Segmentation and motion estimation of stent grafts in abdominal aortic aneurysms

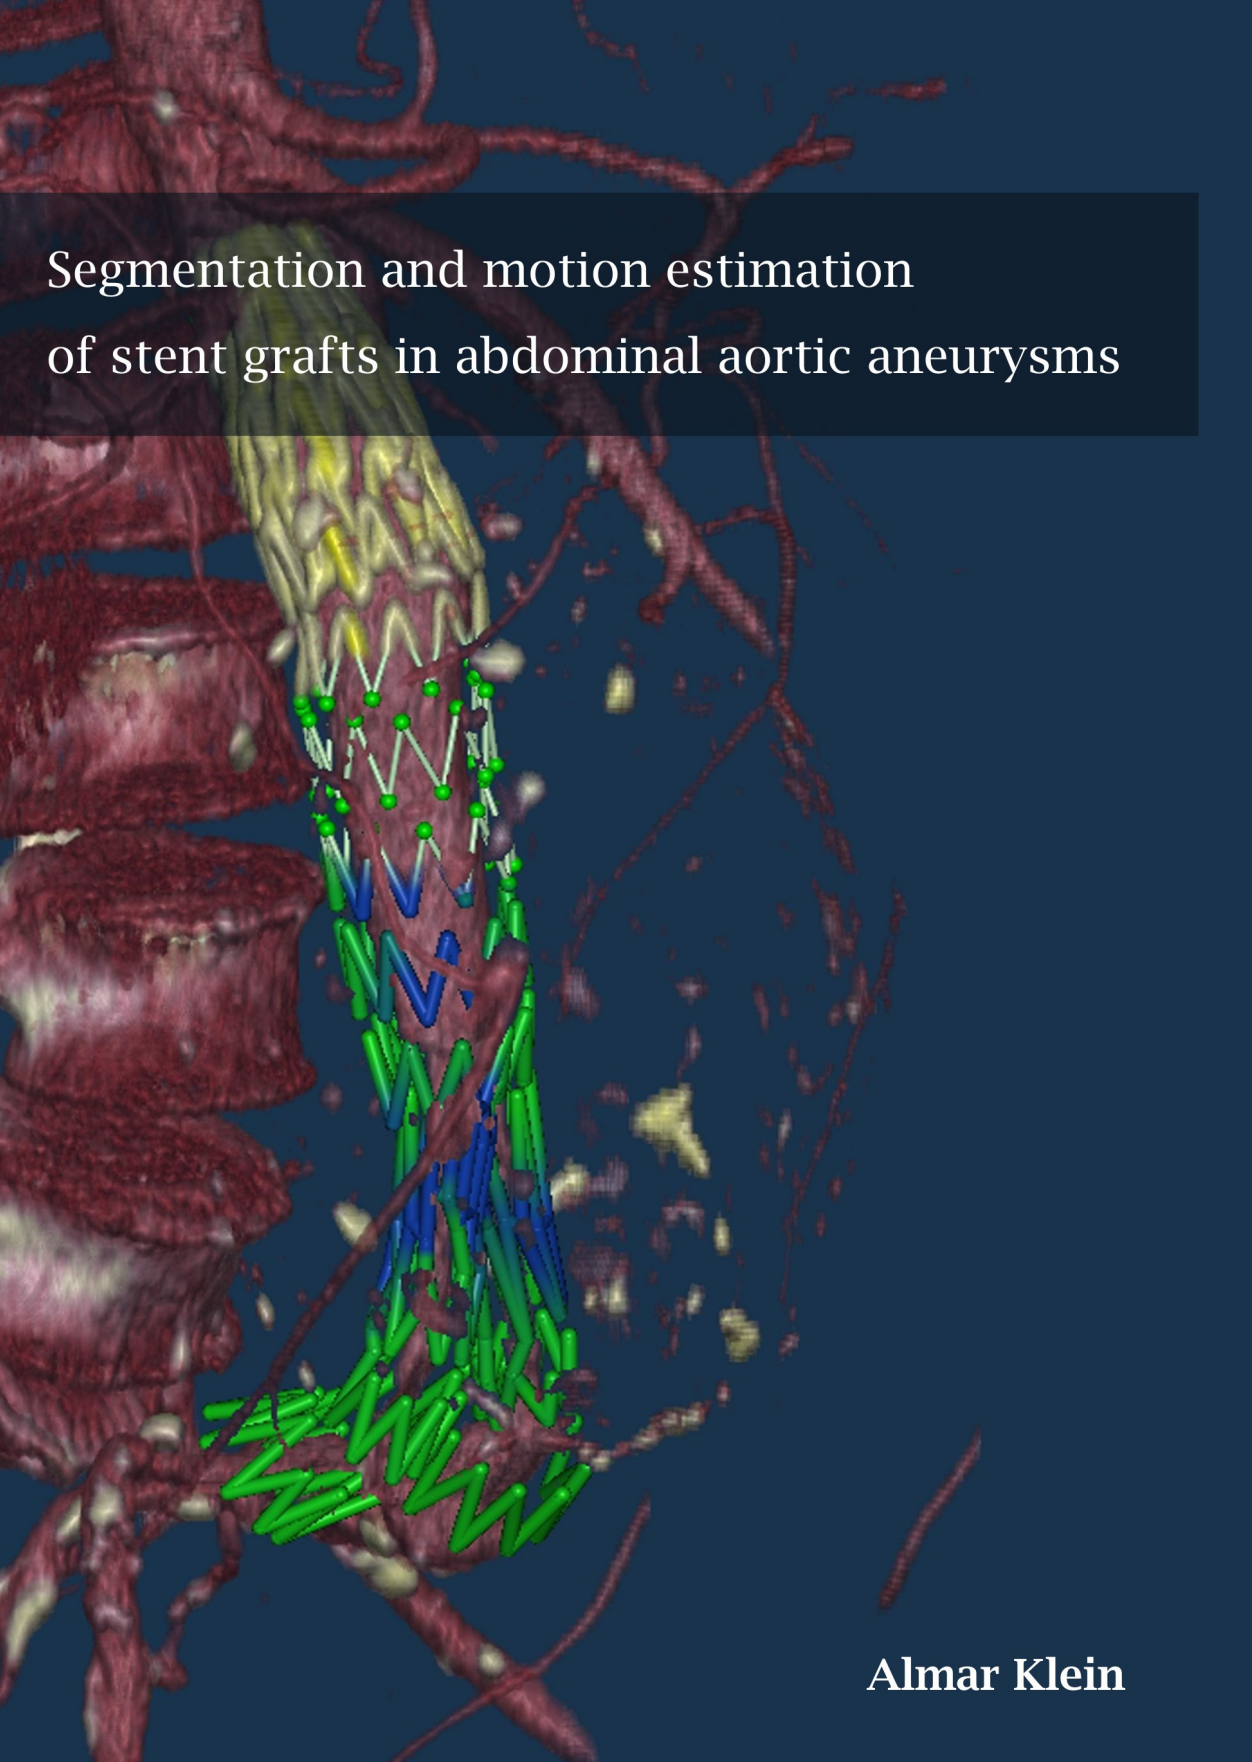

Almar Klein

Segmentation and motion estimation of stent grafts in abdominal aortic aneurysms

Almar Klein

De promotiecommissie:

voorzitter en secretaris:	
Prof.dr.ir. A.J. Mouthaan	Universiteit Twente
promotoren:	
Prof.dr.ir. C.H. Slump	Universiteit Twente
Prof.dr. L.J. Schultze Kool	Radboud University Nijmegen Medical Centre
assistent promotor:	
Dr.ir. W.K.J. Renema	Radboud University Nijmegen Medical Centre
leden:	
Prof.dr.ir. P.H. Veltink	Universiteit Twente
Prof.dr. D.B.F. Saris	Universiteit Twente
Prof.dr. M. Prokop	Radboud University Nijmegen Medical Centre
Dr. Y. Hoogeveen	Radboud University Nijmegen Medical Centre

Signals & Systems group,
EWI Faculty, University of Twente
P.O. Box 217, 7500 AE Enschede, the Netherlands

Print: Gildeprint B.V.
Typesetting: Lyx/LaTeX2e

ISBN 978-90-365-3266-2

SEGMENTATION AND MOTION ESTIMATION OF STENT GRAFTS IN ABDOMINAL AORTIC ANEURYSMS

PROEFSCHRIFT

ter verkrijging van
de graad van doctor aan de Universiteit Twente,
op gezag van de rector magnificus,
Prof.dr. H. Brinksma,
volgens besluit van het College voor Promoties
in het openbaar te verdedigen
op 22 november 2011 om 14:45h

door

Almar Klein
geboren op 4 mei 1983
te Apeldoorn

Dit proefschrift is goedgekeurd door:

De promotoren:	Prof.dr.ir. C.H. Slump
	Prof.dr. L.J. Schultze Kool
De assistent promotor:	Dr.ir. W.K.J. Renema

Contents

1 Introduction

1.1 Motivation

Endovascular aneurysm repair (EVAR) is a technique which uses stent grafts to treat aortic aneurysms in patients at risk of aneurysm rupture. Although this technique has been shown to be successful on the short term, the long term results are less optimistic due to failure of the stent graft. The pulsating blood flow applies stresses and strains to the stent graft, which can cause problems such as breakage, leakage, and migration.

Therefore it is important to gain more insight into the in vivo motion behavior of these devices. If we know more about the motion patterns in well-behaved stent grafts as well as devices with problems, we expect to be better able to distinguish between these type of behaviors. We hope that these insights will enable us to detect stent-related problems and might even be used to predict problems beforehand. Ultimately, these insights may help in designing the next generation stent grafts.

Several patients were asked to participate in this study by having an ECG-gated CT scan instead of a regular one. For patient safety, the dose was kept similar to a normal CT scan.

1.2 Purpose

The purpose of this thesis is to present a method that enables quantitative measurements on the motions of stent grafts from ECG-gated CT data. The proposed method consists of two parts: segmentation and registration. In the segmentation part the stent is detected from the data and a geometric model is produced that describes the stent in a concise way. In the registration part the deformation field of the data is calculated and used to incorporate motion in the geometric model. We distinguish the following research questions:

- Is the data obtained with ECG-gated CT suitable for measuring the motions expected for stent grafts in AAA?
- Can ECG-gated CT replace the regular CT study that patients currently have, without any negative effects on the clinical procedure?
- Can we segment the stent graft from these (noisy) data, and with what accuracy?
- Is it possible to measure the motion of the stent graft from these data, and with what accuracy?
- What kind of motions do stent grafts make inside the human body?

1.3 Outline of this thesis

In the next chapter, we will first give an overview of the work described in this thesis. That chapter also discusses the clinical background in more detail. The remainder of the thesis is divided in three parts:

Part A: ECG-gated CT — In Chapter 3 we discuss the experiments that we have performed to study the possibilities and limitations of ECG-gated CT. In Chapter 4 an experiment is described to compare the quality of the data obtained with ECG-gated CT and regular CT.

Part B: Segmentation — In Chapter 5 we discuss the first part of an approach to segment the stent graft in 2D slices sampled orthogonal to the centerline. In Chapter 6 the tracking part of this approach is described. The fundamental limitations of this approach are discussed and the accuracy of this method is determined in an experiment. In Chapter 7 three variants of an algorithm to track the wires of the stent in 3D are compared. This can be seen as the predecessor of the segmentation algorithm described in Chapter 8, of which we demonstrate that the produced geometric model has a high correspondence with expert annotations.

Part C: Registration — In Chapter 9 we present a groupwise registration algorithm which produces diffeomorphic deformation fields (i.e. no folding). We demonstrate its applicability to different kind of data and its robustness for intensity differences between the images. In Chapter 10 different registration methods are evaluated for data containing a stent graft. The best algorithm is applied to the patient data and the resulting deformation fields are applied to the geometric models to obtain a dynamic model of the stent.

2 Overview

This chapter is based on the book chapter entitled "Motion calculations on stent grafts in AAA" published in the open access book "Diagnosis, Screening and Treatment of Abdominal, Thoracoabdominal and Thoracic Aortic Aneurysms"[64].

2.1 Introduction

2.1.1 Abdominal Aortic Aneurysm

An abdominal aortic aneurysm (AAA) is a dilation of the vessel wall of the aorta, usually between the renal arteries and the illiac arteries (Figure 2.1b). When such an aneurysm ruptures, death follows within several minutes [114]. It is a leading cause of death that affects mainly older white men (between 65 and 74 years of age). Due to the aging population, the incidence and prevalence of AAA is expected to rise [114]. Symptoms can consist of pain in the abdomen, back or groin, although AAA is asymptomatic for most patients. The cause of AAA is multifactorial (for instance cigarette smoking, genetic influence, atherosclerosis) and is related to weakening of the aortic wall.

When detected in time, the aneurysm can be repaired. The norm for repair has been that the aneurysm diameter should be larger than 5.5 cm, although it has been shown that this is not an optimal criterion [111, 108].

2.1.2 Aneurysm repair

The conventional technique to repair AAA (since the 1950s) is to replace the unhealthy aorta with an artificial graft by open surgery (Figure 2.1c). Although this approach has good long-term results, the intervention is associated with high risks and has a 5% mortality rate [12, 53].

Endovascular aortic replacement (EVAR) is a minimal invasive technique (approved by the U.S. Food and Drug Administration in 1999) that uses stent grafts to

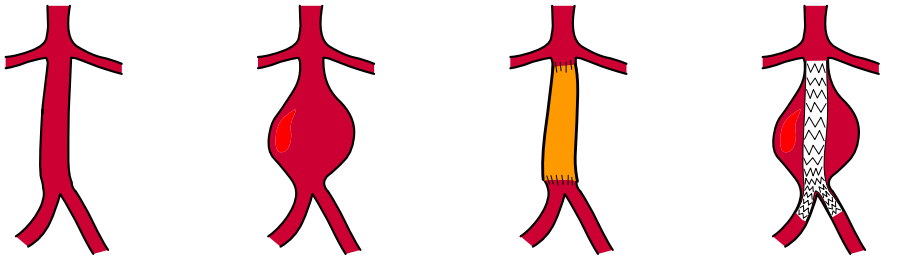

(a) Healthy aorta (b) Aneurysm (c) Open repair (d) EVAR

Figure 2.1: Illustration of AAA and two repair techniques: open repair (c), and EVAR (d).

treat AAA. A stent graft (also known as an *endograft*) is a metallic frame surrounded by a cloth graft. The aneurysm is accessed via the lumen of the blood vessels in the groin. After deployment of the stent graft the blood will flow through the stent graft, and the aneurysm outside the graft is excluded from the circulation. Hereby the force exerted by the pulsating blood on the aneurysm is reduced (Figure 2.1d). This will significantly reduce the chance of aneurysm rupture, and causes the aneurysm to shrink in size over time [114].

With a mortality rate of 2% the technique has been proven to be successful [125, 12]. However, due to the need for reintervention EVAR does not have a significant advantage over open repair on the long term [53, 27]. Late stent graft failure is therefore a serious complication in endovascular repair of aortic aneurysms [16, 28, 54, 80, 88, 100]. Examples are metal fatigue, stent graft migration [80, 70], and the formation of endoleaks (blood flow into the aneurysm sac) which can result in aneurysm expansion and rupture [81, 82, 106].

2.1.3 Motion of stent grafts

The long-term durability of stent grafts is affected by the stresses and hemodynamic forces applied to them, which may be reflected by the movements of the stent graft itself during the cardiac cycle. Studying the dynamic behavior of stent grafts can therefore give a better understanding of their motion characteristics, and can give insights into how these motion characteristics relate to certain stent-related problems. This information will be beneficial for designing future devices and can be valuable in predicting stent graft failure in individual patients [74].

Motions of (stent grafts in) AAA can be measured using fluoroscopic roentgenographic stereophotogrammetric analysis (FRSA) [71], dynamic magnetic resonance imaging [52], and ECG-gated CT [108]. Although ultrasound is also used [115], it does not produce the three dimensional images that are required for the quantitative analysis of the whole stent. ECG-gated CT has the advantage of having high contrast for metal objects. Furthermore, ECG-gated CT is widely available, easily accessible, and can easily be applied in a post-operative setting.

To study the motions quantitatively, and to process the large datasets associated with ECG gating, automated processing is required. We divide the processing in two

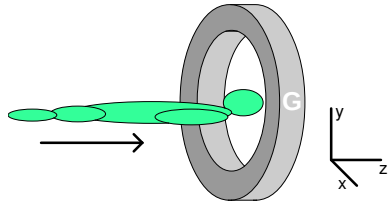

Figure 2.2: Illustration of the orientation of the patient with respect to the CT scanner. The ring indicated by 'G' represents the gantry of the CT scanner.

steps: segmentation of the stent, and calculating the motions of the stent[1].

2.2 ECG-gated CT

In computed tomography (CT) a three-dimensional image of an object is constructed by a computer from a series of images obtained using roentgen radiation (Figure 2.2, Figure 2.3). In current CT scanners the x-ray source rotates around the object while the object is moved through the scanner in the z-direction. This enables scanning the complete object in one continuous (helical) motion [58].

In recent years there have been major advancements in CT. Shorter rotation times and the development of multi detector CT (MDCT) enabled the technique of ECG gating, often referred to as cardiac CT [45]. With this technique, the patient's ECG signal is measured during the scan. It is then possible to divide the raw scan data into bins that correspond to consecutive phases of the heart beat. The data in each bin is then reconstructed into a three-dimensional image (i.e. a volume), and the final result is a sequence of volumes, each corresponding to a different phase of the heart cycle (Figure 2.4). This allows 4D visualization of the scanned object and enables investigation to its temporal behavior [45, 94]. ECG-gated CT enabled measuring motions that are synchronous with the patient's heart beat; other motions, such as those caused by breathing result in motion artifacts. The number of volumes that is reconstructed per scan is in the order of 8-20 [108, 49].

