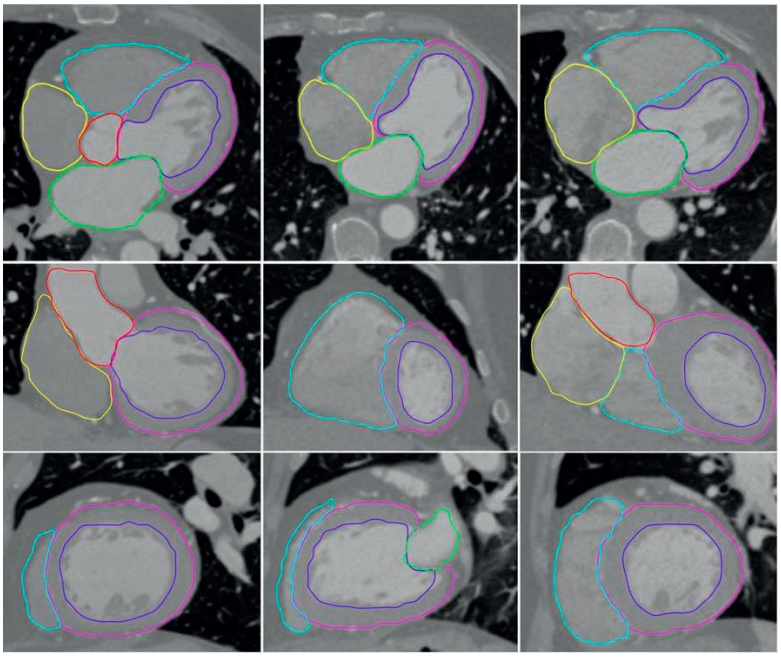

Figure 8.3 (see page 125): *Fully automatic segmentation results from three patients out of the 1380 CTA datasets. Axial view (top), coronal view (middle) and sagittal view (bottom). Segmentations: epicardium left ventricle (pink), endocaridum left ventricle (purple), right ventricle (blue), left atrium (green), right atrium (yellow) and aorta (red).*

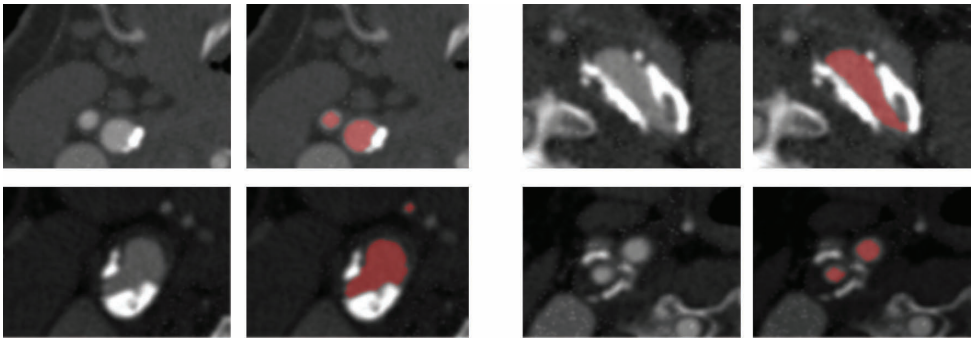

Figure 8.6 (see page 127): *Right: Representative examples of the segmentation results obtained by the proposed method in case of the presence of severe calcified and non-calcified (soft) plaque (rows 1 to 3) and stenosis (bottom row).*

Samenvatting voor niet-ingewijden

Introductie

De kransslagaders voorzien ons hart van bloed. Ziekte in deze bloedvaten kan daarom ernstige problemen veroorzaken. Op dit moment is kransslagaderziekte, ook wel coronaire atherosclerose genoemd, één van de belangrijkste doodsoorzaken. Wereldwijd sterven ongeveer 1 op de 5 mensen aan deze aandoening.

Dit proefschrift gaat over de ontwikkeling van beeldverwerkingsmethoden die gebruikt kunnen worden om snel en nauwkeurig atherosclerose te diagnostiseren op basis van medische beelden. Het grootste deel van dit proefschrift staat in het teken van automatische kwantificatie technieken met computer tomografie angiografie (CTA), een beeldvormende techniek voor 3-dimensionaal bloedvatonderzoek. Daarnaast is er aandacht voor verbeterde visualisatie van CTA-beelden. In deze samenvatting beschrijf ik, na een korte introductie, de bijdragen en uitkomsten van de verschillende hoofdstukken.

Atherosclerose van de kransslagaders

De wand van de kransslagaders bestaat uit drie lagen; een elastische buitenlaag, een gespierde tussenlaag en een dunne binnenwand. Ziekte aan de kransslagaders ontstaat door een chronische ontsteking tussen de binnen- en buitenwand. Door deze ontsteking hopen vetcellen zich op in de vaatwand. In een later stadium van de ziekte hoopt ook calcium zich op; daarom wordt coronaire atherosclerose ook wel aderverkalking genoemd. Na verloop van tijd kan de ophoping van plaque (vet en calcium) resulteren in een vernauwing van het bloedvat (zie figuur 1) en daarmee vermindert bloedvoorziening van het hart. Dit kan, met name bij inspanning, een drukkend en zwaar gevoel in het midden van de borst veroorzaken. Wanneer het bloedvat zover verstopt raakt dat het achterliggende hartspierweefsel niet meer van voldoende

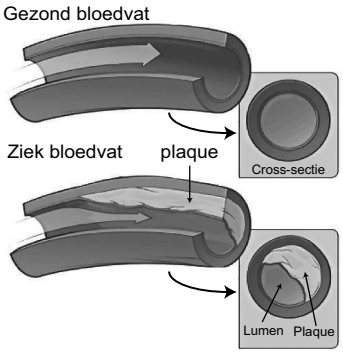

Figuur 1: *Schematische illustratie van atherosclerose. Boven: een gezond bloedvat. Onder: een bloedvat na plaque-ophoping.*

bloed wordt voorzien, kan dit resulteren in een hartinfarct (het afsterven van een deel van de hartspier). Het is ook mogelijk dat plaque zich ophoopt zonder dat het vat problematisch vernauwt. Als bij zo'n plaque-ophoping de binnenwand scheurt komt er ineens plaque in het bloedvat en dit resulteert vaak in een complete verstopping van het bloedvat en in een acuut hartinfarct.