2.2.1 Dose and the noisy nature of CT data

One of the major downsides of CT in general is the exposure of the patient to ionizing radiation, which can have negative effects on the long term health of the patient [98, 39]. The dose should therefore be kept as low as reasonably achievable. However, this results in higher noise levels and more image artifacts, which can cause problems for automatic image analysis algorithms that often need high quality data to operate. Algorithms that can perform their task on low dose data can therefore contribute to better patient safety.

In ECG-gated CT, multiple volumes are produced from the same amount of raw data. Assuming that the dose is kept the same, the amount of noise in each volume is therefore significantly larger than in volumes reconstructed using conventional CT.

[1] A stent graft consists of a metal frame surrounded by blood-proof material (the graft). When we only use the word "stent", we refer to the metallic frame: the graft is not visible on a CT scan.

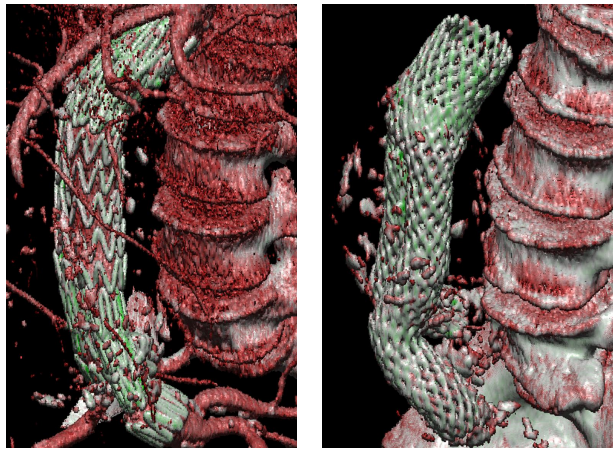

Figure 2.3: Illustration of iso surfaces rendered from CT data of two types of stent grafts.

2.2.2 Combining the volumes

The clinic sometimes also requires the result of a non-gated scan because of its lower levels of noise. Unfortunately, not all scanners are capable of producing a non-gated three-dimensional image in case ECG-gated scanning was used. Scanning patients twice is not an option considering the extra dose this would imply.

Averaging the data of the volumes off-line (i.e. not on the scanner's reconstruction computer) also produces a 3D dataset. This is a straightforward process, yet fundamentally different from combining the raw data (sinogram) before the filtered back-projection reconstruction (as happens for a non-gated scan). Due to non-linearities in the reconstruction process of the scanner, the results may be similar but will never be exactly the same.

In a study on phantom data acquired with a 64-slice Siemens Somatom CT scanner it was found that averaging the volumes in this way does not have negative effects on image quality in terms of noise, frequency response and motion artifacts [63]. Rather, the noise was found to be slightly lower, and motion artifacts were found to be less severe.

For the purpose of segmentation, combining the volumes can also be advantageous. It has been shown that combining a subset of all volumes in the sequence can produce better results due to a more optimal compromise between noise and motion blur [66].

2.2.3 The effect of the patient's heart rate

While the patient is moved through the scanner (i.e. along the z-axis), data is collected and the patient's ECG-signal is measured (Figure 2.2). To construct a single volume with full coverage in the z-direction, data is collected from multiple heart beats (Figure 2.4). The table speed, rotation time of the scanner, and the heart beat of the patient together determine the amount of overlap in the z-direction. Negative overlap signifies a volume gap (Figure 2.4b), which is expressed as extremely noisy bands (Figure 2.5) that propagate through the data (the exact effects can differ per

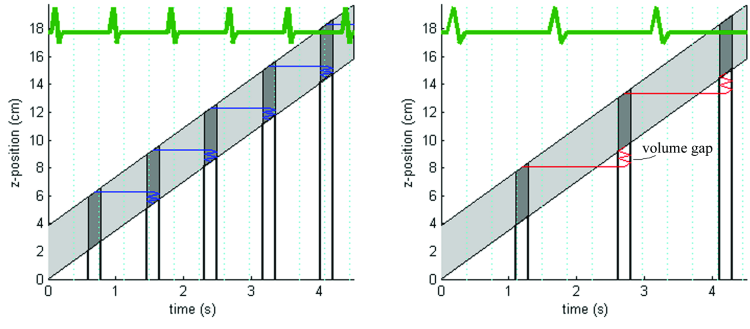

(a) 70 bpm (b) 40 bpm

Figure 2.4: Diagram illustrating the process of ECG-gating. The light grey band indicates the covered z-positions of the detector during the scan. The dark grey patches represent parts of the phase in each heart beat.

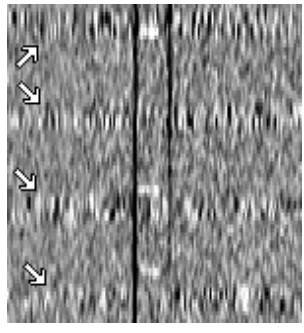

Figure 2.5: Illustration of the noise bands in the CT images caused by the volume gaps due to a too low heart rate (45 bpm) during scanning. Shown is an image of a phantom which has small metal bars embedded at regular intervals. It can be seen how the second bar from the top is hidden by a noise band (i.e. volume gap).

scanner). The data inside these gaps is completely unreliable (even if the scanner tries to interpolate it) because data at these positions is simply not available [62].

It can be shown theoretically, and it has been verified in an experiment [62], that there is a minimum required heart beat in order to obtain images without volume gaps. This minimum heart beat can be calculated as follows:

$$B_{min} = \frac{60 \cdot p}{T_{rot}}, \tag{2.1}$$

where p is pitch factor, T_{rot} the rotation time, and B_{min} the minimum required heart rate in beats per minute. For a typical setup with $T_{rot} = 0.37$ and $p = 0.34$ the minimum required heart beat $B_{min} = 55$ bpm.

It is noteworthy that a too high heart rate should also be avoided, since this leads to increased motion artifacts.

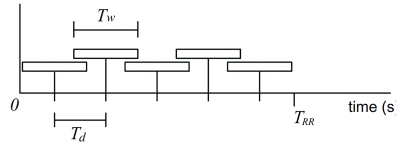

Figure 2.6: Diagram illustrating the two aspects of temporal resolution. T_w is determined by the rotation time and reconstruction algorithm. T_d is determined by the heart rate and the number of reconstructed phases (five in this example).

2.2.4 Temporal resolution

The temporal resolution of the technique of ECG gating consists of two parts (Figure 2.6): the first is the width of each phase T_w, which is fully determined by the rotation time and reconstruction algorithm. Its value determines to what extent motion causes artifacts in the resulting data. Since T_w depends on the applied reconstruction algorithm, which is often chosen by the manufacturer, this value is often unknown. In [62] a simple experiment is described to measure the value of T_w empirically.

The second part is the (temporal) distance between phases T_d, which is determined by the number of phases and the heart rate. It represents the sampling rate of the technique. If more phases are reconstructed, T_d decreases and the overlap between phases increases.

In an ideal scenario, T_w should be as low as possible to be least affected by motion artifacts, and T_d should be approximately equal such that the sampling frequency is high enough to prevent aliasing, with a minimal number of phases.

2.2.5 Application

ECG-gated CT is extensively used in cardiac exams [120, 85, 2], especially for the assessment of coronary arteries [32, 18, 46]. The goal in most of these studies is to limit the effect of motion rather than to examine the motion itself for which the technique can also be utilized.

Recently, ECG-gated CT was used to study the pulsating motion of AAA [108], and the motion of the renal arteries [90].

The abdominal aorta is constantly in motion caused by the pressure waves from the contracting heart. However, the dynamics of this motion are more subtle than the motions present in the heart itself. It has been shown that the order of magnitude of these motions is in the order of 2 mm [91, 108]. It is reported that the limits of the motion that can be detected in clinical practice by ECG gating are slightly less than the spacing between the voxels (usually in the order of 0.5 mm), and that for a typical setup frequency components up to 2.7 Hz can be accurately detected [62]. This makes ECG-gated CT a suitable technique for studying motions in AAA.

2.3 Segmentation of the stent graft

Segmentation of the stent graft is performed on a three-dimensional image. Depending on how the data is processed further, the segmentation is applied to all volumes in the sequence, or to a single volume obtained by combining the volumes in the sequence.

Several studies have been published on the segmentation of blood vessels in 3D, which have correspondences with the wires of the frame of the stent and may therefore be of interest (see [79] and [59] for an overview of vessel segmentation methods). Methods that fit a series of spheres or ellipsoids to the vessel [8, 126], and methods that segment the contour in slices perpendicular to the vessel centerline [50, 76] assume a solid vessel with a diameter of several voxels. Due to the small diameter of stent wires (1 to 3 voxels) and their sharp corners, these methods are not suitable for the segmentation of stents. Region growing methods [13] have problems with leaks and gaps and need a second stage to find the geometry from the segmented voxels. A common method is the two-step approach [44, 57, 84, 123], which first segments the vessel using a vessel measure [43] followed by centerline tracking. This method, however, is known to have difficulties were the structure is not tubular, such as in crossings and sharp corners in the stent's frame.

A related method is used by [75] for segmentation of stent grafts in the aortic arch. Interest points are extracted that are located on the center line of the stent determined by a skeletonization of the volume thresholded at 2000 Hounsfield units (HU) and weighted by its vesselness measure [43]. The result is a dense set of points that lie on the frame of the stent.

Unfortunately, the *quality of the data*—defined as how well the frame of the stent is distinguishable in the data—is not always sufficient for such a method to fully segment the stent's frame [65]. This quality depends on the combination of used dose, stent wire diameter, material properties of the stent (i.e. absorption coefficient), and patient anatomy. The stent can consist of CT values as low as 300 HU [65]. There are also reports of some stent types being barely distinguishable, whereas other stent types are well visible in data obtained using the same scanner settings [66].

In addition to the bad visibility of the frame of the stent, several problems can be identified for (low dose) CT data. Firstly, the data is relatively noisy. Secondly, streak artifacts occur where the stent's metal frame is thick or where a coil is present next to the stent graft. Thirdly, contrast agent injected in the blood results in CT-values close to the range of CT-values seen for most stents. Fourthly, due to image artifacts, the wire of the stent sometimes contains gaps.

In this section, we discuss a way to model the stent graft, and two approaches to obtain such a model from the volumetric CT data in ways that are relatively robust for the aforementioned problems.

2.3.1 Modeling the stent

Most studies related to the motion of stent grafts focus on measuring the stent's diameter changes [51] or determining the motion at a sparse set of points on the stent [74]. A model that enables capturing material properties and high level knowledge regarding the stent graft characteristics would be valuable to gain more insight in the stent's in vivo behavior [74]. Furthermore, such a model can also help in performing more reliable (fluid dynamics) simulations, which is important for improving current stent designs [17, 68].

In [66] a geometric model is proposed that represents the wire frame of the stent as an undirected graph, with nodes placed at the corners and crossings of the frame, and the edges between the nodes representing the wires (Figure 2.7). This model can be applied to different stent types, and represents the topology of the stent's frame

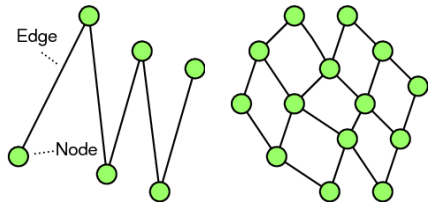

Figure 2.7: Example graphs that describe a geometric model of a stent frame. The edges between the nodes represent the physical wire frame of the stent. Nodes are placed at corners (a) and crossings (b).

in a concise and natural way.

2.3.2 Segmentation of the stent graft via centerline tracking

A stent has a tubular structure, sometimes with branches, and can be approximated by a series of stacked contours which are orthogonal to the centerline (Figure 2.8). An approach published by [65] is to segment the stent in 2D images sampled perpendicular to its centerline. Regions with high CT-values (typically above 500 HU) exist where the metallic frame of the stent penetrates the image. These regions have high CT-values and—due to their "pointy" structure—well suited for point detection.

The approach to segment the stent in these 2D images is to first detect a set of interest points, after which a clustering algorithm is applied to find the points that are on the wire of the stent. This process is then repeated in an iterative fashion, while tracking along the centerline of the stent. At the end of this process, a 3D geometric model of the stent is obtained.

An advantage of this approach is that part of the algorithm is 2D, which makes visualization and algorithm design easier. A disadvantage is that modeling the stent as a series of stacked contours causes difficulties at bifurcations and when parts of the frame of the stent overlap.

2.3.2.1 Point detection

Four different point detection algorithms were taken into consideration and tested in an experiment. An algorithm based on the product of Eigenvalues was found to work the best. This measure, also known as the Gaussian curvature, can be expressed using image derivatives: $\frac{\partial^2 L}{\partial x^2} \cdot \frac{\partial^2 L}{\partial y^2}$, where L is the 2D image. Other methods taken into consideration were a static threshold, a dynamic threshold, and the Laplacian (the sum of Eigenvalues).

2.3.2.2 Clustering

Four different clustering algorithms were taken into consideration and tested in an experiment. The best method was found to be a custom method that uses a virtual stick to select the stent-points in an iterative fashion. By selecting points from the inside of the contour, spurious points outside of the contour are ignored (Figure 2.9).

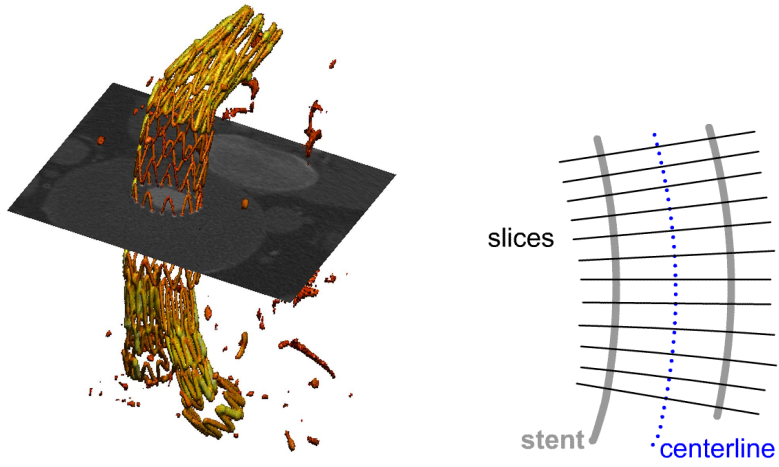

Figure 2.8: Illustration of how the stent can be approximated as a series of stacked contours. (a) shows a volume rendering of the stent and a slice orthogonal to its centerline. (b) shows a schematic illustration of contours perpendicular to the stent's centerline.

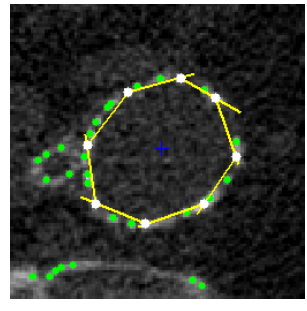

Figure 2.9: Example of the clustering algorithm after finding a coarse contour.

Other clustering methods taken into consideration were circle fitting, ellipse fitting, and GVF snakes [124].

The result of the clustering method is a set of points that represent the contour of the stent. By fitting a circle on these points, an estimate of the radius and center position can be obtained, which are used during centerline tracking.

2.3.2.3 Centerline tracking and modeling

Starting from a manually selected seed point, the algorithm tracks the stent in both directions. Starting from a coarse estimate of the centerline orientation, slices are sampled, to which the aforementioned algorithms are applied. The center position found at each slice is used to estimate the next centerline position. For the method to be less sensitive to noise present on the center estimate (caused by the discrete nature of the contour points), a smoothness constraint is adopted. Bifurcations are detected when a significant change in the diameter estimate is encountered. Subsequently,

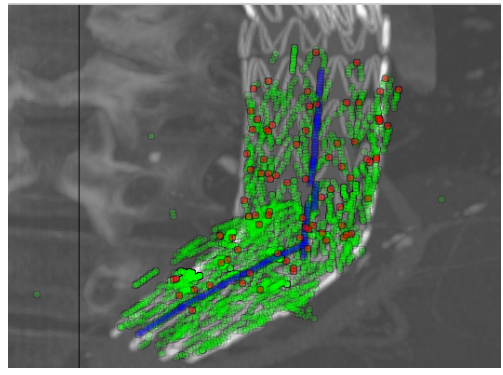

Figure 2.10: Example of the algorithm in progress. The blue dots indicates the found centerline. The green dots indicate the found stent points, and the larger red dots indicate the found nodes which will be connected to form a geometrical model.

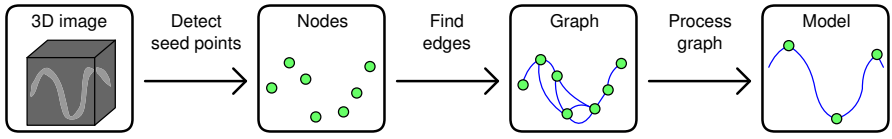

Figure 2.11: Flowchart illustrating the three processing steps to extract a geometrical model from the CT data.

both branches of the bifurcations are tracked individually.

To deal with the gaps between the different parts of the stent graft that are present in some stent types (Figure 2.8a), the tracking will proceed in the last known direction if no contour could be found. When no contour is found along a predefined distance, it is assumed that the end of the stent is reached, and the tracking stops.

During centerline tracking the contour points in the current slice are matched to the contour points of the previous slice. In this fashion the individual wires are tracked too. The positions where two wires meet—which represent the corners and crossings of the stent's frame—are detected, and nodes are created at these positions to build the geometric model (Figure 2.10).

2.3.3 Segmentation of the stent graft via the minimum cost path method

A method to segment the stent graft in 3D by finding the optimal paths between a large set of automatically detected seed points is proposed in [66]. The method can be divided into three steps, which are illustrated in the flow chart in Figure 2.11.

2.3.3.1 Detection of seed points

In the first step, a set of seed points is found by searching the volume for voxels subject to three criteria: 1) The voxel intensity must be a local maximum. 2) The voxel intensity must be higher than a predefined threshold value. 3) The voxel must have a direct neighbor with an intensity also above this threshold value.

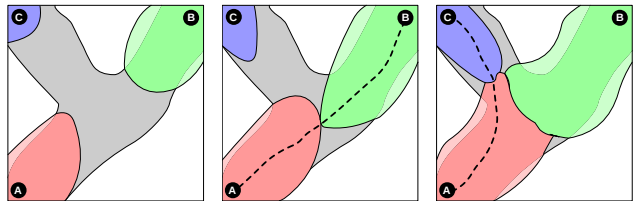

(a) (b) (c)

Figure 2.12: Illustration of three meeting fronts in the MCP algorithm. The black circular shapes indicate seed points A, B and C. In (a) the fronts do not yet meet. In (b) front A meets front B, and the path is traced. A few iterations later, in (d) a third front meets with the first, connection seed points A and C.

2.3.3.2 Finding the optimal paths

In the second step, the seed points are connected using a modified version of the minimum cost path (MCP) method. The MCP method can be used for segmentation of vessels and other structures (e.g. [24, 30, 38, 47, 55, 99, 121]). It is a level set method in which a front is propagated monotonically following a (non-negative) cost function. The advantages of this method are that it can be implemented in a computationally efficient way, and that it can easily be modified to make it more suitable for a specific problem, see for example [60] and [24].

To use the MCP method for stent segmentation, it is modified such that the fronts evolve from all the seed points found in the seed point detection step. Connections between the nodes are detected when two fronts collide, and the paths between the points are found using a backtrace map that is maintained during the evolution of the front.

The result of the MCP algorithm is a graph consisting of nodes (the seed points) connected by edges. Each edge is associated with a path of voxels connecting one node to another. However, many of these edges are false edges and have to be removed.

2.3.3.3 Graph processing

In the third step, the false edges are removed using graph processing techniques. For this purpose, two scalar values are associated with each edge. The first is α, the maximum cumulative cost on the path. It represents the *weakness* (i.e. inverse strength) of the edge. This value is used to establish the order of the edges; a stronger connection (lower α) is preferred over a weaker one. The second scalar value is β, the minimum intensity (the CT-value in Hounsfield Units) on the path. Due to the definition of CT-values (-1000 representing air and 0 representing water) this value has a physical meaning and represents the *quality* of the edge; it is used to determine whether an edge should be removed or not.

The processing of the graph occurs in multiple different passes. Firstly, weak edges are removed based on the expected number of edges for each node. This value depends on the specific stent type being segmented. The weakness value α is used to establish the weakest edges to consider for removal, and the quality measure β is used to determine whether an edge should be removed. Secondly, a clean-up pass is

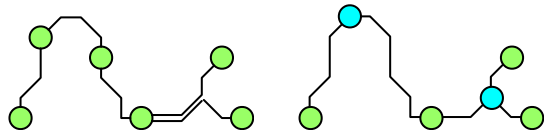

(a) Before (b) After

Figure 2.13: Illustration of adding corners and nodes. In (b) two nodes were removed and a node was placed at the corner. Another node was inserted at the crossing.

performed to remove redundant edges; an edge is found redundant if there is a path of one or two stronger (i.e. lower α) edges that connect the same nodes. Thirdly, corners are detected in the wire, and nodes are placed at the positions that have the highest curvature. Hereafter the graph is cleaned up again. Fourthly, crossings are detected and nodes are added to represent them. Finally, after a final clean-up step, all the paths are smoothed.

2.3.3.4 Experiments and results

To evaluate the quality of the geometric model produced by this method, experiments were performed in which the model was compared with a reference model annotated by three experts. By counting the number of corresponding edges, a similarity measure was obtained. A training set was used to obtain the optimal parameter values of the algorithm, and using a test set the final performance of the method was evaluated.

The algorithm was found to be robust for variations in its parameter values, and for the high noise levels present in the data. The found similarity with the reference data was found to be 96% and 92% for the two stent types considered in the experiments. Visual inspections of the results showed that most errors were present in difficult areas of the stent, such as bifurcations and narrow legs where the wire has relatively sharp corners.

An example of the results after each processing step is shown in Figure 2.14. In Figure 2.15 lit surface renders are shown for the found geometric models of three datasets.

2.4 Calculating motions and forces of the stent graft

When the geometric model of the stent is obtained, it can be used as a tool to study the motions of the stent graft. For this purpose, motion is applied to the model. In the first part of this section we discuss how this can be done, following the ideas presented in [66]. In the second part of this section we discuss an alternative method that uses active shape models to study stent graft motions.

2.4.1 Motion analysis using a geometric model

The motion of interest is obtained from the sequence of CT volumes using a registration algorithm. The purpose of a registration algorithm is to (elastically) align two images; the result is a deformation field that describes how one image should be

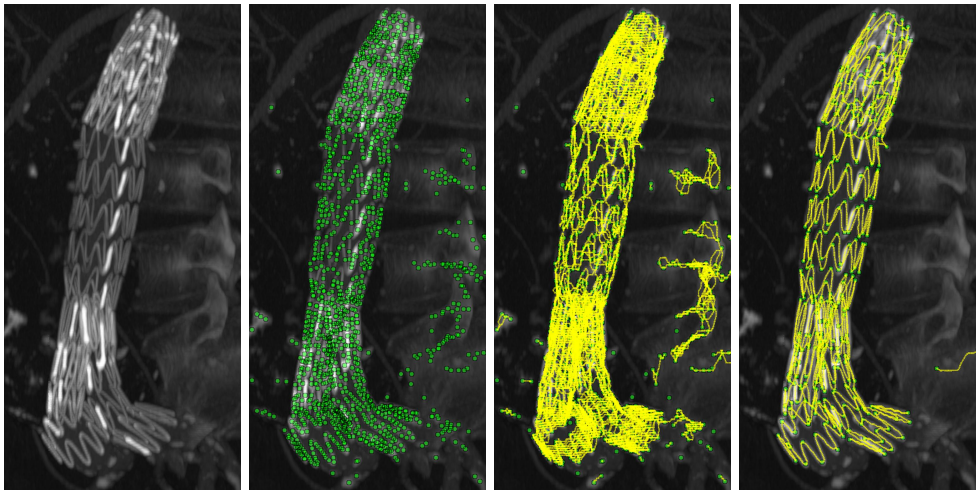

(a) MIP of the CT data (b) 1732 seed points (c) 4963 initial edges (d) 531 final edges

Figure 2.14: Illustration of the different algorithm steps. Shown are a Maximum Intensity Projection (MIP) of the original data (a), the detected seed points (b), the found edges (c), the result after processing the graph (d).

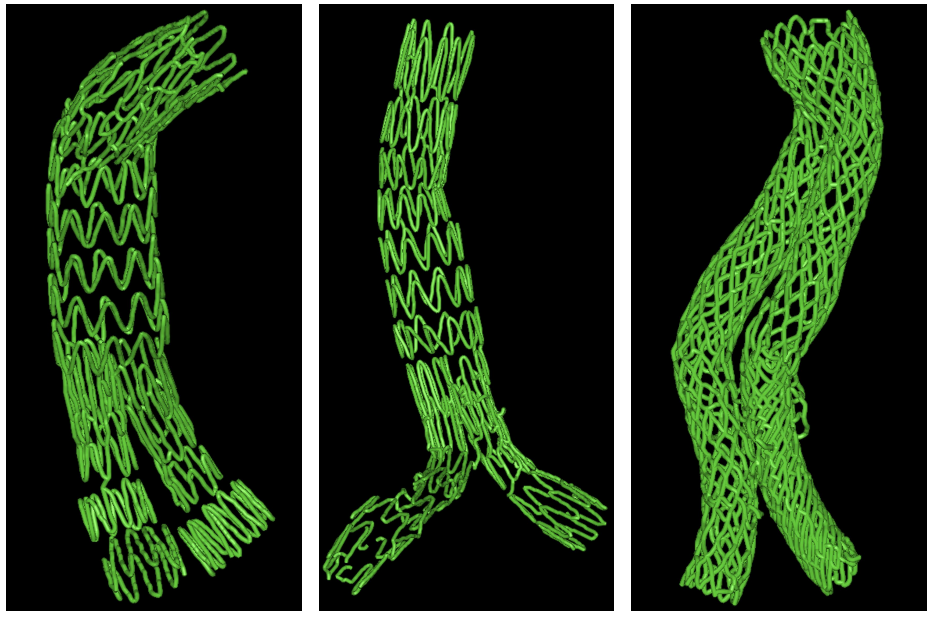

Figure 2.15: Illustration of lit surface renders of the geometric models for three example stent grafts.

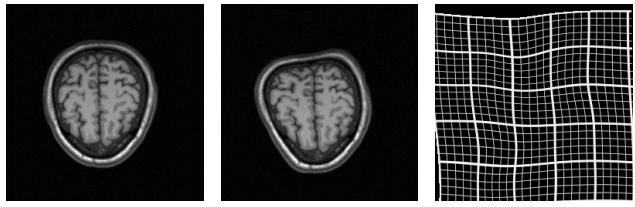

(a) Image A (b) Image B (c) Deformation field

Figure 2.16: An example of registering two images, and the resulting deformation field.

deformed to align it with the second (Figure 2.16). This deformation field can be applied to the geometric model to enable studying the motions.

2.4.1.1 Image registration

The current range of common region based image registration algorithms can be divided into two classes. Both classes usually adopt a multiscale approach in order to prevent finding a local minimum, and to speed up the registration process. The first class employs a B-spline grid to describe the deformation field, which is optimized by minimizing/maximizing a similarity measure. Using Mutual Information (MI) as a similarity measure, these methods have been shown to be robust for differences in the appearance between the images [101, 118, 83]. While the use of a B-spline grid can cause problems when describing rotational deformations [72], it has the advantage that the deformations are described in an efficient way and are physically realistic [93]. Additionally, the deformations can be regularized in various ways, for example by minimizing bending energies or penalizing small Jacobians [101].

The second class uses image forces calculated at the pixels/voxels to drive the registration process. A popular example is the Demons algorithm [109], which is related to optical flow. The deformation is obtained for each pixel individually by calculating image forces, and regularization of the deformation field is performed by Gaussian diffusion. The Demons algorithm is capable of handling extreme deformations, which can also be a downside, since such deformations are usually not physically realistic. Another problem with the Demons algorithm is that it assumes pixel intensities in corresponding regions between images to be similar, which causes problems in images containing much noise or artifacts such as bias fields [93].

It is of importance to select the right registration algorithm for each problem, and to choose the best parameters, of which most registration algorithms have many [67]. In the case of registering the different volumes obtained by ECG-gated CT, the used registration algorithm should be accurate in order to deal with the small motions present in AAA, and should be robust for noise and other artifacts associated with low dose CT. Which algorithms qualify for this task is currently being investigated.

2.4.1.2 Analysis of motion and forces

The result of the registration algorithm is a deformation field that describes the deformation for each voxel in the volume. To study the motions of a stent, the

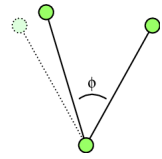

Figure 2.17: Illustration of how the (change of) angle ϕ can be used to a estimate the force present in the node when motion is applied to the model.

deformation field is applied to the nodes of the geometric model: for each node the deformation is applied that corresponds to its location in the volume. This will allow quantitative studies to the motion patterns of individual stents, and allows comparison between patients.

Because the topology of the stent is fully captured by the geometric model, the forces acting on the stent's frame can be estimated by incorporating material properties such as stiffness, and by calculating the change of the angle between two edges (Figure 2.17).

2.4.2 Motion analysis using active shape models

A technique of interest for the evaluation of motions of stent grafts is that of [75]. Their application is for stent grafts in the aortic arch to treat aortic ruptures caused by trauma. In [74] a method for the unsupervised learning of models from sets of interest points was proposed. It is based on minimum description length (MDL) group-wise registration [110]. The global and local deformation are captured using a statistical deformation model that is built during registration of a sparse set of interest points. No a priori annotation, or definition of topological properties of the structure is necessary.

Instead of deforming the whole volume they search for correspondences between finite lists of interest points and local features in the data. This has a few advantages: 1) The algorithm can omit variations that are not relevant to the model. 2) The approach is not constrained to an a priori topological class because it does not rely on a mapping to a reference manifold. 3) No prior segmentation of the object is necessary, only the interest point extraction method has to be chosen according to the structure of interest. A disadvantage of this approach, however, is that it is less accurate than texture-based registration with which registration errors smaller than the voxel size can be obtained [92, 61]. This can be a problem for stent grafts in AAA, because—due to the larger distance from the heart—the motions are smaller than in the aortic arch.

2.4.2.1 MDL registration

First a set of interest points are detected in each volume of the ECG-gated sequence. These points are treated as landmarks candidates; each landmark is associated with a position (x, y, z) and local features (such as image intensity and steerable filters). The registration is initialized by pairwise matching of a subset of the interest points.

Figure 2.18: Global deformation (color-coded) for a few stent grafts. (Image courtesy of G. Langs from [75].)

Starting from these correspondences group-wise registration is performed by minimizing a criterion function that captures the compactness of the model comprising the variation of landmark positions and local feature variation at the landmark positions. The minimum description length criterion accounts for the fact that the landmarks located on the stent move in a highly correlated manner during the cardiac cycle.

The registration is optimized by a combination of k-D trees and genetic-optimization, and is followed by a refinement using a direct search. The optimization process results in a shape variation model, which is then used to study the motions of the stent.

2.4.2.2 Motion analysis

The analysis of the stent deformation during the cardiac cycle is performed using the shape model that results from the group wise registration (Figure 2.18). For each landmark the positions in all volumes in the sequence are known. Three measurements can be obtained for each landmark: 1) The modes of variation of the statistical shape model, which capture the correlation between landmark movements. 2) The displacement of the landmarks, which reflect the absolute movement in the anatomical environment. 3) The compactness of the local shape model build with the closest landmarks, which gives an indication about the complexity of the local deformation. This last measure is particularly of interest, since it is well suited to show regions of potential stress to the material.

2.5 Outlook

An automated method to quantitatively study the motions and forces of stent grafts in vivo enables studying the motion patterns of individual patients, relate them to data of a previous date, or relate them to the motion patterns of other patients.

It would also be interesting to study the range of motion patterns of stent grafts in patients without problems, and compare them to the motions in patients who do have problems. Such studies would, however, require large datasets to incorporate all the variabilities in motion patterns, particularly because problems with stent grafts are relatively rare. Nevertheless, we believe that such studies can help our understanding of the dynamics and failure of stent grafts, and can thereby help in designing better stent grafts in the future. Further, we hope that we are able to correlate certain

distinct motion patterns to specific stent-related problems, so that this technique can be used for diagnostic purposes and prediction of stent failure.

2.6 Conclusion

Using ECG-gated CTA, information about the motion of stent grafts in AAA can be obtained. Using segmentation methods, a geometric model of the stent can be obtained that describes the topology of the stent in a compact way. Using registration techniques, the deformation field can be found, which can then be applied to the found geometric model. Thereby the motions of the stent graft are known in great detail, and enables calculating the forces acting on the stent. Both parameters (motion and force) provide new information that can be used in further analysis of in vivo stent graft behavior and future device design.

Part A

ECG-Gated CT

gating signal: *A digital signal or pulse that provides a time window so that a particular event or signal from among many will be selected and others will be eliminated or discarded.*

3

Detectability of motions in AAA with ECG-gated CTA: a quantitative study

This chapter is published in slightly modified form in Medical Physics [62].

Abstract

ECG-Gated CT is a technique that can be used for evaluating the motions of stent grafts inside abdominal aortic aneurysms (AAA). To be able to reliably quantify the motion, however, it is of importance to know the performance and limitations of ECG-gating, especially when the motions are small, as is the case in AAA. Since the details of the reconstruction algorithms are proprietary information of the CT manufacturers and not in the public domain, empirical experiments are required. The goal of this chapter is to investigate to what extent the motions in AAA can be measured using ECG-gated CT.

The duration of each ECG-gated phase was found to be 185 ms, which corresponds to half the rotation time and is thus in accordance with half scan reconstruction applied by the scanner. By using subpixel localization, motions become detectable from amplitudes as small as 0.4 mm in the x direction and 0.7 mm in the z direction. With the rotation time used in this study, motions up to 2.7 Hz can be reliably detected. The reconstruction algorithm fills volume gaps with noisy data using interpolation, but objects within these gaps remain hidden.

This chapter gives insight into the possibilities and limitations for measuring small motions using ECG-gated CT. From the results we conclude that ECG-gated CTA is a suitable technique for studying the expected motions of the stent graft and vessel wall in AAA.

3.1 Introduction

In recent years there have been major advancements in computed tomography (CT). Shorter rotation times and the development of multi detector CT (MDCT) enabled the technique of ECG gating [45]. ECG gating uses the ECG signal of the patient to divide the raw scan data into bins that correspond to consecutive phases of the heart beat. The data is reconstructed into a number of volumes, each corresponding to a different phase of the heart cycle. This allows 4D visualization of the scanned object and enables the investigation of its temporal behavior [45, 94].

ECG gating is extensively used in cardiac exams [2, 120, 85], especially for the assessment of coronary arteries [32, 18, 46]. The goal in most of these studies is to limit the effect of motion rather than to examine the motion itself for which the technique can also be utilized. Recently, ECG-gated CT Angiography (CTA) was used to study the pulsating motion of abdominal aortic aneurysms (AAA), [91, 108, 119] and the motion of the renal arteries [90]. Finally, ECG gating can also be used to evaluate the motion of implanted abdominal stent grafts [65].

Late stent graft failure is a serious complication in endovascular repair of aortic aneurysms [54, 100, 16, 28]. Better understanding of the motion characteristics of stent grafts will be beneficial for designing future devices. In addition, these data can be valuable in predicting stent graft failure in patients. If detected, these patients will benefit from early reintervention.