Computer Tomografie Angiografie (CTA)

Computer Tomografie Angiografie (CTA) is op dit moment de meest nauwkeurige minimaal-invasieve (geen ingreep vereist) beeldvormende techniek voor de diagnose van coronaire atherosclerose. Het is met deze techniek mogelijk om te meten hoe vernauwd een kransslagader is en wat de samenstelling van de plaque is. Na injectie van contrastvloeistof bij een patiënt maakt een CT-scanner vanuit vele hoeken rond het lichaam röntgenfoto's en combineert die vervolgens met een geavanceerd computerprogramma tot een 3-dimensionaal beeld. Figuur 2 toont een CT-scanner en figuur 3 een voorbeeld van een CTA-beeld van het hart.

Figuur 2: *Een CT-scanner (SOMATOM Definition Flash, Siemens Healthcare).*

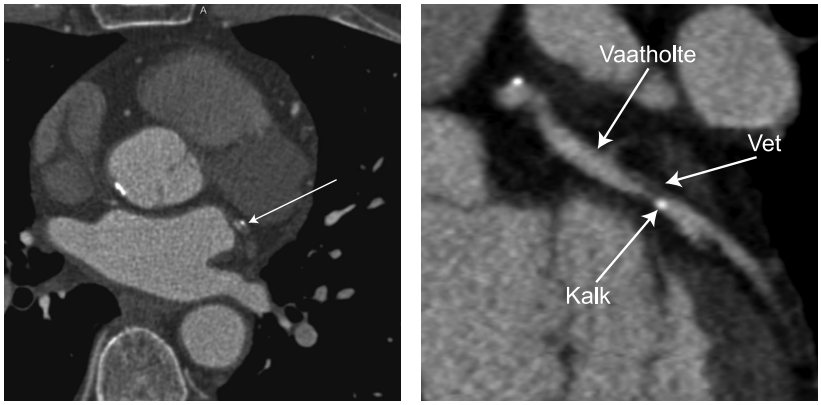

Figuur 3: *Links: een CTA-beeld van het hart. De pijl wijst naar een kransslagader met plaque. Rechts: een uitvergrote (en gedraaide) versie van de kransslagader in het linker beeld.*

Kwantitatieve beeldanalyse van CTA-beelden

Beeldanalyse van CTA kan gebruikt worden in de diagnose van coronaire atherosclerose door de verschijningsvorm van atherosclerose (bijvoorbeeld de mate van vaatvernauwing en de hoeveelheid plaque) te kwantificeren. Het grote voordeel van automatische kwantitatieve beeldanalyse, in vergelijking met manuele kwantificatie, is dat een computermethode altijd hetzelfde rapporteert, niet vermoeid raakt, en mogelijk zelfs nauwkeuriger kan zijn. Hierdoor kan de diagnose sneller en goedkoper worden. Daarnaast is het mogelijk om grote hoeveelheden data snel te analyseren. Hierdoor kunnen bijvoorbeeld relaties gevonden worden tussen bestaande risicofactoren (zoals roken en overgewicht) en coronaire atherosclerose, maar het is ook mogelijk om efficiënt onderzoek te doen naar nieuwe risicofactoren.

In dit proefschrift wordt kwantitatieve beeldanalyse van CTA-scans van het hart in twee stappen gedaan: eerst analyseert een computerprogramma op een grove manier de CTA-scan en bepaalt het de positie van de bloedvaten. Daarna volgt een meer gedetailleerde stap waarin het programma de beeldinformatie rondom de bloedvaten gebruikt om exacte informatie over de bloedvaten te verkrijgen en deze informatie uit te drukken in getallen. Beide stappen zijn complexe taken voor een computer omdat cardiale CTA-beelden door de beweging van het hart en de beperkte hoeveelheid straling die toegediend wordt enigszins uitgesmeerd, ruizig en onscherp zijn. Daarnaast zijn de kransslagaders zeer klein (maar ongeveer 2 tot ongeveer 20 beeldpunten groot) en cardiale CTA-beelden bevatten soms opnamefouten, zoals verschoven beeldinformatie (door een onregelmatige hartslag) en strepen door het beeld (door b.v. pacemakers). Ook wordt er een hoge nauwkeurigheid verwacht van computermethodes. Als twee radiologen de kransslagaders intekenen variëren de intekeningen gemiddeld gezien 0,2 mm (ongeveer een half CTA-beeldpunt) en dat is ook de nauwkeurigheid die verwacht wordt van computermethoden.

Evaluatie van vaatasextractietechnieken

In de grove analyse stap van CTA-beelden worden lijnen gevonden door het midden van de bloedvaten in het beeld; dit wordt ook wel vaatasextractie genoemd. In 2006 waren er ongeveer 40 verschillende methoden voor het vinden van deze vaatassen in CTA-beelden beschreven in de literatuur en daarnaast bestonden er verschillende commerciële programma's voor deze taak. Sommige methoden werden als superieur aangeprezen en van andere had men het vermoeden dat ze niet goed werkten, maar niemand wist precies welke methode het beste was en hoe goed die methodes precies werken.