The abdominal aorta is constantly in motion caused by the pressure waves from the contracting heart. However, the dynamics of this motion are more subtle than the motions present in the heart itself. To be able to reliably quantify these motions, it is of importance to know the capabilities and limitations of the applied ECG gating technique, especially when the motions of interest are small as in the case of AAA (in the order of 2 mm [91, 108]). Several studies have been performed to validate the use of ECG gating for diagnostic purposes [2, 120, 32, 18, 36]. Quantitative performance [41] and simulation [86] studies have also been performed. However, to the best of our knowledge, there are no quantitative studies on the performance of ECG-gated CTA with respect to the detectability of motions in AAA. Such a study is required to be able to distinguish measured motion from measurement errors in studies on motion using ECG-gated CTA, and will help in designing future experiments to study motions of stent grafts in AAA.

The purpose of this paper is to investigate the performance of ECG gating quantitatively in motion detection for AAA. The results are compared to values which are theoretically determined on the basis of the scan parameters. This provides insight into the effects of the complex reconstruction algorithms — and the applied proprietary optimizations and corrections applied by manufacturers — on the motion detectability. Furthermore, the limits of the motion that can be detected in clinical practice by ECG gating is determined. In the present study the research questions can be divided into four topics, which are discussed in the next sections.

3.1.1 Temporal resolution

In ECG-gated CT temporal resolution consists of two parts (Figure 3.1): the first is the width of each phase T_w, which is fully determined by the rotation time and reconstruction algorithm. Its value determines to what extent motion causes artifacts in the resulting data. The second is the (temporal) distance between phases T_d,

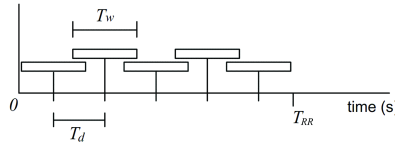

Figure 3.1: Diagram illustrating the two aspects of temporal resolution. T_w is determined by the rotation time and reconstruction algorithm. T_d is determined by the heart rate T_{RR} and the number of phases that we chose to reconstruct (five in this example).

which is determined by the number of phases and the heart rate. If more phases are reconstructed, T_d decreases and the overlap between phases increases.

Since using redundant data degrades the temporal resolution [96], optimal temporal resolution in terms of T_w is achieved by minimizing the number of projections used to reconstruct the image. There are a variety of reconstruction algorithms, which result in different values for T_w. For standard fan beam reconstruction, for example, the minimum range of projections is 180 degrees plus the fan beam [96]. For parallel beam reconstruction, however, temporal resolution of half the rotation time can be achieved [96, 94, 45, 86, 1]. Multi segment reconstruction can result in an even higher temporal resolution for some heart rates by reconstructing a volume using the raw data from different heart cycles [46, 1, 32, 37]. To be able to use an N-segment reconstruction, the spiral pitch factor (or pitch) has to be low enough and the heart rate high enough such that every z location is imaged during at least N heart beats. Because lowering the pitch generally results in a higher exposure, the technique can only be used at high heart rates.

Since T_w depends on the applied reconstruction algorithm, which is often chosen by the manufacturer and of which the details are not in the public domain, an experiment was designed to determine it empirically.

3.1.2 Amplitude

The motions typically seen in (stent grafts inside) AAA are in the order of 2 mm [91, 108]. The detectability of these motions depends on the localization accuracy of the object in each phase. The accuracy can be increased by using fitting techniques to find the non-integer (subpixel) location between two voxels of an object. However, the localization fit suffers from errors in the found location (localization noise), which can be larger than the motion itself if the motion's amplitude is low. We investigated the amplitude limit below which motions cannot be detected.

3.1.3 Frequency

The data collected during the time T_w results in a single phase. Due to this averaging effect there is an upper limit f_1 on the frequencies that can be measured. For a sinusoidal motion this limit is:

$$f_1 = \frac{1}{2T_w}. \tag{3.1}$$

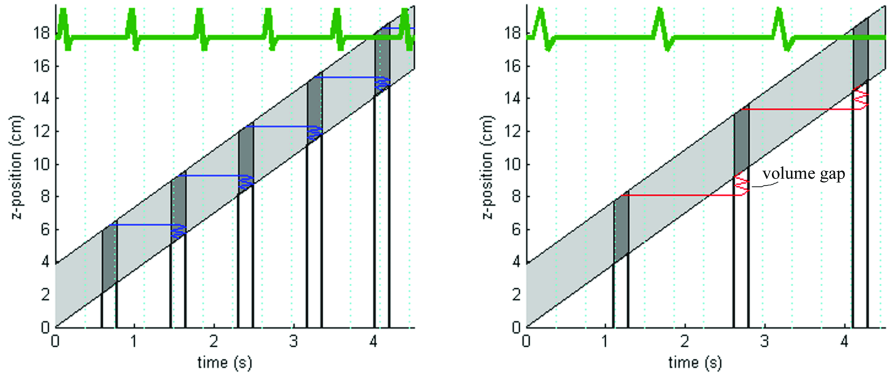

(a) 70 bmp (b) 40 bpm

Figure 3.2: Diagram illustrating the process of ECG-gating. The light grey band indicates the covered z positions of the detector during the scan. The dark grey patches represent parts of the phase in each heart beat. The horizontal lines that connect the patches indicate the measured z-position in subsequent parts of the same phase. They show the overlap (as in a) or volume gap (as in b) between the patches that belong to the same phase. The dotted vertical lines indicate the time at which the gantry is at zero degrees.

In order to measure motions accurately, a sufficient sampling rate is required (Nyquist frequency). This introduces a second upper limit f_2 for the measurable frequencies:

$$f_2 = \frac{1}{2T_d} = \frac{N_{phases}}{2T_{RR}} = \frac{N_{phases} \cdot B}{120}, \tag{3.2}$$

where T_{RR} is the time of one heart cycle, and B the beats per minute. Hereby is shown that patient's heart rate has a linear relation with the maximum frequency, and consequently, may affect the detectability of motion.

Motions in the abdominal aorta are produced by the pressure wave of the blood induced by the pumping of the heart. It has been shown that this pressure has a relatively simple shape: the pressure first increases quickly in around 200 ms and then decreases slowly until the next pressure wave [107, 80, 49]. When the heart rate increases, the shape of the pressure increase is approximately constant.

In the present study the aim is to determine which frequencies can be reliably detected before evident motion artifacts occur. To evaluate whether this is sufficient to reliably measure the motions as they occur in a clinical setting, the result is compared with the frequency components present in a pressure profile measured in vivo in the aortic artery, published in a study by Hazer et al. [49]. Additionally, we investigate whether certain motions, like motions synchronous with the rotation of the scanner, can result in unexpected behavior.

3.1.4 Minimum required heart rate

Figure 3.2 shows a diagram that illustrates the process of ECG gating. The dark patches represent a part of the phase in each heart beat. The overlap in z direction

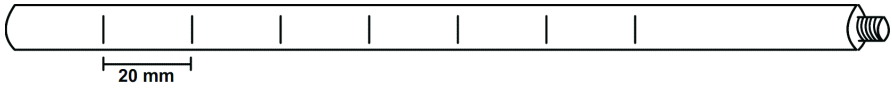

Figure 3.3: Schematic drawing of the phantom used for detecting motion. The phantom consists of a PMMA cylinder with stent wire fragments embedded at 20 mm intervals.

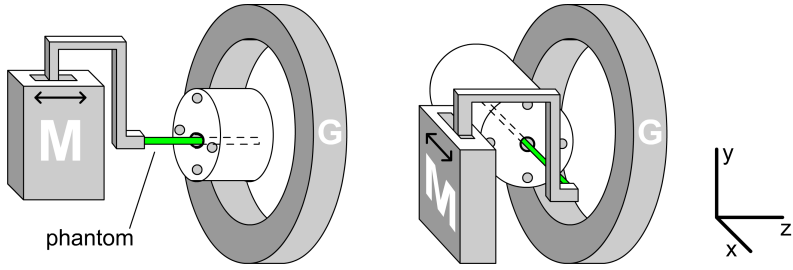

Figure 3.4: Schematic drawing of the setup. The motion unit (M) drives the phantom (Figure 3.3) inside the CTDI phantom's center hole, which is depicted in front of the gantry (G). The left and right hand side show the setup for measuring in the z- and x-direction, respectively.

between these patches depends on the patient's heart rate. Increasing the number of phases will result in subsequent phases being closer together (in time), but two patches of the same phase will remain at equal distance (both in time as in z location).

In Figure 3.2b it is shown that the time between two subsequent heart beats is too large for 40 bpm: the z coverage for two subsequent heart beats does not overlap, but shows a volume gap. To prevent this, the table displacement is limited to the nominal beam width during one heart cycle. The minimum heart rate B_{min} required to prevent volume gaps is given by [96, 58]:

$$B_{min} = \frac{60 \cdot p}{T_{rot}}, \tag{3.3}$$

with p the pitch and T_{rot} the rotation time.

To lower the minimum heart rate the pitch should be reduced, resulting in a longer scan time. Increasing the rotation time is not an option as it would increase motion blur (except for multi segment reconstruction at a certain heart rate). Equation 3.3 shows that if the rotation time is reduced, the pitch should be reduced accordingly. Since the number of photons that contribute to the reconstructed image depends on the rotation time and the tube current only, decreasing the rotation time requires an increase in tube current for the noise to remain the same. Because lowering the rotation time requires also lowering the pitch, the total exposure is increased. A faster rotation thus leads to a higher temporal resolution at the cost of increased exposure [86, 96].

It is of importance to verify the above theoretical limit and to know how the scanner performs in the presence of volume gaps since this can occur in clinical practice.

	A0	A1	A2	A3	A4	A5	A6
# periods per phase	1	1	1	1	1	1	1
heart rate	60	60	60	60	60	60	60
frequency (Hz)	1.0	1.0	1.0	1.0	1.0	1.0	1.0
amplitude (mm)	0.2	0.4	0.7	1.2	2.0	3.0	4.0

	B0	B1	B2	B3	B4	B5	B6
# periods per phase	2	2	2	2	3	3	3
heart rate	45	54	56	60	54.05	60	80
frequency (Hz)	1.5	1.8	1.87	2.0	2.7	3.0	4.0
amplitude (mm)	3.0	3.0	3.0	3.0	3.0	3.0	3.0

Table 3.1: The motion patterns used in the experiments. A0 trough A6 vary in amplitude, while B0 through B6 vary in frequency.

3.2 Materials and Methods

All experiments were performed on a Siemens Somatom 64-slice CT scanner (Siemens Medical Solutions, Erlangen, Germany) with a rotation time of 0.37 seconds, a pitch of 0.34 and $2 \times 32 \times 0.6$ mm collimation. An effective tube current time product of 180 mAs was used at a tube voltage of 120 kVp. The same parameters are used in the clinic, with the exception of the automated tube current modulation, which was turned off for our experiments. Retrospective gating was applied to obtain ten (equal distant) cardiac phases, unless stated otherwise. Each volume was reconstructed using the B36f reconstruction filter and resulted in approximately 80 slices of 512×512 voxels. The slices (thickness 2 mm) were spaced 1 mm apart, and the spacing between voxels in the xy-plane was approximately 0.5 mm.

To quantitatively study motions in ECG-gated CT, a device capable of moving in a predetermined pattern was used (PC Controlled Phantom Device, QRM, Möhrendorf, Germany). It consists of a motion unit that moves a lever, to which a phantom can be attached. The phantom (constructed in-house) consisted of a cylinder made of PMMA (length 160 mm, diameter 10 mm) in which pieces of nitinol wire were embedded at 20 mm intervals (Figure 3.3). The wires (length approximately 6 mm, diameter 0.2 mm) were cut from the framework of a stent graft, and resulted in highly localized points (with a full width at half maximum of approximately 2-3 voxels in the xy-plane) . A standard CTDI body phantom (32 cm in diameter) was used to provide a tissue-like medium and functioned as a guide for the cylindrical phantom to move in (Figure 3.4). To drive the motion unit, seven amplitude patterns and seven frequency patterns were designed (Table 3.1). Triangular motion patterns were used such that the resulting motion was linear and could easily be mathematically described. No assumptions about the shape of the motion were made, since the goal of the present study was to investigate the motion's frequency components individually. To realize a higher frequency, some patterns consisted of multiple triangular periods per cardiac phase. An ECG signal was provided by the motion unit during the scan.

Using profiles A0-A6 (Table 3.1), measurements of the motion amplitude were performed. Note that with amplitude, we refer to the peak-to-peak value of the motion. For each profile one scan was performed in the x direction and one in the z direction (Figure 3.4). The detectability in the x and y direction can be assumed

equal due to the scanner geometry. To quantify the effects of the amplitude on the detectability of motion, the bright spots where the nitinol wires penetrate the slice are detected. Next, a triangular shape with the appropriate amplitude is fitted through the points, and the localization errors of the detected points are calculated by subtracting the fitted triangle from the found locations.

Profiles B0-B6 were designed to investigate the frequency characteristics and were set up to go from low to above the expected maximum measurable frequency. Profiles B0-B3 have a T_{RR} around the expected minimum required heart rate, which is (according to Equation 3.3), 55.1 beats per minute. Profile B4 is designed to move synchronous with the rotation of the scanner. All measurements with profiles B0-B6 were done in the z direction because of practical considerations concerning the setup. To measure the detectability of motion as a function of frequency, the same approach was used as for the amplitude measurements. To be able to compare the results of the frequency measurements with the frequency characteristics present in a clinical setting, the reported pressure profile published by Hazer et. al. [49] was used. The spectrum of the profile was obtained using the fast Fourier transform.

To measure the temporal resolution (T_w) the uniform module of the Catphan phantom (The Phantom Laboratory, Salem, USA) was scanned with ECG gating using the simulation ECG signal of the scanner at 70 bpm. This single scan was then reconstructed eight times with the number of phases ranging from 3 to 10. For 3 phases, there was no overlap between two subsequent phases. For higher number of phases, the overlap between two subsequent phases increases. We measured the correlation coefficient for a set of voxels in two subsequent phases: $\rho_{a,b} = E((A - \mu_a)(B - \mu_b))/(\sigma_a \sigma_b)$, with E the expected value operator, A and B the voxel data of the two phases, μ the mean, and σ the standard deviation. The resulting number (between zero and one) indicates to what extent the noise is correlated (i.e. coming from the same source), and is a measure for the overlap between the two phases. The point at which there is just no overlap between subsequent phases is the point where T_d and T_w are equal. Estimating this point gives us T_w.

We developed algorithms in Python[1] to process the data on a PC. To process the results of the moving phantom scans, the slices penetrated by the nitinol wires in the phantom were manually selected. To compensate for the noise in scans in which motion in the z direction was measured, multiple slices were averaged. Next, the locations of the stent graft wires in the phantom were automatically detected by finding the voxel with maximum intensity in a region where the wire is expected, and the subpixel location is estimated using a polynomial fit. Would a 2D quadratic fit be used, the system of equations is over-determined (nine equations and five unknowns) and the result would be a least squares solution, which is non-interpolating and can therefore deviate more than half a pixel from the detected integer location. Therefore, two 1D quadratic polynomials were used to fit the x and y subpixel location independently.

3.3 Results

To measure T_w, the overlap of the different phases was determined from the correlation of the noise. The correlation between two subsequent phases is shown in Figure 3.5.

[1] www.python.org

For a low number of phases the correlation is zero. From the point where T_w equals T_d the correlation increases as the number of phases increases. It can be seen in Figure 3.5 that this occurs after approximately four phases (for 70 beats per minute). The dotted line represents the correlation of two phases that are separated by one phase, in which case the correlation starts to rise after 8 or 9 phases. In Figure 3.5b the correlation is plotted as a function of temporal distance T_d. The lines incline in a linear fashion, which enables fitting a line through the points and finding the point (on the x-axis) where T_d is equal to T_w. The dashed line shows the fit, which intersects with zero correlation at $T_d = 186.6 \pm 2.4$ ms, which corresponds to half the rotation time as used in the experiment. Hence the exact value of T_w can be assumed to be 185 ms.

The motion detectability was derived from the measured positions (with the mean subtracted) of the detected points and is shown for different amplitudes in Figure 3.6. The scans contain four or more points in each of the ten phases, resulting in at least forty data points per scan. The triangular shape becomes more apparent as the amplitude increases. The absolute error as a function of amplitude (after subtraction of the known triangular shape) is illustrated in Figure 3.7.

To determine which frequencies can be reliably detected, the absolute error was calculated for different frequencies (Figure 3.8). Figure 3.9 shows an example of a detected motion for profile B5 (3Hz). In the introduction we discuss the possibility of unexpected results for motions synchronized with the gantry rotation. This was investigated (using the scan with motion profile B4), but no differences compared to the other scans were detected. Figure 3.10 shows the shape and spectrum of a pressure profile measured in vivo in the aortic artery. It can be seen that the spectrum contains several higher harmonics.

The minimum required heart rate was determined by examining four slices through the phantom at heart rates around the minimum required theoretical heart rate of 55.1 bpm. Figure 3.11 illustrates four example slices at profiles B0-B3. From the lowest heart rate in Figure 3.11 one can clearly observe the noisy bands due to the volume gap, which propagate from top to bottom for increasing phase number. At 54 bpm the bands are still visible, but very thin. For 56 bpm, which is just above the theoretical limit, a band can be observed in some phases (near the top of the shown image for example) on close examination. For 60 bpm, however, the images contain no noise bands. In Figure 3.11a four bars of the phantom can be seen, of which the first, third and fourth from the top are clearly visible. The second, however, seems to have disappeared, while it is clearly visible in the other phases and in the other examples.

3.4 Discussion

In the current study several experiments were performed to evaluate temporal resolution, the effect of amplitude and frequency on the detectability of motion, and the minimum heart rate.

3.4.1 Temporal resolution

The value of T_w was found to be 185 ms which corresponds to half the rotation time. This result strongly suggests that the scanner used a half scan reconstruction

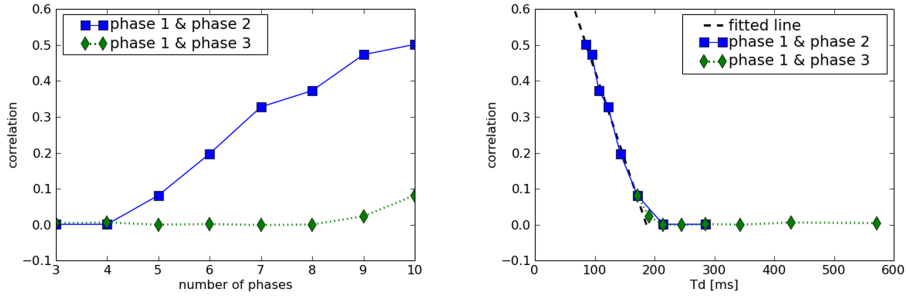

(a) (b)

Figure 3.5: Illustration of the correlation between subsequent phases against the number of phases (a) and against the time between phases (b). The dashed line in b is a linear fit through the seven data points left of the 180 ms mark.

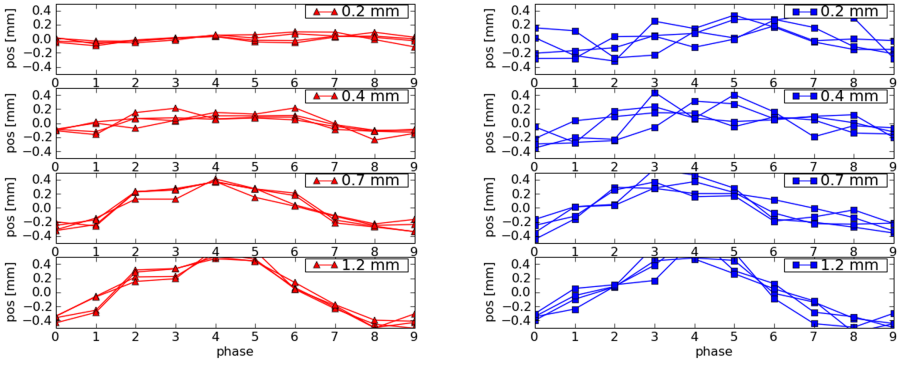

(a) x direction (b) z direction

Figure 3.6: Illustration of the moving position of the points for different amplitudes.

algorithm. However, above a certain heart rate some scanners might switch to multi segment reconstruction, which results in higher temporal resolution [46, 37].

The number of phases to reconstruct should be chosen such that there is overlap between subsequent phases ($T_d < T_w$) even for patients with low heart rates. For our settings and a heart rate of 50 bpm ($T_{RR} = 1.2$ s) this is $1.2/0.185 = 7$ phases. Using more phases results in a higher temporal resolution (in terms of T_d). However, because more than 50% overlap between subsequent phases results in redundant data, a maximum number can also be calculated: for a heart rate of 50 bmp this is achieved at $2 \times 1.2/0.185 = 13$ phases. We can thus conclude that for ECG-gating (on our scanner type) using eight to twelve phases is a reasonable choice.

The described experiment enables measuring the temporal resolution in a generic and reliable way, is applicable to other scanner types, and can be performed on any phantom with a uniform volume.

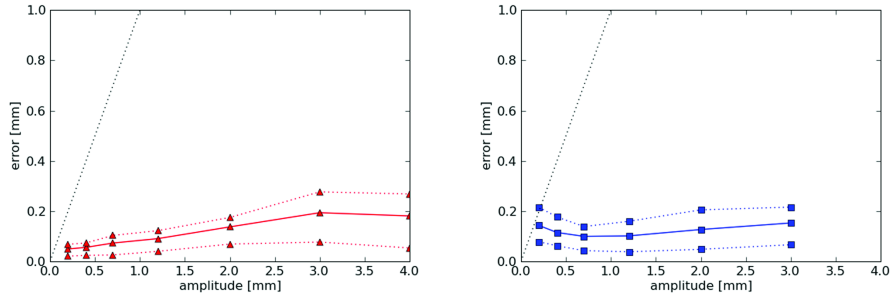

(a) x direction (b) z direction

Figure 3.7: Illustration of the error versus amplitude. The solid line represents the mean absolute error. The dotted lines are the 25 and 75 percentile of the sorted absolute error of the 40+ datapoints in each experiment. The dotted 45 degrees line indicates where the error and amplitude are equal.

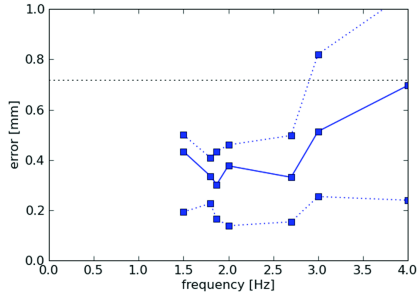

Figure 3.8: Illustration of the error versus frequency. The solid line represents the mean absolute error. The dotted lines are the 25 and 75 percentile of the sorted absolute error.

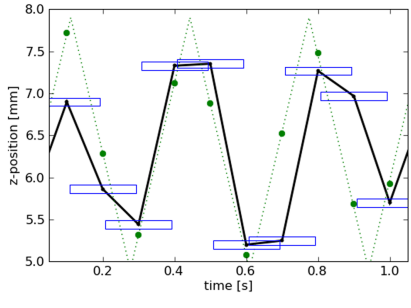

Figure 3.9: Example of the detected motion (solid) of a point at 3.0 Hz and the known profile (dotted). The horizontal boxes indicate the temporal width T_w of 185 ms.

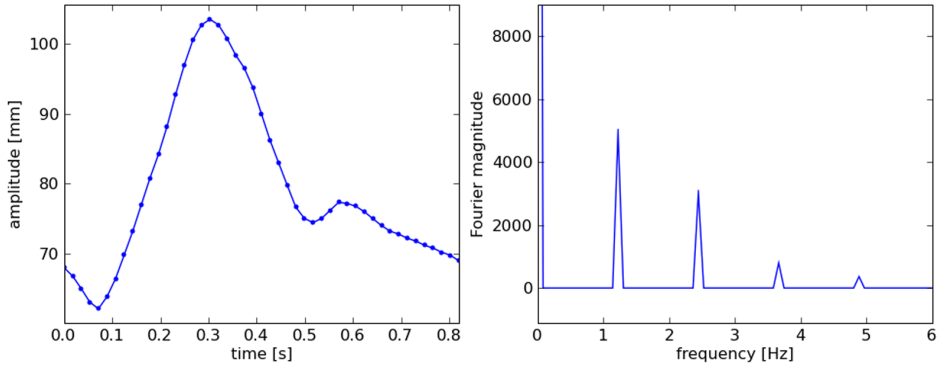

Figure 3.10: Illustration of the shape (left) and the Fourier response (right) of a pressure profile in the aortic artery, as reported in the literature [49].

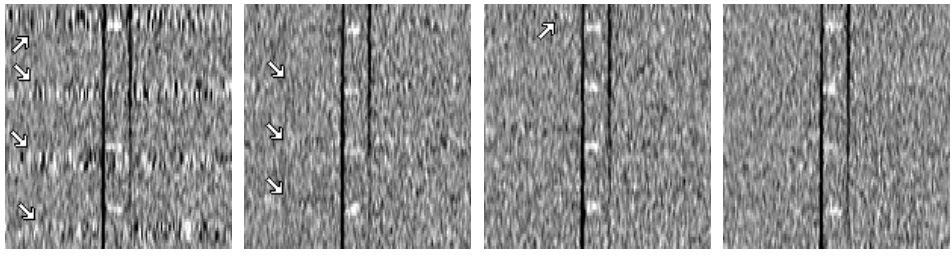

(a) B0 (45 bpm) (b) B1 (54 bpm) (c) B2 (56 bpm) (d) B3 (60 bpm)

Figure 3.11: Illustration of the of noise bands in the CT images, caused by the volume gaps due to a too low heart rate during scanning. At 45 bpm the (horizontal) noise bands are clearly visible (indicated by the arrows). It can be seen how it hides the second bar from the top. At 54 bpm the noise bands are very thin. At higher heart rates no noise bands can be detected.

3.4.2 Amplitude

Figure 3.6 and Figure 3.7 show that, as expected, the error in localization is higher in the z direction because the voxel size is approximately twice as large as in the x direction (1.0 mm versus approximately 0.5 mm). From Figure 3.7 it can be seen that, as anticipated, the error is nearly constant. The slight slope is probably due to the effect of motion artifacts, which become more prominent as the amplitudes increases. In Figure 3.7a and Figure 3.7b the amplitude exceeds the noise level when the error is to the right of the dotted 45 degrees line. Naturally, this is not an abrupt process: the motion will emerge from the noise with increasing amplitude. Nonetheless, from Figure 3.6 it can be seen that amplitudes as small as 0.4 mm in the x direction and 0.7 mm in the z direction can be detected.

In the experiments for the amplitude measurements it is of importance that the phantom moves accurately according to the intended profile. Two sources of error can be distinguished. First, the motion unit. According to its specification, the precision of the start position of the motion unit is better than 0.2 mm and the reproducibility

of motion profiles in quasi-stationary state better than 1%. This suggests that the device may introduce small errors for the lower amplitudes. Second, the transfer of the motion to the phantom. In the z direction this transfer was realized by directly attaching the phantom to the lever. The lever was then fixed in such a position (using bolts) that the phantom moved smoothly in its guide. Due to mechanical restrictions for the x direction, however, the motion had to be transferred via a corner piece, which allows some minor bending. By ensuring the phantom moved smoothly in its guide, the friction was reduced as much as possible. The fact that the accuracy was found to be better in the x direction suggests that the measurements were not unduly affected by the latter source of possible error.

Noise in the image data causes errors in the subpixel localization. Thus it is expected that when images with less noise would be produced, the accuracy will increase. However, to realize the latter, exposure will have to be increased. In our experiments we used an exposure comparable with that used in the clinical setting.

3.4.3 Frequency

The B0 measurement (1.5 Hz) had a relatively large error due to the heart rate that was too low during this measurement (Figure 3.8). The next three measurements had a relatively low error, and from 2.7 Hz and up the error increases.

With ten phases and a heart rate of 50 bpm, $T_d = 120$ ms and the Nyquist frequency is found at 4.2 Hz (Equation 3.2), which is well above the frequencies examined. Since higher heart rates give rise to even higher Nyquist frequencies, the Nyquist criterion is of minor importance for our experiments.

Due to the temporal width of the phases, the effect of motion artifacts increases with increasing frequency, until they reach an upper limit, above which the motion should not be measurable. Given $T_w = 185$ ms, the upper limit is found at 2.7 Hz (Equation 3.1). However, the results suggest that motion at 2.7 can be measured relatively well. To study this in more detail, the detected motion of a point at 3 Hz is shown in Figure 3.9, indicating the temporal phase width T_w using horizontal boxes of 185 ms. (Half scan reconstruction was assumed, as the bpm in this measurement was lower than that used in the temporal resolution experiment.) As 185 ms is over half the period of motion (167 ms), we would expect the estimated motion to be poor, yet (for most points) the estimated locations are good. This surprising result can be explained by the way the data—acquired during half a rotation—is processed; because of the reconstruction, motion during the acquisition results in highly localized motion artifacts and not necessarily blurring. Therefore, the location of the detected point can still be relatively accurate.

Figure 3.10 shows that the pressure (and thus the motion) in the aortic artery contains frequency components higher than 2.7 Hz. This will express itself in motion artifacts in the phases acquired during the sudden rise of pressure at the start of the cardiac cycle. Consequently, these phases can be rendered useless if the motion artifacts are too strong, in which case the motion needs to be estimated from the other phases. The cause of this problem is the width of the phase T_w, which can be reduced by using smaller rotation times or using dual source CT. However, as we have discussed in section 3.1.4, this requires an increase in exposure in order for the minimal required heart rate to remain equal.

3.4.4 Minimum required heart rate

The results in Figure 3.11 shown that the minimum required heart rate is close to the theoretical limit of 55.1 bpm calculated in section 1.4. In a volume gap, the voxel data is undefined. Nevertheless, the scanner attempts to fill in the gap; it can be seen that the black vertical lines (the air between the phantom and its guide) are continued in the bands, which suggests interpolation. The fact that the bar in Figure 3.11a was completely hidden implies that the scanner fills in the bands to some extent, but the data is very noisy and not reliable.

3.5 Conclusions

We performed experiments to investigate the effect of amplitude and frequency on the detectability of small motions in ECG-gated CT. Also investigated were temporal resolution and minimum required heart rate. The experimental methods can be applied to CT scanners of other manufacturers.

The experiment designed to measure the temporal resolution empirically clearly showed that the duration of each ECG-gated phase is 185 ms for our scanner and settings, which corresponds to half the rotation time. The other experiments showed that motions become detectable from amplitudes as small as 0.4 mm in the x direction and 0.7 mm in the z direction. Motions up to 2.7 Hz can be accurately detected. Volume gaps caused by a too low heart rate are expressed in noisy bands in the data that propagate in the z direction. The reconstruction algorithm uses some form of interpolation, but cannot prevent objects in volume gaps becoming hidden.

This study gives insight into the possibilities and limitations for measuring small motions using ECG-gated CT. From the results we conclude that CTA is a suitable technique for studying the expected motions in AAA.

Acknowledgment

The authors would like to thank Daan van der Vliet for the fruitful discussions.

4

Diagnostic quality of time-averaged ECG-gated CT data

This chapter is published in slightly modified form at the SPIE Medical Imaging conference in Orlando, USA (2009) [63].

Abstract

In Chapter 3 we found that ECG-gated CTA is indeed a suitable technique for studying the motions of stent grafts. However, the individual phases have a lower signal to noise ratio (SNR) than data obtained using non-gated CT. In order not to obstruct the clinical care for the patients involved in our study, data should be available for diagnosis that has at least the quality as a non-gated CT scan (without having to increasing the dose).

Some common CT scanners cannot reconstruct a non-gated volume from ECG-gated acquired data. In order to obtain the same diagnostic image quality, we propose off-line temporal averaging of the ECG-gated data. This process, though straightforward, is fundamentally different from taking a non-gated scan, and its result will certainly differ as well. The purpose of this study is to quantitatively investigate how good off-line averaging approximates a non-gated scan.

The experiments show that the spatial frequency content is not affected by the averaging process. The minor differences observed for the noise properties and motion artifacts are in favor of the averaged data. Therefore the averaged ECG-gated phases can be used for diagnosis. This enables the use of ECG-gating for research on stent grafts in AAA, without impairing clinical patient care.

4.1 Introduction

In recent years major advancements are made in computed tomography (CT). Shorter rotation times and the development of multi detector CT (MDCT) enable the technique of ECG gating, often referred to as cardiac CT. [45] ECG gating uses the ECG signal of the patient to divide the raw scan data into bins that correspond to consecutive phases of the heartbeat. The data is reconstructed to a number of volumes, corresponding to the pertinent phase of the heart cycle. This allows 4D visualization of the scanned object and enables investigation to its temporal behavior.[45, 94] ECG gating is used extensively in cardio CT, and is increasingly popular for research to aortic abdominal aneurysms (AAA)[65, 90, 91, 108, 119]. Our research focuses on the possibilities and limitations of ECG-gated CTA for investigating motions in AAA.

In our institution, when patients are recruited for research to stentgraft motion in AAA, the non-ECG-gated scan that the patient normally receives is replaced by an ECG-gated scan. A drawback of ECG gating is that less data is available per volume, which results (as the dose is kept the same for both protocols) in relatively noisy data and artifacts being more common. Hence, for diagnostic purposes, the clinic requires the result of a non-gated scan. Unfortunately, not all scanners are capable of producing a non-gated three dimensional scan in case ECG-gated scanning was used. Scanning patients twice is not an option considering the extra dose this would imply.

We propose to average the data of the phases off-line (i.e. not on the scanners reconstruction computer) to create the 3D dataset required. This is a straightforward process, yet fundamentally different from combining the sinogram data before the filtered backprojection reconstruction (as happens for a non-gated scan). Due to non-linearities in the reconstruction process of the scanner, the results may be similar, but will never be exactly the same.

The purpose of the experiments described in this document is to quantitatively investigate how good off-line averaging approximates a non-gated scan. A positive outcome of our research implies that studies using ECG gating can be performed without obstructing patient care, also on scanners that do not support producing regular scan data from an ECG-gated scan. As to our knowledge, this has not yet been investigated.

A comparison is made between three datasets: the non-gated scan, a single phase of the ECG-gated scan, and the averaged phases of the ECG-gated scan. The quality of the different datasets is assessed by comparing noise properties, frequency response, and motion artifacts. Noise properties are best compared based on the Noise Power Spectrum (NPS): it provides a better method for investigating noise properties compared to calculating the standard deviation, because it shows in what frequencies the noise is expressed.[14] Frequency response is compared based on the Modular Transfer Function (MTF). For ECG-gated studies, motion is generally present (or ECG gating would not be necessary). Therefore it is of importance to know the effect of motion artifacts and how they are expressed in an averaged ECG-gated dataset compared to a dataset obtained with a non-gated scan.

4.2 Methods

Experiments have been performed using a non-ECG-gated protocol and an ECG-gated protocol. The settings of the first were chosen as similar to the ECG-gated

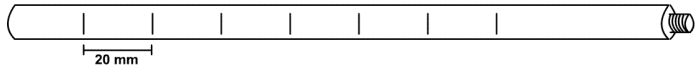

Figure 4.1: Diagram of the phantom used for detecting motion. The phantom consists of a perspex cylinder with nitinol wire fragments embedded at 20 mm intervals.

protocol as possible. Both experiments were performed on a Siemens Somaton 64 CT scanner (Siemens Medical Solutions, Erlangen, Germany). A fixed effective mAs of 180 was used at 120 kVp, and with $2 \times 32 \times 0.6$ collimation. The slices had a thickness of 2 mm and were spaced 1 mm apart. For the ECG-gated protocol a rotation time of 0.37 seconds and a pitch of 0.34 were used. For the non-ECG-gated protocol the rotation time and pitch were chosen as close to the ECG-gated protocol as the scanner allowed: 0.5 s and 0.45 respectively.

Data from the ECG-gated protocol was reconstructed to ten cardiac phases. A third dataset was created by averaging the ten phases of the ECG-gated scan, thus producing a three dimensional dataset comparable to the data of the non-gated scan.