Daarom hebben wij een gestandaardiseerde evaluatie methode voor coronaire vaatasextractietechnieken geïntroduceerd. **Hoofdstuk 2** beschrijft deze ontwikke-

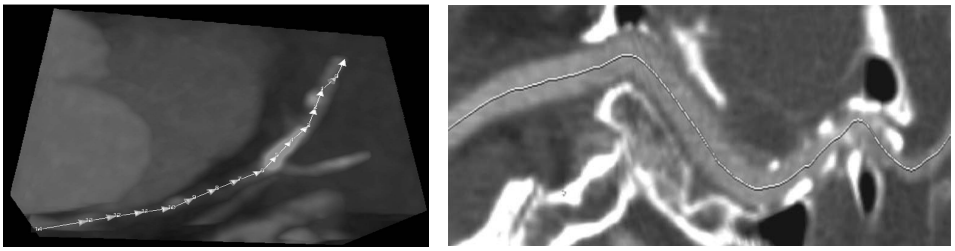

(a) (b)

Figuur 4: *Links: een visualisatie van een gevonden vaatas door een kransslagader met een complete verstopping. Rechts: een visualisatie van een gevonden vaatas door een halsslagader.*

ling. We hebben 32 CTA datasets van het hart (en dus de kransslagaders) verzameld en deze op een website (http://coronary.bigr.nl) openbaar beschikbaar gesteld. Ook hebben we representatieve evaluatiematen ontwikkeld. Ontwikkelaars van vaatasextractie methoden kunnen de beelden downloaden, hun computerprogramma op de beelden uitvoeren en de resultaten terugsturen naar de website. Deze extractieresultaten worden dan geanalyseerd en evaluatieparameters worden teruggestuurd naar de ontwikkelaar. De gestandaardiseerde resultaten kunnen geïnspecteerd worden op de website en ze kunnen opgenomen worden in een wetenschappelijke publicatie. Op deze manier kan objectief beoordeeld worden op welk aspect welke methode goed (of slecht) werkt.

Op dit moment is onze evaluatiemethode de standaard voor de evaluatie van nieuwe coronaire vaatasextractiemethoden. In nieuwe publicaties wordt nu altijd ons raamwerk gebruikt en er zijn al ongeveer 20 publicaties verschenen met resultaten van onze evaluatiemethode. Dit heeft ervoor gezorgd dat nu alle vaatasextractiemethoden goed met elkaar te vergelijken zijn.

Probabilistische vaatasextractietechniek

In **hoofdstuk 3** presenteren we een nieuwe generieke vaatasextractie methode. De voornaamste kracht van de methode is dat zij probabilistisch werkt. Dat betekent dat tijdens het vinden van de optimale vaatas met vele mogelijke vaatassen rekening gehouden wordt. Dit in tegenstelling tot andere methoden die maar met één vaatas of een beperkte set rekening houden. De nieuwe methode voorkomt op een efficiënte manier dat een lokale opnamefout, vaatverstopping of ruis het vinden van vaatassen verstoort. Daarnaast kan de methode zeer gemakkelijk aangepast worden voor andere bloedvaten en andere beeldvormende technieken. We hebben de methode met succes toegepast op synthetische computerbeelden, CTA-beelden van de halsslagaders en CTA-beelden van de kransslagaders. Figuur 4 toont hiervan twee voorbeelden.

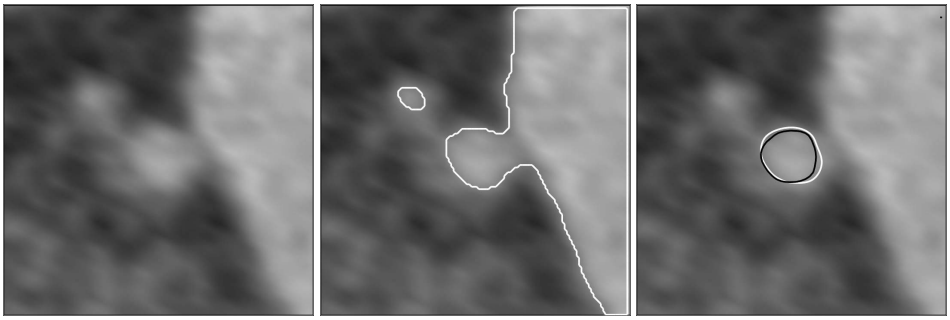

(a) (b) (c)

Figuur 5: *(a) Een gedeelte van een CT-scan met in het midden een kransslagader. (b) Het resultaat na de minimale snede stap. (c) In wit de resulterende vaatholtesegmentatie met daarbij in zwart de manuele intekening van een expert.*

Segmentatietechnieken voor de kransslagadervaatholte

Na de extractie van de vaatassen volgt de tweede stap: de gedetailleerde analyse. In dit proefschrift hebben we ons gericht op de kwantificatie van de vorm van de vaatholte (het gebied waar het bloed stroomt) van de kransslagaders. Deze kwantificatie begint met het uittekenen (segmentatie) van de vaatholte in het CTA-beeld. Met behulp van de vaatassen uit de eerste stap tekent een computerprogramma automatisch een oppervlak rond de vaatholte. Met deze uittekening kan het computerprogramma vervolgens de vorm van de vaatholte kwantificeren en bijvoorbeeld aangeven hoe dik het bloedvat is op een bepaald punt.

Hoofdstuk 4 presenteert een nieuwe techniek voor de segmentatie van de vaatholte van de kransslagaders. De techniek combineert twee relatief nieuwe technieken in de medische beeldverwerking; de zogenaamde minimale-snede graafanalyse ('graph-cut') techniek en robuuste kernelregressie. Met de inimale-snede graafanalyse techniek kan, op basis van bijvoorbeeld de sterkte van intensiteitsovergangen, een beeld zeer efficiënt in twee gedeeltes verdeeld worden; in ons geval de beeldpunten die horen bij de kransslagader en die van omliggend weefsel. Daarna volgt robuuste kernelregressie. Dit is een techniek uit de statistiek die gebruikt kan worden om mogelijke fouten van de minimale-snedesegmentatie te detecteren en te verwijderen. De combinatie van deze twee methoden resulteert in een efficiënte en nauwkeurige segmentatiemethode voor de kransslagadervaatholte. Figuur 5 toont een illustratie van het segmentatieproces.

In **Hoofdstuk 5** presenteren we een tweede aanpak voor de segmentatie van de kransslagadervaatholte. Deze techniek is, in tegenstelling tot de techniek gepresenteerd in hoofdstuk 4, niet regelgebaseerd, maar zelflerend. Met regelgebaseerd wordt bedoeld dat wij alle keuzes van de software expliciet geprogrammeerd heb-

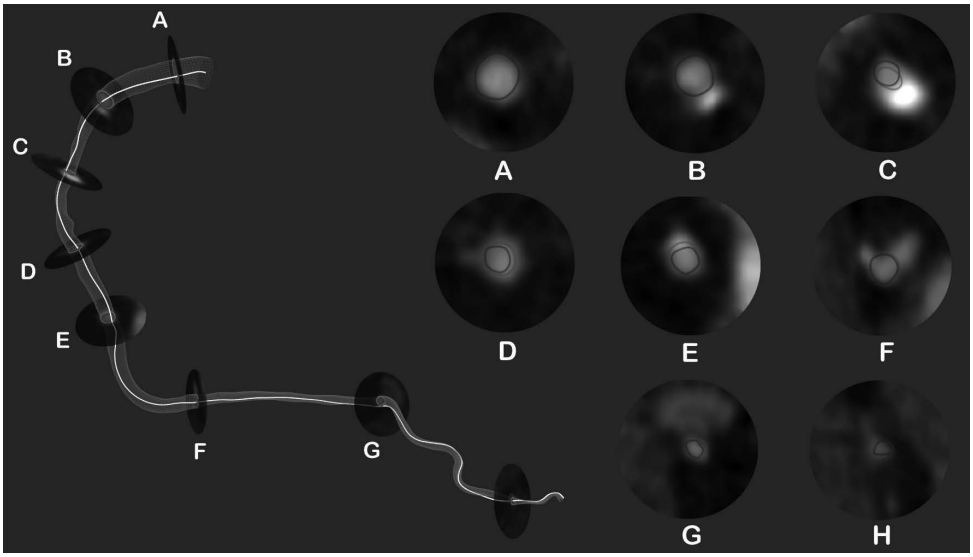

Figuur 6 (zie pagina 149 voor een kleurenversie): *Een voorbeeld van een 3D segmentatie van de coronaire vaatholte: een kransslagader en 8 geselecteerde dwarsdoorsnedes met daarop de automatische segmentatie in rood en een intekening van een expert in blauw.*

ben. We hebben de methode in hoofdstuk 4 gemaakt, getest op beelden en daarna aangepast tot zij optimaal werkte. De methode in **hoofdstuk 5**, daarentegen, leert van segmentaties gemaakt door medische experts. De computersoftware analyseert de beelden en de manuele intekeningen en leert op basis daarvan hoe een kransslagader er uit ziet in een CTA-beeld. Deze kennis gebruikt de methode volgens om de kransslagaders in een nieuw CTA-beeld te segmenteren. Een groot voordeel van deze methode is dat het, door haar zelflerende karakter, zeer makkelijk toepasbaar is op andere bloedvaten en beelden van andere beeldvormende technieken.

De methode is geëvalueerd door automatische intekeningen te vergelijken met intekeningen van experts. Daarnaast hebben we ook automatische kwantificaties van de (dwarsdoorsnede) oppervlakte van de vaatholte vergeleken met die van experts. Het bleek dat de methode vaak resultaten oplevert die nauwkeuriger zijn dan die van experts. In de toekomst kan deze methode daarom gebruikt worden tijdens de diagnose van coronaire atherosclerose om automatisch en nauwkeurig de mate van vaatvernauwing te kwantificeren. Figuur 6 toont een voorbeeld van een vaatholtesegmentatie met deze methode. De omslag van dit proefschrift toont dezelfde segmentatie met daarbij een aantal kwantificaties van de dwarsdoorsnede vaatholteoppervlaktes.

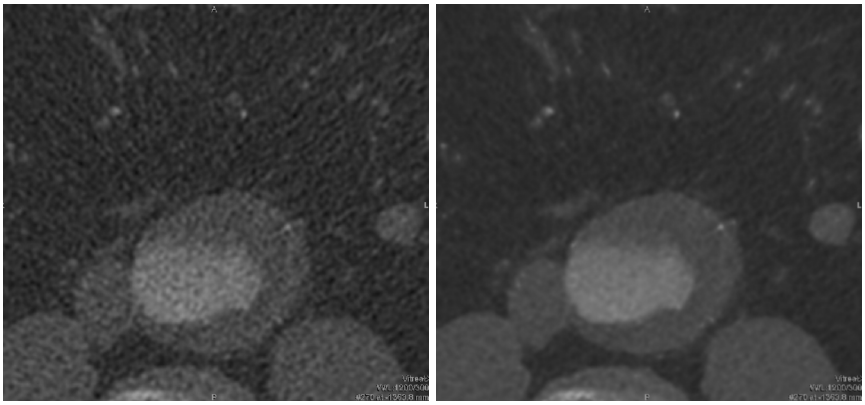

Figuur 7: *Links: Een gedeelte van een CTA-scan van een verwijde buikslagader. Rechts: Hetzelfde gedeelte van de CTA-scan na ruisreductie met de nieuwe methode. De ruis is onderdrukt en hierdoor zijn randen en kleine details beter zichtbaar.*

Geavanceerde visualisatie van CTA-beelden

Naast de kwantificatie van beeldparameters kan beeldverwerking ook gebruikt worden om de visualisatie van medische beelden te verbeteren. In **hoofdstuk 6** wordt een techniek gepresenteerd om de diagnostische kwaliteit van CTA-beelden te verbeteren door de ruis in deze beelden te onderdrukken. De methode verwijdert de ruis uit het beeld, terwijl het de diagnostisch belangrijke structuren zo goed mogelijk probeert te behouden. De methode is geëvalueerd op CTA-beelden van de buikslagader en het bleek beter te werken dan andere veelgebruikte ruisreductie methoden. Figuur 7 toont een voorbeeld van de resultaten.