To study the noise and frequency response, the Catphan500 phantom (The Phantom Laboratory, Salem, USA) was scanned. From the resulting datasets the Noise Power Spectrum (NPS) and Modular Transfer Function (MTF) were determined. For this purpose automated software developed in-home was used. The NPS was calculated from the uniformity module of the Catphan phantom. The MTF was calculated from the two beads present inside the linepair module. Both the NPS and MTF were calculated via a 2D FFT using a method similar to that described by Boedeker et al.[14]

To study motion artifacts for the different protocols, a device capable of moving in a predetermined pattern was scanned (PC Controlled Phantom Device, QRM, Möhrendorf, Germany). It consists of a motion unit that moves a lever, to which a phantom can be attached. The used phantom (which was developed in-home, see Figure 4.1) consists of a perspex cylinder (length 160 mm, diameter 10 mm) in which pieces of nitinol wire (cut from a stent) are embedded at 20 mm intervals. A standard dose phantom (the CTDI phantom) was used to provide a tissue-like medium and functioned as a guide for the cylindrical phantom. The device was set to produce triangular motions at 60 beats per minute with a (peak-to-peak) amplitude of 3 mm. The setup was such that the motion was lateral (in the x-direction). The resulting motion artifacts were quantified by measuring the full width at half maximum (FWHM) of the nitinol wires as they appear in the scanned data.

4.3 Results

The NPS signals were determined for the non-gated, single phase, and averaged phases (Figure 4.2). The standard deviation of the three datasets were calculated by taking the square root of the mean of the NPS (Table 4.1). Furthermore, the average of the MTF signals of both beads was determined (Figure 4.3).

In the three volumes resulting from the motion experiment, the slices penetrated by the nitinol wires were selected manually (Figure 4.4), and a profile was selected through the peaks (Figure 4.5). The FWHM of the peaks are shown in Table 4.2.

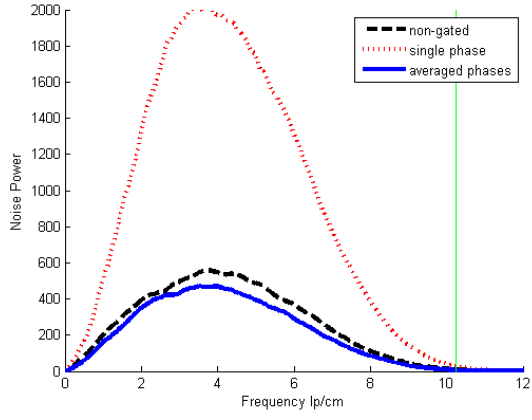

Figure 4.2: Illustration of the Noise Power Spectrum (NPS) of each dataset. The straight vertical line indicates the Nyquist frequency.

Method	SD
non-gated	13.6
single phase	26.8
averaged phases	12.5

Table 4.1: The standard deviations calculated from the NPS.

Method	FWHM
non-gated	6.6
single phase	4.8
averaged phases	5.9

Table 4.2: The average full width at half maximum for the different scans, measured from the wires in the moving phantom.

4.4 Discussion

The result of the NPS (Figure 4.2) shows that, as expected, for the single ECG-gated phase the noise is high compared to both other measurements. It also shows that averaging the ECG-gated phases leads to a slight reduction in noise, especially for frequencies around 5 line pairs per centimeter. Using a two-sample t-test on measurements from five slices in the uniformity module of the phantom, the standard deviation is indeed found lower ($p<0.0005$). The measured standard deviations show the same relation between the three datasets. The standard deviation of the averaged phases is around 10% lower than the standard deviation of the non-gated scan. As the effective mAs (the tube current divided by the pitch) is equal for both scans, the radiation dose is equal as well. The precise reason for the reduced noise is un-

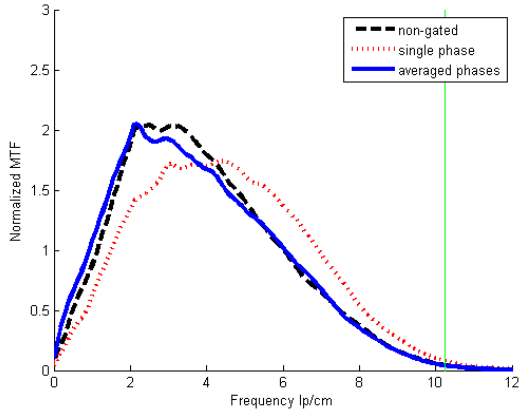

Figure 4.3: Illustration of the Modulation Transfer Function (MTF) of each dataset. The straight vertical line indicates the Nyquist frequency.

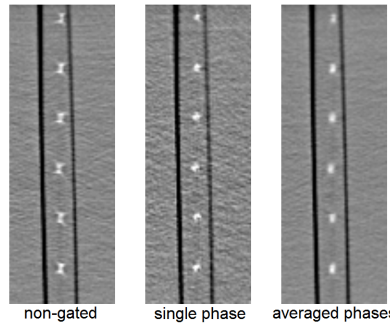

Figure 4.4: Illustration of the slices (from the three datasets) containing a moving object.

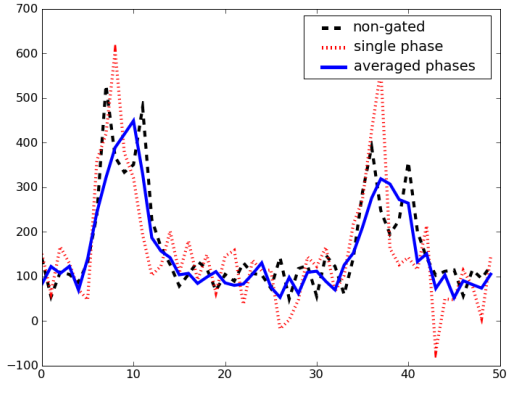

Figure 4.5: Profiles through the slices of Figure 4.4.

known, but probably can be found in the way that data is combined (averaged) in the reconstruction of a non-gated scan.

The approach for measuring the MTF from the beads differs from Boedekers [14] in the normalization step. When the value at 0 lp/cm is used to normalize, the normalization proofed not to be consistent, and comparison of the three methods not reliable. To solve this problem, we have chosen to normalize by the mean of the MTF, taking into account all frequencies below the Nyquist frequency. The motivation is that an MTF can only show the relative power for each frequency. In our approach the integral over all frequencies (up to the Nyquist frequency) is one.

The results of the MTF (see Figure 4.3) show that the MTF of the single phase differs from that of a non-gated scan. The MTF of the averaged phases is similar to the non-gated dataset, which suggests that the temporal averaging did not have any negative effects on the spatial frequencies of the data.

From the partial slices in Figure 4.4 and the profiles in Figure 4.5 it can be seen that for the single phase the peaks are more narrow because the data acquisition time is lower compared to both other datasets. Due to the motion blur that occurs when multiple phases are averaged in which the object is present at slightly different locations, the profile for the averaged phases is wider. For the non-gated scan, however, the data is "combined" before the reconstruction and thus leads to a different kind of artifact, as can be seen from its profile. From the measured FWHM values (the averages are shown in Table 4.2), it follows that the artifacts of the averaged phases are smaller compared to the non-gated dataset ($p<0.001$).

In these experiments the pitch and rotation time of the non-gated scan could not be set the same as the ECG-gated protocol. While the effective mAs, and therefore the dose, are the same, different results may be found when different scanner settings are used. Larger differences in the results are probably found on different scanner types as their internal reconstruction algorithms will differ.

We can conclude that—on our scanner—the proposed method is better than taking a normal CT-scan. However, we suspect the method is not feasible to replace a normal CT-scan; while the proposed method can be applied in a matter of seconds, reconstructing the data and sending it to a dicom node will take ten times as long (in the case of ten phases). Additionally, the patients ECG must be measured during the scan, and the patient has to hold his/her breath longer. Our result does show, however, that the current reconstruction techniques for non-gated CT are not optimal.

4.5 Conclusions

The temporal averaging of ECG-gated CTA data does not have negative effects on image quality in terms of noise, frequency response and motion artifacts. The minor differences observed for the noise properties and motion artifacts are in favor of the averaged data. Therefore the averaged ECG-gated phases can be used for diagnosis. This enables the use of ECG-gating for research on stentgrafts in abdominal aortic aneurysms, without impairing clinical patient care.

Part B

Segmentation

image segmentation*: the process of partitioning a digital image into multiple segments. The goal of segmentation is to simplify and/or change the representation of an image into something that is more meaningful and easier to analyze.*

5

Segmentation of stent grafts using a 2D approach

This chapter is published in slightly modified form as "A segmentation method for stent grafts in the abdominal aorta from ECG-gated CTA data" at the SPIE Medical Imaging conference in San Diego, USA (2008) [65].

Abstract

This chapter describes a segmentation approach in which the stent is segmented in 2D slices sampled orthogonal to the centerline. In Chapter 6 it is described how this can be used to segment the 3D stent by means of tracking.

Segmentation of the stent graft is performed by examining slices perpendicular to the centerline. Regions with high CT-values exist at the locations where the metallic frame penetrates the slice. These regions are well suited for detection and sub-pixel localization. Spurious points can be removed by means of a clustering algorithm, leaving only points on the contour of the stent. We compare the performance of several different point detection methods and clustering algorithms. The position of the stent's centerline is calculated by fitting a circle through these points.

The proposed method can detect several stent graft types, and is robust against noise and streak artifacts.

5.1 Introduction

Endovascular Aortic replacement (EVAR) for the treatment of aortic aneurysms in patients at risk of aneurysm rupture is an established endovascular technique [97]. However, the development of the stentgraft design is ongoing in order to prevent late stentgraft failure (e.g. graft breakage, graft displacement). The stresses applied to the stentgraft by the high physiological forces and stresses in the aorta have an effect on the durability and functioning of the stentgraft. The stresses and forces that occur during movement may be reflected by the motion of the stentgraft itself during the cardiac cycle. Knowing the intrinsic properties of the stentgraft, it should be possible to evaluate these stresses and forces by calculating its displacement and deformation.

We apply a cardiac CT technique, ECG-gated CTA (computed tomography angiography), to measure the stentgraft's movements during the cardiac cycle. With a normal CT scan, a 3D volume is imaged. This new method of applying ECG-gating results in ten CT volumes, which each represent a different phase of the cardiac cycle, allowing 4D visualization of the stentgraft.

5.1.1 Purpose

The purpose of our research is to gain more insight in the dynamic behaviour of the stentgraft in vivo. Knowledge about the stresses acting on the stentgraft will help understand the functionality and failure of different types of stentgraft over the long term. These insights can help in stent design. Furthermore, the information on the behaviour (during the cardiac cycle) and the long term functioning of different types of stentgrafts can aid the treating physician in deciding which stent-type to use for a specific patient.

5.1.2 Problem Statement

The movements of a stentgraft can only be measured if we know its exact position. Therefore, we must first accurately segment the stentgraft. One of the challenges of ECG-gated CTA is analysing the large dataset acquired: in the order of 4000 image slices per patient. Automatic detection of the metal frame of the stentgraft will improve and speed up data processing.

Unfortunately, a simple thresholding algorithm is not sufficient for detecting the stentgraft. The first reason is that the data is relatively noisy. For patient safety, the total dose is kept similar to a normal CT scan. However, each volume is reconstructed from the fraction of the total dataset related to the specific phase of the cardiac cycle. Therefore the dose per volume is less, resulting in a lower SNR. Second, streak artifacts occur where the stentgraft's metal frame is thick or where a coil is present next to the stentgraft; due to the lower SNR, their effect is larger than in regular CT. Third, contrast agent injected in the blood results in CT-values close to the range of CT-values seen for most stentgrafts. Fourth, different stentgraft types are made from different materials (with different absorption coefficients) and differ in the thickness of their frames. In some images, stentgrafts are seen to have CT-values as low as 300 HU (Hounsfield Units).

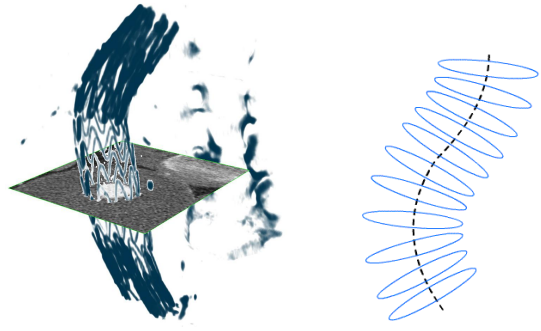

(a) (b)

Figure 5.1: Volume rendering of a stent and a plane perpendicular to the stent's centreline (a), and an illustration of contours perpendicular to a centreline (b).

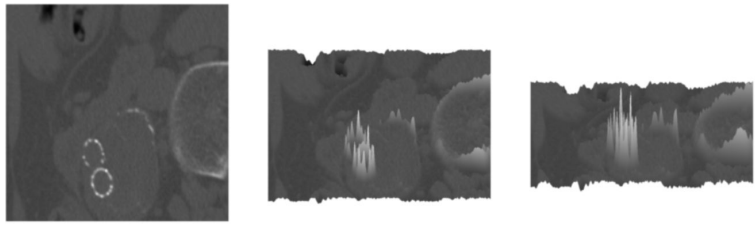

Figure 5.2: The surface plot of a slice shown at different angles to illustrate the pointy structures that form the contour of the stent.

5.1.3 Approach

A stent has a tubular structure, sometimes with branches and can be modelled by a series of stacked contours which are orthogonal to the centreline (see Figure 5.1) [42]. Our approach is to segment the stent in images sampled perpendicular to its centreline, as shown in Figure 5.1a. Regions with high CT-values (typically above 500 HU) exist where the metallic frame of the stentgraft penetrates the image. These regions have high CT-values and are "pointy" (see Figure 5.2), which makes them well suited for point detection. By segmenting the stent's contours and iteratively estimating its centreline the stent can be tracked, similar to how vessels are tracked in a recent paper by Lee et al. [76]. An advantage of this 2D approach is that visualization and algorithm design is easier in 2D than it is in 3D. In this document, we discuss the process of finding the best way to detect the contour of the stent in images sampled perpendicular to its centreline.

5.1.4 Overview

In our research we have considered various methods and combinations thereof. In section 5.2 we will discuss the different methods. In section 5.3 we compare the methods with an experiment to be able to select the best method.

5.2 Methods

In this paper we focus on segmenting the stent in a slice perpendicular to its centreline. Our approach consists of two tasks:

- Detect points on the contour of the stent.
- Find the cluster of points that really belong to the stent, removing spurious points.

In this paper we will consider four point detection methods and three clustering methods. Many of these methods use algorithms for specific basic tasks. These basic algorithms are first discussed in subsection 5.2.1. In the six subsections after that we will discuss the following point detection and clustering methods:

Point detection methods		**Clustering methods**	
static threshold	(5.2.2)	fitting	(5.2.5)
dynamic threshold	(5.2.3)	snake	(5.2.6)
Laplacian curvature	(5.2.4)	chopstick	(5.2.7)
Gaussian curvature	(5.2.4)		

5.2.1 Basic Algorithms

In this subsection we discuss the algorithms that are used as part of the methods described in this section.

5.2.1.1 Removal of metal artifacts

Metal artifacts are characterized by streaks of alternating high and low CT-values. The high streaks can cause false points to be detected. The low streaks, on the other hand, are often below 0 HU. Physiologically, tissue with negative CT-values is very unlikely to occur in the local environment of the abdominal aorta because it suggest tissue less dense than water.

We remove spurious points due to streak artifacts in all point detection methods discussed in this paper. For example in the methods of 5.2.3 and 5.2.2 such points can be found by checking for abnormally low CT-values in one of the direct neighbours of a pixel under examination.

5.2.1.2 Circle Fit Algorithm

The stent is designed to have the shape of a blood vessel and is thus approximately circular. Because the stent is easily deformed, however, the stent is compressed in a shape that is in general *not* a perfect circle at all. By fitting the data with a perfect circle, we generate a robust measure for the centre and radius of the stent, even when the points are not distributed in a full circle. This provides a valuable tool that is used by all of our clustering methods.

We have implemented the Modified Least Squares algorithm of Umbach and Jones [112]. Their method is a fast closed form solution and performs well even with noisy data or when only points on part of the circle are present. For details on the algorithm, we refer to their paper [112].

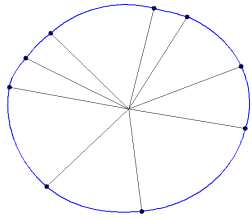

Figure 5.3: Illustration of the way contour points define the circular spline.

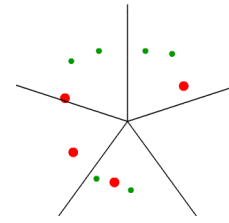

Figure 5.4: Illustration of the working of the centre estimate convergence algorithm.

5.2.1.3 The Circular Spline

The clustering methods find points that lie on the contour of the stent. To produce a more continuous contour on which we can perform calculations, intermediate points are interpolated. We choose to interpolate in the polar coordinate system. This will produce a smooth circular shape even when large gaps are present. The centre point required to convert the points (x, y) to (r, ϕ) is estimated using the circle fit algorithm of 5.2.1.2. In the polar coordinate system, we use bicubic interpolation (using a cardinal spline [113] with tension $-\frac{1}{4}$) to connect the points. Figure 5.3 illustrates the results after the spline is converted back to Euclidean space.

5.2.1.4 Centre Estimate Convergence Algorithm

Given an initial estimate of the centre of the stent[1], and the set of points found by the detection method, this algorithm calculates a better estimate of the centre and the radius. The initial estimate can be near the edge of the contour. It is used by all clustering methods discussed in this paper.

The algorithm first divides the image in five segments emanating from the centre estimate (see Figure 5.4). Next, for each segment (or direction), the closest point to the centre is selected. Then the algorithm discards points that are too far away from the centre (two times the distance of the third closest selected point). The selected set thus has a maximum of five points, which are distributed in multiple directions. Finally, the circle fit algorithm is used to estimate the centre. This process is repeated until the centre position does not change. By using five directions the algorithm will succeed (find at least three points) with gaps of over 180 degrees.

5.2.2 Point Detection: Static Threshold

The simplest method to detect points is to detect local maxima: points which have a larger pixelvalue than their eight neighbours. A threshold will prevent points to be found in noise. This threshold is the main parameter of this method and should be low enough to detect most points, but high enough to prevent too many spurious points, for instance in contrast fluids. The threshold was empirically determined at 650HU.

[1]This centre is now manually selected, but will later be calculated from the estimated centreline found in previous slices.

5.2.3 Point Detection: Dynamic Threshold

This method too, detects local maxima and requires pixelvalues to exceed a certain threshold. In this case however, the threshold is dynamic: it depends on the local environment of the pixel under examination. The threshold is applied to the pixelvalue subtracted by the average of a local patch around it. The threshold for this method is set at 260HU, which is—as expected—lower than for static thresholding. This method allows us to find parts of the stentgraft that make a small contribution to the pixel intensity, while minimizing the amount of spurious points, especially in scans where contrast fluid was used.

5.2.4 Point Detection: Hessian Based

Differential geometry (in image processing) means extracting information (structure) using the derivatives of images [69]. As can be seen in Figure 5.2, the spots where the stent penetrates the slice are very intense, but what is even more distinguishing is their pointy structure. This aspect can be detected using second order derivatives. The Hessian matrix is an often used concept to describe second order properties in 2D or 3D. In 2D it is given as:

$$H(L) = \begin{bmatrix} \frac{\partial^2 L}{\partial x^2} & \frac{\partial^2 L}{\partial xy} \\ \\ \frac{\partial^2 L}{\partial xy} & \frac{\partial^2 L}{\partial y^2} \end{bmatrix}, \tag{5.1}$$

with $L(x, y)$ the image. In a pixel, the Hessian matrix can be used to calculate the second order derivative in any direction. In the spots on the contour of the stent, the second order derivatives are very high compared to the rest of the image. Therefore, we consider the following measures:

- The sum of the Eigenvalues. Also known as the Laplacian, or the trace of the Hessian matrix: $\frac{\partial^2 L}{\partial x^2} + \frac{\partial^2 L}{\partial y^2}$.
- The product of the Eigenvalues. Also know as the Gaussian curvature or $\frac{\partial^2 L}{\partial x^2} \cdot \frac{\partial^2 L}{\partial y^2}$.

The functions $\frac{\partial^2 L}{\partial x^2}$ and $\frac{\partial^2 L}{\partial y^2}$ are calculated by convolving the image with (second order) Gaussian derivative kernels with $\sigma = 1.6$. For a pixel to be classified as a candidate point it has to be a local maximum and the Laplacian or Gaussian curvature in that point has to be larger than a certain threshold. This threshold is the main parameter of the methods and are empirically determined at 95 and 2000 for the Laplacian and Gaussian curvature respectively, where the pixelvalues are scaled at one unit per CT-value.

5.2.5 Clustering: Fitting

The first proposed method to segment the stent is to iteratively fit circles and ellipses to the stent, assuming that its shape does not deviate too much from an ellipse. The method starts by running the algorithm of 5.2.1.4 to produce an initial circle estimate. Next, the algorithm calculates which points are closer than $2 \cdot t_{\text{tol}}$ to this

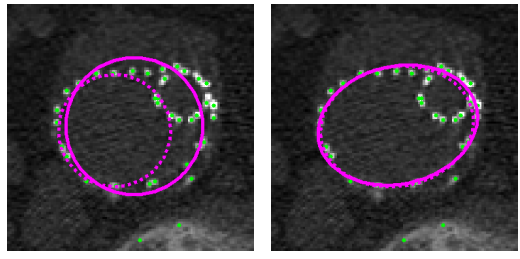

(a) initial and converged circle (b) initial and converged ellipse

Figure 5.5: Illustration of the fitting method applied to an example.

circle (where t_{tol} is the tolerance measure) and selects these points. A new circle is fit to the selected points and the process repeats itself until the circle does not change. Then the method will fit ellipses rather than circles and points must now be closer than t_{tol} to the ellipse. When the ellipse estimate is converged the selected points are the contour. See Figure 5.5 for an example. The tolerance measure t_{tol} is the single parameter of this method; we used a value of $t_{\text{tol}} = 0.1mm$.

Fitting an ellipse and calculating the distance of a point to an ellipse are much more complex then for a circle. The algorithm we use to fit ellipses is proposed by Halir et al. [48] in 1998, and an improvement on the original algorithm of Fitzgibbon et al. [40] To calculate the distance of a point to an ellipse we applied the algorithm proposed by Eberly [34].

5.2.6 Clustering: Snake

Active contours (or snakes) are often used in medical imaging for segmentation. A general challenge of active contours is to influence the snake where it is in a homogeneous region. We have chosen to solve this problem by implementing the method of Xu and Prince: gradient vector flow [124] (GVF). Another solution is to use a balloon force such as in Cohen et al. [23] and more recently in Dedbleds-Mansard et al. [29] We have seen, however, that in our setting the balloon force can force the snake over the points, which allows the snake to grow unlimited until iteration stops.

The GVF method iteratively generates a vector field (the GVF field) to guide the snake. This field enables the snake to move into boundary concavities, gives the GVF snake a large capture range, and makes it insensitive to initialization. Especially the latter two features are expected to make this type of snake very robust and work well in our setting.

We generate a black image with white pixels on the locations of the detected points, and create a GVF field from this image. We initialize the snake by sampling points on the circle estimated by the algorithm of 5.2.1.4. This makes the method invariant to spurious points inside the contour, see Figure 5.6 for an example. After the snake has converged, all points closer than 1.4 pixels to the snake are selected as the contour.

One drawback of snakes in general is their large amount of parameters. In addition to *elasticity*, *rigidity*, *viscosity*, *external force weight* we can also change the sample

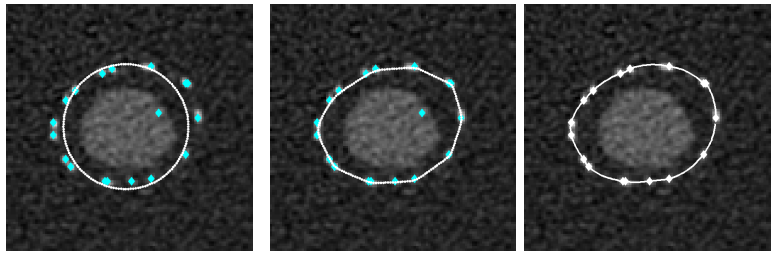

(a) initial snake (b) converged snake (c) circular spline

Figure 5.6: Illustration of the snake method applied to an example.

distance of the snake. On top of that, the resolution (number of pixels per mm) of the created image also has influence on its behaviour. In our experiments we used the slice's resolution because the resolution among different datasets only differs by a few percent. We used a relatively large rigidity to prevent the snake from leaking through the larger gaps between points. The appropriate rigidity value was found to be 20.

5.2.7 Clustering: Chopstick

We propose an algorithm that uses the internal geometry of the detected points to select the points on the contour of the stent. It is designed for this clustering problem in specific and has the advantage of an (in our opinion) intuitive working.

The first step consists of selecting a coarse set on the contour, which we will explain using a virtual stick. The stick is attached to a point found by the algorithm of 5.2.1.4 and aimed at the centre. We rotate the stick around the initial point until it touches another point, which we then take into our selection. We attach the stick to the new point and the process is repeated until we reach a point that we already selected. In Figure 5.7 an example result is illustrated.

The next step consists of adding and removing points. Consider two points n_1 and n_2 that are connected by the stick in the first step. The algorithm searches for a point in between n_1 and n_2 that is suitable for selection. If such a point is present, it is taken into the selection; the stick is chopped in two pieces to connect n_1 and n_2 to the new point. After searching in between all points, we examine each point using its two direct neighbours (n_1 and n_2), and remove the point if it is not acceptable, using the same method. This process of adding and removing points is repeated until the selected set has converged.

The method to test the suitability of points is based solely on (the geometry of) the two reference points (n_1 and n_2) and the centre point. It adopts two simple criteria which are illustrated by three examples in Figure 5.8. Consider c the centre point, p the candidate point, and $\overrightarrow{n_1\, n_2}$ and $\overline{n_1\, n_2}$ the vector and distance from n_1 to n_2 respectively. We calculate a point n_3 which lies exactly in between n_1 and n_2, and a point n_4 which lies on the normal of $\overrightarrow{n_1\, n_2}$ (see Figure 5.7a).

The first criterion requires that the point must be in between the lines $\overrightarrow{c\, n_1}$ and $\overrightarrow{c\, n_2}$. This intuitively focusses at points "in line of sight". In the examples of Figure 5.7

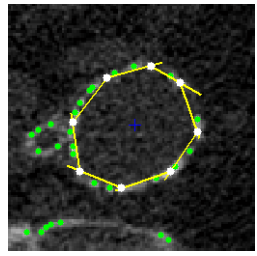

Figure 5.7: Example result after finding a coarse contour.

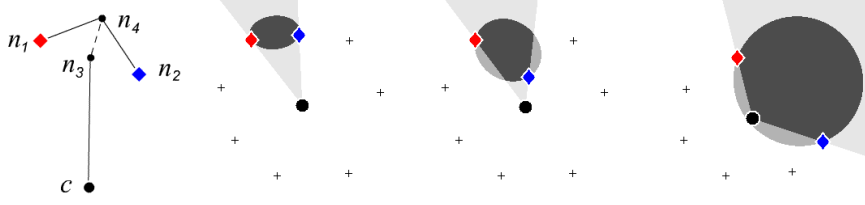

(a) Overview (b) Example 1 (c) Example 2 (d) Example 3

Figure 5.8: Illustration of the region in which points are found suitable by the chopstick method. The light, medium and dark shades of grey correspond to the first, second and combined criteria respectively.

this is illustrated by the light grey area.

The second criterion requires the point to be inside an ellipse spanned between n_1 and n_2. An ellipse in general can be defined by two points and a constant D_{ellipse}. A point p is inside the ellipse if the distance from one ellipse-point to p and back to the other ellipse-point is smaller then D_{ellipse}. We position the two ellipse-points on the lines $\overrightarrow{n_1\ n_4}$ and $\overrightarrow{n_2\ n_4}$ respectively, and chose D_{ellipse} such that the ellipse just reaches n_1 and n_2. The distance from n_4 is equal for both ellipse-points and determines the eccentricity of the ellipse. We use a measure f to describe this distance, where $f = 0$ means both ellipse-points are in n_4 (the ellipse is a circle) and $f = 1$ means the points are in n_1 and n_2 respectively (the ellipse is maximally elongated). We want the shape of the ellipse to depend on $\overline{c\ n_3}$, being close to a circle when c is close, and elongated when c is far. We propose for f:

$$f = 1 - e^{-\overline{c\ n_3}\ /\ \overline{n_1\ n_2}} \tag{5.2}$$

The displacement of n_4 relative to n_3 is applied because points are more likely to be found on the outer side of the line $\overrightarrow{n_1\ n_2}$ due the circular nature of the contour. The distance of n_4 to n_3 is varied by the factor f of (5.2):

$$\overline{n_3\ n_4} = \frac{1}{4} \cdot \frac{\overline{c\ n_3}}{f} \tag{5.3}$$

We will shortly discuss the examples in Figure 5.8 for clarification. Example 1 shows eight points on the contour and the region between n_1 and n_2 where points are justified according to the described criteria. In example two n_2's location is different;

this point will be removed in a later stage of the algorithm. Example three illustrates a large gap between n_1 and n_2. It can be seen how the displacement of n_4 relative to n_3 moves the centre of the ellipse outward allowing points to be included that would not be found suitable if n_3 was the centre of the ellipse.

5.3 Experiments

5.3.1 Setup

We have performed experiments to compare the methods described in the previous section. For this purpose, we randomly selected 200 slices from a set of scans of 23 patients. Each slice is centred at the stent, and tilted so that the slice is approximately perpendicular to the stent's centreline. In cases where two "legs" of the stent were visible, we randomly selected either the left or the right. All slices were then stored in a database file.

We build a MATLAB application to annotate the slices by hand. With this tool, we were able to load a database of slices, and select points on each slice through which a circular spine (see 5.2.1.3) was drawn to show the contour of the stent. An experienced radiologist annotated the 200 slices. We used the annotation data as our "gold standard" to compare the performance of the methods. Additionally, we created a training set of 50 slices which we used to determine the best parameter values of the methods. All the values of parameters and thresholds mentioned in the previous section have been determined using the trainingset.

The algorithms were applied to the slices in the database and compared to the annotation data. Both the annotated data and the algorithms produced a set of points on the contour of the stent. A circular spline was drawn through the points to make the contours continuous, which enabled us to calculate the area of the segmented stent, as well as compare the areas of the calculated contour and the annotated contour. We used the following performance measure:

$$r = \frac{A \cap M}{A \cup M}, \tag{5.4}$$

where A is the annotated area and M is the area that resulted from the selected method. We have calculated the ratio r of each annotated slice in the database and so produced an array of ratios.

5.3.2 Results and Discussion

In Figure 5.9a the histograms of all 12 methods are shown to illustrate the distribution of r, the performance. From the overall shape it can be seen that the performance is mainly above 85%. In Figure 5.9b an example result is shown against the annotation for a case with $r = 93\%$. It is clear that the result is good, which indicates that the (7%) error is due to inaccuracy of the annotation.

The error which we want to measure is not the accuracy of the points—the algorithm would always be more accurate than an annotator—but the *selection* of points. Because such errors are relatively rare, we chose not to simply average the array of ratios to compare the methods. Instead, we take the worst 30% and average it. In this way the difference in performance between the methods is much more clear.

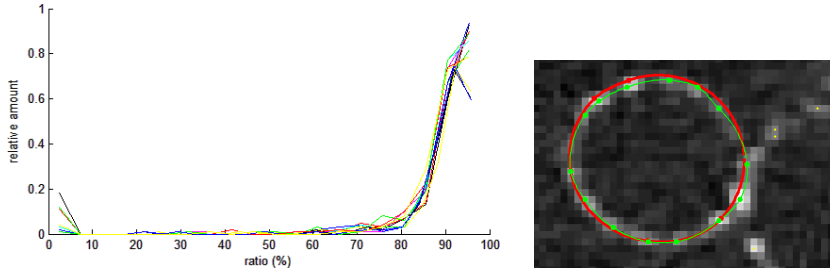

(a) (b)

Figure 5.9: **a:** Histogram illustrating the distribution of the ratios. **b:** An example result. The thick red line is the annotation. The green thin line the result for Gaussian curvature with chopstick.

In Figure 5.10 the resulting performances are shown in a bar plot. We will first discuss the results of the point detection methods and then the results for the clustering methods.

5.3.2.1 Point Detection Methods

It is clear that the performance of static thresholding is the worst. Although dynamic thresholding significantly improves the results, both Hessian based methods perform even better. It appears that the pointy nature of the structures on the contour of the stent is a very characteristic feature. We have also seen in experiments that the amount of spurious points on bone tissue or calcifications is significantly reduced for the Laplacian and Gaussian curvature compared to the two thresholding methods. We suspect that the characteristic property of the criterion is the cause of this.

Initially, we also considered the first Eigenvalue of the Hessian matrix as a measure, as it is the second order derivative in the direction where it it is largest. However, experiments revealed at an early stage that the other two Hessian based methods outperform this measure. It appears the Gaussian curvature detection method performs best. Although the Laplacian scores better in combination with the chopstick clustering method, the difference is so small that the significance is probably neglectable. It does suggest, however, that in combination with chopstick, Laplacian point detection works as well as Gaussian curvature point detection. It is interesting to see that the difference between the two Hessian based methods is so pronounced when combined with snake clustering.

5.3.2.2 Clustering Methods

It is clear that the fitting method performs poorly compared to the other methods. This suggests that the assumption that the stent's contour is shaped like an ellipse is insufficient. There are, however, specific situations (as in the example of Figure 5.5) where fitting succeeds and the other clustering methods fail.

For each point detection method, the snake's performance is lower than that of chopstick. Although the difference is neglectable in combination with Gaussian curvature point detection, the difference is larger in combination with the other point detection methods. In combination with the static threshold the snake method is

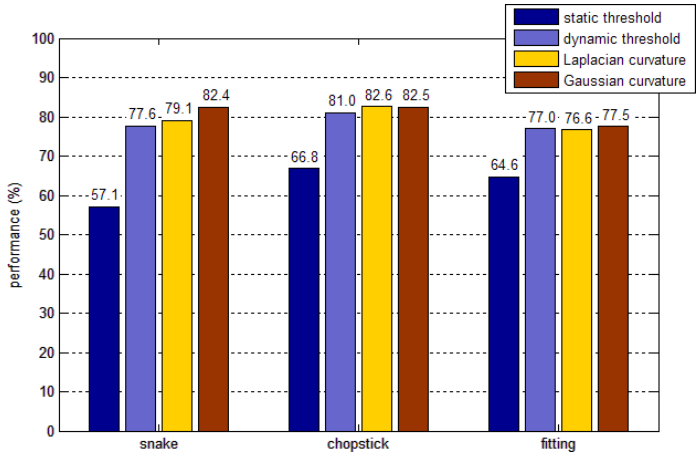

Figure 5.10: Bar plot illustrating the performance of all combinations of methods

even outperformed by the fitting method. We suspect that although the snake *can* perform good, it is sensitive for spurious points and missing points. It could be that the performance can improve if more care is taken to tune all parameters, but as this is very time-consuming, we have not been able to do that.

We propose the combination of Gaussian curvature and the chopstick method, but remark that Laplacian point detection with chopstick clustering and Gaussian curvature point detection with snake clustering perform equally well.

5.4 Conclusions and Future work

We have proposed to segment stents in ECG-gated CTA data by detecting the stent in slices perpendicular to its centreline, by first detecting points and then select the points belonging to the stent by means of a clustering method.