Additioneel CTA-beeldverwerkingsonderzoek

Naast het werk wat ik hierboven heb beschreven heb ik ook meegewerkt aan ander beeldverwerkingsonderzoek met CTA. **Hoofdstuk 8** geeft een korte samenvatting van vier van deze CTA projecten: (i) De ontwikkeling en evaluatie van een specifieke methode om de coronaire vaatassen te vinden in een CTA-beeld van het hart, (ii) de ontwikkeling en evaluatie van een methode om automatisch alle hartkamers uit te tekenen in een CTA-beeld van het hart, (iii) onderzoek naar een nieuwe methode om CTA te maken met meerdere röntgenbronnen, en (iv) de ontwikkeling en evaluatie van een methode om de halsslagaders te segmenteren in CTA-beelden van de nek.

Dankwoord - Acknowledgement

I would like to thank all my colleagues, friends, and family for their assistance, support, or for just for being there, in happy and difficult times of my PhD period. Some people I would like to thank specifically for their contribution to this thesis.

Allereerst mijn promotor, professor Niessen. Wiro, enorm bedankt voor al je inzichten, je snelle feedback op mijn artikelen en je vertrouwen. Ik ben dankbaar voor de vrijheid die jij mij hebt gegeven en de mogelijkheden die je voor mij hebt gecreëerd. Het was mijn doel om breed geschoold dit promotieonderzoek af te ronden en dat is mede dankzij jou gelukt. Ik zie uit naar onze verdere samenwerking.

Mijn co-promotor, Theo van Walsum wil ik ook graag bedanken. Theo, je was een fijne begeleider en ik heb veel gehad aan jouw hulp, kritische vragen en snelle reacties. Je bent enorm hulpvaardig en je was er op de momenten dat ik een luisterend oor nodig had. Bedankt.

Graag wil ik ook mijn promotiecommissie bedanken; professor de Feijter, dr. Boudewijn Lelieveldt en dr. Bram van Ginneken, zeer bedankt voor het lezen van mijn manuscript. Bram, bij jou in de groep ben ik begonnen met medische beeldverwerking en ik vind het erg leuk dat je in mijn leescommissie hebt gezeten. Ik ben onder de indruk van jouw visie op het gebied van evaluatie en ik vond het erg leuk om jouw succesvolle evaluatie workshop een vervolg te geven. Verder wil ik professor Krestin en professor van der Steen vriendelijk bedanken voor het zitting nemen in mijn promotiecommissie. Dr. Angelini, Elsa, thank you for being part of my PhD committee. I'm very honored that you are willing to come to Rotterdam for my defense.

Coert, we zijn verschillend, maar juist dat heeft voor mij onze samenwerking bijzonder gemaakt. Ik ben onder de indruk van jouw doelgerichtheid. Onze samenwerking bij het organiseren van de 'grand challenge' workshop kon volgens mij niet beter en ook bij alle andere projecten hebben we erg goed samengewerkt. Zonder jouw inspanningen was mijn publicatielijst een stuk minder uitgebreid geweest. En bedankt voor de gezelligheid op onze kamer de afgelopen jaren!

Renske, jij hebt ervoor gezorgd dat ik als 'vatentrekker' ook nog een beetje breinonderzoek heb mogen doen. Ik heb genoten en veel geleerd van onze samenwerking. Ook wil ik je bedanken voor alle gezelligheid en vriendschap; we hebben veel plezier gehad en je was er voor me op de lastige momenten. Bedankt!

Ihor, thanks for the discussions on particle filtering and your help with the research on probabilistic vessel tracking. I'm impressed by your theoretical insights

and your knowledge of the literature. Also, I really enjoyed being your paranimf!

Marleen, bedankt voor de leuke samenwerking, het brainstormen en je feedback op ons artikel. Hopelijk werken we in de toekomst vaker samen.

Stefan, jouw theoretische kennis en kritische houding in combinatie met de pragmatische 'het werkt toch?' instelling is indrukwekkend. Bedankt voor de leuke samenwerking. Jammer dat het geplande artikel uiteindelijk net niet is gelukt. Ooit komt er nog een Klein & Schaap paper over een bizar en mooi theoretisch idee (en met een zinnige toepassing).

Lisan, Annick, Ermanno, Nico and Aad, bedankt dat ik jullie CTA-data mocht gebruiken voor mijn onderzoek en bedankt voor alle manuele annotaties!

Rashindra, ik kijk met veel plezier terug op onze samenwerking en ik heb genoten van onze discussies en gesprekken over beeldverwerkingsonderzoek, onderzoek in het algemeen, en compleet niet-werkgerelateerde dingen.

Karel, Arnold, Evert-Jan en Mathias, bedankt voor jullie bijdrage aan het ruisreductie werk. Het onderzoek met jullie was de basis van Hoofdstuk 6.

All the participants of the 'grand challenge' workshop, because of you the evaluation framework had such a great start. Thank you!

Marcel, ik ben blij dat jij als CT-specialist betrokken bent bij BIGR. Bedankt dat je mij vroeg om mee te werken aan het dual-energy project; ik vond het erg leuk om er aan bij te dragen. Ook kijk ik met plezier terug op ons reisje naar München.

Alina, Frank en Jolanda, bedankt voor de leuke samenwerking en voor de leerervaringen over stromingsleer en IVUS.

Hortense, Rahil, Hui, Nora, thanks for letting me be part of your research and thanks for making the 23rd floor such a nice place to work.

Fedde, ik kijk met veel plezier terug op onze discussies en hoop dat er nog vele mogen volgen. Ook vond ik het erg leuk om jouw paranimf te zijn. En bedankt voor het achterna brengen van mijn poster in Boston; zonder jou had ik misschien tijdens het IPMI-congres voor een leeg posterbord gestaan.

All other people from the Departments of Medical Informatics and Radiology, I really enjoyed all PhD drinks, dinners, cakes, daily lunches, and coffee breaks. Thanks!

Harry, heel erg bedankt voor al je correcties en je hulp bij het maken van de omslag van dit proefschrift!

Freek en Jord, super leuk dat jullie mijn paranimf zijn. Bedankt voor jullie adviezen over het werk in dit boekje, het promotietraject, en over niet-promotiegerelateerde zaken. Ik geniet van jullie, vaak verbazingwekkend scherpe, inzichten en vragen.

En als laatste, Anne en Leen, zonder jullie was ik er niet geweest, maar was ik vooral niet gekomen waar ik nu ben. Bedankt voor alles!

Publications

Journal Papers

- **M. Schaap**, A.M.R. Schilham, K. J. Zuiderveld, E. P. A. Vonken, M. Prokop and W.J. Niessen, Fast noise reduction in Computed Tomography for improved 3D visualization, *IEEE Transactions on Medical Imaging*, 2008

- **M. Schaap**, C.T. Metz, T. van Walsum, A.G. van der Giessen, A.C. Weustink, N.R. Mollet, C. Bauer, H. Bogunović, C. Castro, X. Deng, E. Dikici, T. O'Donnell, M. Frenay, O. Friman, M. Hernández Hoyos, P.H. Kitslaar, K. Krissian, C. Kühnel, M.A. Luengo-Oroz, M. Orkisz, Ö. Smedby, M. Styner, A. Szymczak, H. Tek, C. Wang, S.K. Warfield, S. Zambal, Y. Zhang, G.P. Krestin, W.J. Niessen, Standardized evaluation methodology and reference database for evaluating coronary artery centerline extraction algorithms, *Medical Image Analysis*, 2009

- C.T. Metz, **M. Schaap**, A.C. Weustink, N.R.A. Mollet, T. van Walsum and W.J. Niessen, Coronary centerline extraction from CT coronary angiography images using a minimum cost path approach, *Medical Physics*, 2009

- A.G. van der Giessen, **M. Schaap**, F.J.H. Gijsen, H.C. Groen, T. van Walsum, N.R.A. Mollet, J. Dijkstra, F. van de Vosse, W.J. Niessen, P. de Feyter, A.F.W. van der Steen and J.J. Wentzel, 3D fusion of Intravascular Ultrasound and Coronary Computed Tomography for in-vivo wall shear stress analysis: a feasibility Study, *International Journal of Cardiovascular Imaging*, 2009

- R. Manniesing, **M. Schaap**, S. Rozie, K. Hameeteman, D. Vukadinovic, A. van der Lugt and W.J. Niessen, Robust CTA lumen segmentation of the atherosclerotic carotid artery bifurcation in a large patient population, *Medical Image Analysis*, 2010

- H.A. Kirisli, **M. Schaap**, S. Klein, S.L. Papadopoulou, M. Bonardi, C.H. Chen, A.C. Weustink, N.R.A. Mollet, E. P. A. Vonken, R.J. van der Geest, T. van Walsum and W.J. Niessen, Evaluation of a multi-atlas based method for segmentation of cardiac CTA data: a large-scale, multi-center and multi-vendor study, *Medical Physics*, 2010

- C.T. Metz, S. Klein, **M. Schaap**, T. van Walsum and W.J. Niessen, Nonrigid registration of dynamic medical imaging data using nD+t B-splines and a groupwise optimization approach, *Medical Image Analysis*, 2010

- **M. Schaap**, T. van Walsum, M. de Bruijne, L. Neefjes, C. Metz, E. Capuano, and W.J. Niessen, Robust shape regression for supervised vessel segmentation and its application to coronary segmentation in CTA, *Submitted*, 2010

- R. de Boer, **M. Schaap**, F. van der Lijn, H.A. Vrooman, M. de Groot, A. van der Lugt, M.A. Ikram, M.W. Vernooij, M.M. Breteler, and W.J. Niessen, Statistical analysis of minimum cost path based structural brain connectivity, *Submitted*, 2010

- M. van Straten, **M. Schaap**, M.L. Dijkshoorn, M.J. Greuter, A. van der Lugt, G.P. Krestin, and W.J. Niessen, Automated bone removal in CT Angiography: Additional value of tin filtration in dual-energy based methods and comparison to single energy methods, *Submitted*, 2010

- A.A. Isola, C.T. Metz, **M. Schaap**, S. Klein, M. Grass, and W.J. Niessen, Cardiac motion-corrected iterative cone-beam CT reconstruction using a semi-automatic minimum cost path-based coronary centerline extraction, *Submitted*, 2010

- K. Hameeteman, M.A. Zuluaga, M. Freiman, L. Joskowicz, O. Cuisenaire, L. Florez Valencia, M.A. Gulsun, K. Krissian, J. Millei, W.C.K. Wong, M. Orkisz, H. Tek, M. Hernandez Hoyos, F. Benmansour, A.C.S. Chung, S. Rozie, M. van Gils, L. van den Borne, J. Sosna, P. Berman, N. Cohen, P.C. Douek, I. Sanchez, M. Aissat, **M. Schaap**, C.T. Metz, A. van der Lugt, G. Krestin, W.J. Niessen, T. van Walsum, Evaluation framework for carotid bifurcation lumen segmentation and stenosis grading, *Submitted*, 2010

Papers in Conference Proceedings

- **M. Schaap**, I. Smal, C.T. Metz, T. van Walsum and W.J. Niessen, Bayesian tracking of elongated structures in 3D images, *Information Processing in Medical Imaging*, 2007

- **M. Schaap**, R. Manniesing, I. Smal, T. van Walsum, A. van der Lugt and W.J. Niessen, Bayesian tracking of tubular structures and its application to carotid arteries in CTA, *Medical Image Computing and Computer-Assisted Intervention*, 2007

- **M. Schaap**, L. Neefjes, C.T. Metz, A.G. van der Giessen, A.C. Weustink, N.R.A. Mollet, J.J. Wentzel, T. van Walsum and W.J. Niessen, Coronary lumen segmentation using Graph Cuts and Robust Kernel Regression, *Information Processing in Medical Imaging*, 2009

- **M. Schaap**, C.T. Metz, T. van Walsum, A.G. van der Giessen, A.C. Weustink, N.R.A. Mollet, G. P. Krestin and W.J. Niessen, On the evaluation of coronary artery centerline extraction algorithms - A standardized methodology and reference database, *Invited paper. World Congress on Medical Physics and Biomedical Engineering*, 2009

- C.T. Metz*, **M. Schaap***, A.G. van der Giessen, T. van Walsum and W.J. Niessen, Semi-automatic coronary artery centerline extraction in computed tomography angiography data, *IEEE International Symposium on Biomedical Imaging: Macro to Nano*, 2007 *: Joint first authorship

- T. van Walsum, **M. Schaap**, C.T. Metz, A.G. van der Giessen and W.J. Niessen, Averaging centerlines: Mean Shift on paths, *Medical Image Computing and Computer-Assisted Intervention*, 2008

- C.T. Metz, **M. Schaap**, T. van Walsum and W.J. Niessen, Two point minimum cost path approach for CTA coronary centerline extraction, *3D Segmentation in the Clinic: a Grand Challenge II. MICCAI 2008 workshop*, 2008

- C.T. Metz, **M. Schaap**, S. Klein, A.C. Weustink, N.R.A. Mollet, C. Schultz, R.J. van Geuns, P.W. Serruys, T. van Walsum and W.J. Niessen, GPU accelerated alignment of 3-D CTA with 2-D X-ray data for improved guidance in coronary interventions, *IEEE International Symposium on Biomedical Imaging*, 2009

- C.T. Metz, **M. Schaap**, S. Klein, L. Neefjes, E. Capuano, C. Schultz, R.J. van Geuns, P.W. Serruys, T. van Walsum and W.J. Niessen, Patient specific 4D coronary models from ECG-gated CTA data for intra-operative dynamic alignment of CTA with X-ray images, *Medical Image Computing and Computer-Assisted Intervention*, 2009

- H.A. Kirisli, **M. Schaap**, S. Klein, L. Neefjes, A.C. Weustink, T. van Walsum and W.J. Niessen, Fully automatic cardiac segmentation from 3D CTA data: a multi-atlas based approach, *SPIE Medical Imaging*, 2010

- R. Shahzad, **M. Schaap**, T. van Walsum, S. Klein, A.C. Weustink, L. J. van Vliet and W.J. Niessen, A patient-specific coronary density estimate, *IEEE International Symposium on Biomedical Imaging: from Nano to Macro*, 2010

- R. de Boer, **M. Schaap**, F. van der Lijn, H. Vrooman, M. de Groot, M. W. Vernooij, M. A. Ikram, E.F.S. van Velsen, A. van der Lugt, M. M. B. Breteler and

W.J. Niessen, Statistical analysis of structural brain connectivity, *Medical Image Computing and Computer-Assisted Intervention*, 2010

- C.T. Metz, N. Baka, H.A. Kirisli, **M. Schaap**, T. van Walsum, S. Klein, L. Neefjes, N.R.A. Mollet, B.P.F. Lelieveldt, M. de Bruijne and W.J. Niessen, Conditional shape models for cardiac motion estimation, *Medical Image Computing and Computer-Assisted Intervention*, 2010

- H. Tang, R. S. van Onkelen, T. van Walsum, K. Hameeteman, **M. Schaap**, F.L. Tori, Q.J.A. van den Bouwhuijsen, J.C.M. Witteman, A. van der Lugt, L. J. van Vliet and W.J. Niessen, A semi-automatic method for segmentation of the carotid bifurcation and bifurcation angle quantification in Black Blood MRA, *Medical Image Computing and Computer-Assisted Intervention*, 2010

- H. Tang, T. van Walsum, R. S. van Onkelen, S. Klein, K. Hameeteman, **M. Schaap**, Q.J.A. van den Bouwhuijsen, J.C.M. Witteman, A. van der Lugt and L. J. van Vliet, Multispectral MRI centerline tracking in carotid arteries, *SPIE Medical Imaging*, 2011

Workshop Proceedings

- C.T. Metz*, **M. Schaap***, T. van Walsum, A.G. van der Giessen, A.C. Weustink, N.R.A. Mollet, G. P. Krestin and W.J. Niessen, Editorial: 3D Segmentation in the Clinic: A Grand Challenge II - Coronary Artery Tracking. *Medical Image Computing and Computer-Assisted Intervention*, 2008 *: Joint first authorship

Patents

- **M. Schaap** and K. J. Zuiderveld, Image enhancement using anisotropic noise filtering, filed 2005, issued 2010

- C.T. Metz, **M. Schaap**, T. van Walsum and W.J. Niessen, Image processing method and system, filed 2009

Conference Abstracts

- C.T. Metz, **M. Schaap**, N.R.A. Mollet, T. van Walsum and W.J. Niessen, Semi-automatic multi planar reformatting of the coronary arteries from computed tomography angiography data, *European Congress of Radiology*, 2008

- R. Manniesing, D. Vukadinovic, S. Rozie, **M. Schaap**, A. van der Lugt and W.J. Niessen, Automated CTA lumen segmentation of the atherosclerotic carotid

artery bifurcation in a large patient population, *Radiological Society of North America, Annual Meeting*, 2009

- C.T. Metz, **M. Schaap**, S. Klein, E. Capuano, L. Neefjes, T. van Walsum and W.J. Niessen, Extracting coronary artery motion from ECG-pulsed and gated CTA data, *Radiological Society of North America, Annual Meeting*, 2009
- M. van Straten, **M. Schaap**, M. Dijkshoorn, M. Greuter, A. van der Lugt and W.J. Niessen, Improved Dual-Energy-based bone removal technique for Dual-Source CT, *Radiological Society of North America, Annual Meeting*, 2009
- M. van Straten, **M. Schaap**, M. Dijkshoorn, A. Moelker, G. P. Krestin and W. J. Niessen,, The trouble with calcium blooming at low voltage CT angiography, *World Congress on Medical Physics and Biomedical Engineering*, 2009
- A.A. Isola, C.T. Metz, **M. Schaap**, S. Klein, W.J. Niessen and M. Grass, Coronary segmentation based motion corrected cardiac CT reconstruction, *IEEE Nuclear Science Symposium and Medical Imaging Conference*, 2010

Phd portfolio

Courses	**Year**
Cardiovascular Image Analysis, COEUR, Rotterdam	2006
Heart Failure, COEUR, Rotterdam	2006
Introduction to Data Analysis, NIHES, Rotterdam	2006
Principles in Research in Medicine, NIHES, Rotterdam	2009
INFOBIOMED Training Challenge, Universitat Pompeu Fabra, Barcelona, Spain	2006
MevisLab introduction course, MeVis, Bremen, Germany	2006
Angina Pectoris, Erasmus MC, Rotterdam	2007
Front-end Vision and Multiscale Image Analysis, ASCI, University of Eindhoven, Eindhoven	2006
Presentations at international conferences	
International Conference on Information Processing in Medical Imaging, Kerkrade (Oral presentation)	2007
International Conference on Medical Image Computing and Computer-Assisted Intervention, Brisbane, Australia (Workshop, oral presentation)	2007
IEEE International Symposium on Biomedical Imaging, Boston, US (Poster presentation)	2007
International Conference on Medical Image Computing and Computer-Assisted Intervention, New York, US (Poster presentation)	2007
International Conference on Information Processing in Medical Imaging, Williamsburg, US (Poster presentation)	2009
World Congress on Medical Physics and Biomedical Engineering, München, Germany (Oral presentation)	2009
Other presentations	
Medical Informatics PhD days, Breukelen	2009
ASCI conference, Lommel, Belgium	2006
Research Seminars: Biomedical Imaging Group Rotterdam (4x), Departments of Radiology (2x) and Medical Informatics (3x)	2006-2010

Attending international conferences

International Conference on Information Processing in Medical Imaging, Kerkrade	2007
International Conference on Medical Image Computing and Computer-Assisted Intervention, Brisbane, US	2007
IEEE International Symposium on Biomedical Imaging, Washington, US	2007
International Conference on Medical Image Computing and Computer-Assisted Intervention, New York, US	2007
International Conference on Information Processing in Medical Imaging, Williamsburg, US	2009
World Congress on Medical Physics and Biomedical Engineering, München, Germany	2009
IEEE International Symposium on Biomedical Imaging, Boston, US	2009
IEEE International Symposium on Biomedical Imaging, Rotterdam	2010

Attending international workshops

Feature Extraction and Classification, Washington, US	2007
3D Segmentation in the Clinic: A Grand Challenge, Brisbane, Australia	2007
Statistical Registration: Pair-wise and Group-wise Alignment and Atlas Formation, Brisbane, Australia	2007
Manifold Learning in Medical Imaging, New York, US	2008
Statistical Shape Analysis: Theory, Software, and Applications, Boston, US	2009
Machine Learning for Biomedical Image Analysis, Rotterdam	2010

Teaching

Teaching Assistant, Introduction to Image Processing for Medical Students	2006-2009

Other

Organizing the International workshop '3D Segmentation in the Clinic: a grand challenge II', New York, US	2008
Reviewing for IEEE Transactions on Medical Imaging, Medical Physics, and international conferences	2006-2010

About the author

Michiel Schaap was born on April 30, 1979 in Hilversum, the Netherlands. After finishing his HAVO education at the A. Roland Holst College in Hilversum, he started with the study of Computer Engineering at the Hogeschool Utrecht. During this study he worked as a freelance computer programmer for a variety of companies. He obtained his Bachelor degree in Computer Engineering with a graduation project in the field of medical image processing (thorax CT-scan registration) at the Image Sciences Institute at the Academic Medical Center Utrecht. After his Bachelor graduation, Michiel started with the Master's program Biomedical Image Sciences at Utrecht University. During this study he worked as a scientific programmer for the Radiology department of the Academic Medical Center Utrecht. His graduation project involved the development of a noise reduction algorithm for medical images. Part of his graduation project was carried out at Vital Images in Minnetonka, VS. He graduated in 2005 (cum laude).

From 2006 to 2010 he worked as a PhD student at the Biomedical Imaging Group Rotterdam (BIGR), part of the Departments of Medical Informatics and Radiology. The result of this work is described in this thesis. Since September 2010 Michiel works at the Biomedical Imaging Group Rotterdam on a project to develop and evaluate a fully-automatic coronary stenoses detection and quantification system based on CT angiography.