We have discussed four point detection methods and three clustering methods and evaluated the performance of all combinations with an experiment. We have proposed a simple yet effective method for streak artifact invariance which is applied in the four point detection methods. Contrast fluids sometimes cause spurious points, but the clustering algorithms are all to some extend robust for such points. The proposed method is to detect points using the Gaussian curvature of the image, and to cluster the points using the proposed chopstick algorithm.

In future research we plan to use the proposed method and extend our approach to segment the stent in 3D by following the centreline of the stent. The set of segmented contour points can be tracked in the volumes that represent the different phases of the cardiac cycle and will enables us to measure the surface and diameter change during the cardiac cycle.

Additionally, we plan to research the possibilities in segmenting the stent directly from the 3D data. In recent years, vessels have successfully been segmented [43] using 3D differential geometry; the wires of a stent's metallic framework have similar geometrical structure to a narrow vessel. Therefore we expect that it is possible to segment the stent using a similar method.

6

A comparison of two methods to segment stent grafts in CT data

This chapter will be published in slightly modified form at the SPIE Medical Imaging conference in San Diego, USA (2012).

Abstract

This chapter presents the tracking part of the approach to segment the stent in 2D slices sampled orthogonal to its centerline (Chapter 5). Further, the method is compared with the approach based on the MCP method presented in Chapter 8.

Using annotated reference data both methods were evaluated in an experiment. The results show that the centerline-based method and the MCP-based method have an accuracy of approximately 65% and 92%, respectively. The difference in accuracy can be explained by the fact that the centerline method makes assumptions about the topology of the stent which do not always hold in practice. This causes difficulties that are hard and sometimes impossible to overcome. In contrast, the MCP-based method works directly in 3D and is capable of segmenting a large variety of stent shapes and stent types.

6.1 Introduction

6.1.1 Purpose

Endovascular aneurysm repair (EVAR) is an established technique, which uses stent grafts to treat aortic aneurysms in patients at risk of aneurysm rupture [27, 53, 125]. The long-term durability of these stent grafts is affected by the applied stresses and hemodynamic forces, which may be reflected by the movements of the stent graft itself during the cardiac cycle. Late stent graft failure is a serious complication in endovascular repair of aortic aneurysms [16, 28, 54, 80, 88, 100], such as the formation of endoleaks (blood flow into the aneurysm sac) that can result in aneurysm expansion and rupture [81, 82, 106]. Better understanding of the motion characteristics of stent grafts will likely be beneficial for designing future devices. In addition, these data can be valuable in predicting stent graft failure in individual patients [74, 75].

Applying ECG-gated CTA [45] provides three dimensional datasets at different phases of the cardiac cycle. This allows 4D visualization of the scanned object and enables the investigation of its temporal behavior. ECG-gated CT has been used to study the motions of aneurysms [49, 91, 119] and stent grafts [74, 108]. In recent work [62] it was found that ECG gating is a suitable technique for studying the expected motions in the stent graft and vessel wall in abdominal aortic aneurysm (AAA).

A model that enables capturing material properties and high level knowledge about the stent graft would be a valuable tool to gain more insight into the stent's in vivo behavior [74]. Such a model can also help in performing more reliable (fluid dynamics) simulations, which is important for improving current stent designs [17, 68].

6.1.2 Contribution

In [65] we proposed a method to segment the stent graft in slices sampled orthogonal to the stent's centerline. In the current work we complete this approach by applying the method recursively while tracking along the centerline of the stent. In [66] we proposed a method to segment the stent graft using a more direct approach by applying the minimum cost path method in 3D.

In the current work we discuss the differences between the two methods. In addition, we performed experiments in which the produced geometric models were compared with a reference model annotated by three experts.

6.2 Method

6.2.1 Modeling the stent

The result of both algorithms is a geometric model that represents the wire frame of the stent as an undirected graph, with nodes placed at the corners and crossings of the frame, and the edges between the nodes representing the wires (Figure 6.1). This kind of model (also known as a spatial graph [6]) can be applied to different stent types, and represents the topology of the stent's frame in a concise and natural way. In the current work we only focus on the Zenith stent graft (Cook, Bloomington, USA), which has a zigzag pattern (Figure 6.2a).

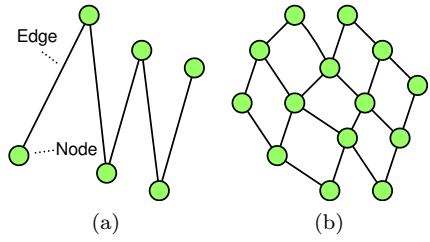

Figure 6.1: Example graphs that describe a geometric model of a stent frame. The edges between the nodes represent the physical wire frame of the stent. Nodes are placed at corners (a) and crossings (b).

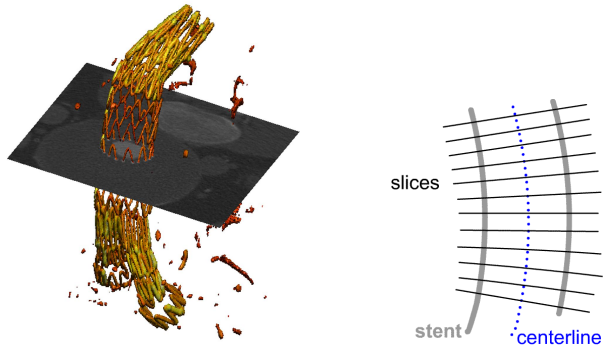

Figure 6.2: Illustration of how the stent can be approximated as a series of stacked contours. (a) shows a volume rendering of the stent and a slice orthogonal to its centerline. (b) shows a schematic illustration of contours perpendicular to the stent's centerline.

6.2.2 Segmentation via centerline tracking

A stent has a tubular structure, sometimes with branches, and can be approximated by a series of stacked contours which are orthogonal to the centerline (Figure 6.2). In the 2D images sampled perpendicular to its centerline, regions with high CT-values (typically above 500 HU) exist where the metallic frame of the stent penetrates the image. These regions have high CT-values and—due to their "pointy" structure—well suited for point detection.

The approach to segment the stent in these 2D images is to first detect a set of interest points, after which a clustering algorithm is applied to find the points that are on the wire of the stent. This process is then repeated in an iterative fashion, while tracking along the centerline of the stent. At the end of this process, a 3D geometric model of the stent is obtained.

6.2.2.1 Segmenting stent points in 2D slices

The method for segmenting the stent in 2D slices was first proposed in [65]. First, the pointy structures in the image are improved by taking the product of Eigenvalues.

This measure, also known as the Gaussian curvature, can be expressed using image derivatives: $\frac{\partial^2 L}{\partial x^2} \cdot \frac{\partial^2 L}{\partial y^2}$, where L is the 2D image. A threshold is applied to the resulting image to find a set of points. Next, the stent points are selected by means of a clustering algorithm. A custom method was developed that uses a virtual stick to select the stent points in an iterative fashion. The result of the clustering method is a set of points that represent the contour of the stent. By fitting a circle on these points, an estimate of the radius and center position are obtained, which are used during centerline tracking.

6.2.2.2 Centerline estimation

An initial centerline estimate is obtained by letting the user annotate two seed points. Hereafter, the centerline estimate is updated using the center positions obtained by the 2D stent point segmentation algorithm. Unfortunately, this center position is a rather coarse estimate because it depends on the contour points only, and can differ a few mm between subsequent slices. Therefore a global centerline tensor is maintained, which is updated with the center positions using a predetermined weight factor. This causes the centerline to be rather rigid, which is why extra seed points are sometimes necessary to handle sharp bends in the stent.

6.2.2.3 Centerline tracking

To obtain an initial set of contour points, the 2D segmentation algorithm is applied to a slice sampled perpendicular to the centerline. From the number of detected contour points, the expected number of corners in the current "ring" is calculated (Figure 6.2a).

A new slice is sampled a small distance further along the centerline. The new found center position is used to update the centerline estimate, and the detected contour points are processed (section 6.2.2.5). If the new radius is smaller than 75% of the current radius, there is potentially a bifurcation (section 6.2.2.4).

This process is repeated until all expected corner points are detected. Next, the same process is performed, but in the opposite direction. When finished, a single ring of the stent is segmented (Figure 6.3b).

Next, the algorithm "jumps" from the upper nodes in the direction of the centerline estimate, over a distance d_{ring}. If this new position is inside the stent, a new ring is segmented. In this fashion all rings of the stent are segmented until a bifurcation or the end of the stent is reached. In case of a bifurcation, both legs are processed individually.

6.2.2.4 Bifurcations

A potential bifurcation is found when the new found radius is smaller than 75% of the previous radius. To verify that that there is indeed a bifurcation, the radius calculated a certain distance further along the centerline should be smaller than 75% as well.

The detected center position is located in one of the branches. The center of the other branch can be estimated by mirroring the center position with respect to the centerline. These two points are used as seed points to segment the individual branches.

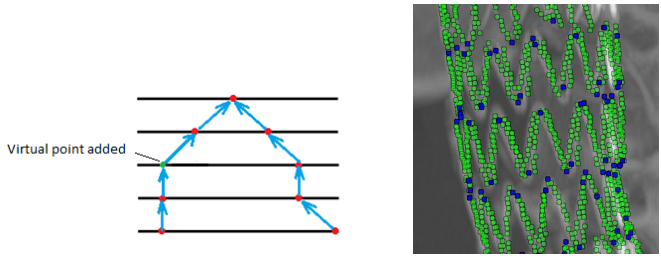

(a) (b)

Figure 6.3: Illustration of how the contour points are detected (a), and an example result (b). In (b) one can observe the corner points which will form the nodes of the geometric model.

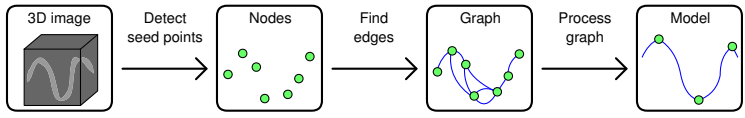

Figure 6.4: Flowchart illustrating the three processing steps to extract a geometrical model from the CT data.

6.2.2.5 Stent modeling

In each slice a set of potential stent points is found. By connecting points in consecutive slices, a model of the stent can be realized. For every point in the new slice the distances to all points in the previous slice are calculated. If this distance is less than a d_{point} amount of pixels, a connection is made (Figure 6.3a). As the data is relatively noisy, the point detection algorithm can sometimes miss points. To deal with this, the concept of virtual points is introduced. When a stent point is not present in the new slice, a virtual point is created instead. Adding virtual points can be done multiple times, until a real point is found at the approximate location. If this does not happen in n_{virtual} slices, the point is discarded as an artifact.

In cases where a single point in the new slice is connected to two points in the previous slice, a corner point is detected (Figure 6.3a). This point is marked as a node (Figure 6.3b) and is used to build the geometric model. After processing a ring, a node is placed at the end of any string of points that is a "loose end". In this way, corner points can be detected even if one of its "legs" was not detected.

6.2.3 Segmentation via the minimum cost path method

The second method segments the stent graft in 3D by finding the optimal paths between a large set of automatically detected seed points. This method was first proposed in [66]. The method can be divided into three steps, which are illustrated in the flow chart in Figure 6.4. For completeness, the method is briefly discussed in the following subsection.

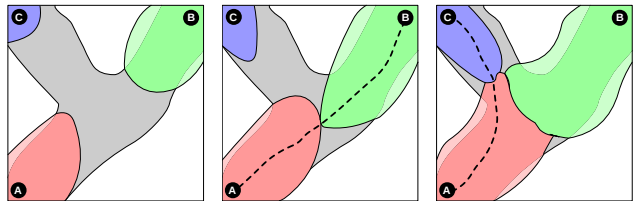

(a) (b) (c)

Figure 6.5: Illustration of three meeting fronts in the MCP algorithm. The black circular shapes indicate seed points A, B and C. In (a) the fronts do not yet meet. In (b) front A meets front B, and the path is traced. A few iterations later (c) a third front meets with the first, connecting seed points A and C.

6.2.3.1 Detection of seed points

In the first step, a set of seed points is found by searching the volume for voxels subject to three criteria: 1) The voxel intensity must be a local maximum. 2) The voxel intensity must be higher than a predefined threshold value. 3) The voxel must have a direct neighbor with an intensity also above this threshold value.

6.2.3.2 Finding the optimal paths

In the second step, the seed points are connected using a modified version of the minimum cost path (MCP) method. The MCP method can be used for segmentation of vessels and other structures (e.g. [24, 30, 38, 47, 55, 99, 121]). The advantages of this method are that it can be implemented in a computationally efficient way, and that it can easily be modified to make it more suitable for a specific problem, see for example [60] and [24].

To use the MCP method for stent segmentation, it is modified such that the fronts evolve from all the seed points found in the seed point detection step. Connections between the nodes are detected when two fronts collide, and the paths between the points are found using a backtrace map that is maintained during the evolution of the front.

The result of the MCP algorithm is a graph consisting of nodes (the seed points) connected by edges. Each edge is associated with a path of voxels connecting one node to another. However, many of these edges are false edges and have to be removed.

6.2.3.3 Graph processing

In the third step, the false edges are removed using graph processing techniques. For this purpose, two scalar values are associated with each edge. The first is α, the maximum cumulative cost on the path. It represents the *weakness* (i.e. inverse strength) of the edge. This value is used to establish the order of the edges; a stronger connection (lower α) is preferred over a weaker one. The second scalar value is β, the minimum intensity (the CT-value in Hounsfield Units) on the path. Due to the definition of CT-values (-1000 representing air and 0 representing water) this value has a physical meaning and represents the *quality* of the edge; it is used to determine whether an edge should be removed or not.

The processing of the graph occurs in multiple different passes. Firstly, weak edges are removed based on the expected number of edges for each node, which is two for the Zenith stent graft. The weakness value α is used to establish the weakest edges to consider for removal, and the quality measure β is used to determine whether an edge should be removed. Secondly, a clean-up pass is performed to remove redundant edges; an edge is found redundant if there is a path of one or two stronger (i.e. lower α) edges that connect the same nodes. Thirdly, corners are detected in the wire, and nodes are placed at the positions that have the highest curvature. Hereafter the graph is cleaned up again. Fourthly, crossings are detected and nodes are added to represent them. Finally, after a final clean-up step, all the paths are smoothed.

6.3 Experiments and results

To quantify the accuracy of the two algorithms, the geometric models obtained by the algorithms were compared with annotated reference data. The segmentation of the stent contour points in 2D slices was already quantitatively studied in earlier work [65]. For the MCP-based method, the performance was already studied extensively in [66].

6.3.1 Materials

ECG-gated CT data of 6 different patients was obtained using a Siemens Somatom 64-slice CT scanner (Siemens Medical Solutions, Erlangen, Germany). A rotation time of 0.37 s, a pitch of 0.34, and $2 \times 32 \times 0.6$ mm^2 collimation were used. The effective tube current (per rotation) time product was 180 mAs and the tube voltage was 120 kVp. Each volume was reconstructed using the B36f reconstruction filter and resulted in approximately 300 slices of 512×512 voxels. The slices (with a thickness of 2 mm) were spaced 1 mm apart, and the resolution in the xy plane was approximately 0.5 mm. Retrospective gating was applied to obtain 10 (equal distant) cardiac phases. The seven most diastolic phases were averaged to obtain a single volume for the segmentation algorithms to operate on [66]. The resulting data were manually cropped to $256 \times 256 \times 256$ voxels to reduce the memory requirements.

For the purpose of creating a reference graph, two dimensional images of unfolded pieces of stent were obtained via radial raycasting. Virtual cylinders were manually aligned with the stent's centerline, and given a diameter slightly larger than that of the stent. Rays were projected (using maximum intensity projection) from the axis onto the face of the cylinder to obtain an unfolded image of the stent (Figure 6.6a). Since a bifurcation cannot be unfolded in this manner, only tubular sections without bifurcations were used. This resulted in 12 unfolded stent images.

An annotation tool was developed which could be used to annotate a graph in the 2D image of the unfolded stent. Using the mouse, nodes could be created, repositioned and removed, and edges between the nodes could be defined. Three trained observers (an experienced radiologist, an experienced vascular surgeon and a computer scientist) annotated the stent's frame using the annotation tool. Because the (x, y, z) location in the original volume was known for each pixel in the unfolded images, a three dimensional graph could be obtained from the experts' 2D annotations (Figure 6.6b), which served as a reference to evaluate the two algorithms. In earlier work [66] it was found that the consensus between the different annotators is as high as 98%, which

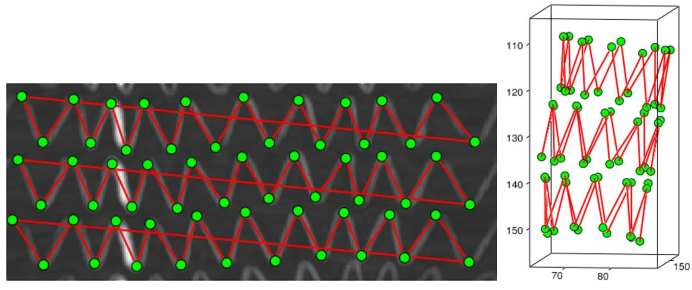

(a) 2D image and annotations (b) Reference graph

Figure 6.6: Illustration of an example image used for the manual annotations (a), obtained using radial raycasting, and the resulting 3D reference graph (b).

suggests that the reference data is a good approximation of the real topology of the stent.

To quantify the correspondence between two graphs, the edit distance [15] was used; the score is based on the number of operations to transform one graph into the other. The final score is expressed using the Dice Similarity Coefficient (DSC) [26]:

$$DSC = \frac{2 \cdot N_{match}}{(N_{match} + N_{miss}) + (N_{match} + N_{wrong})} \cdot 100\% \tag{6.1}$$

where N_{match} represents the number of true positives, and N_{wrong} and N_{miss} are the number of false positives and false negatives, respectively.

6.3.2 Results

Multiple experiments were performed using the annotated reference data. The reference data was divided into a training set containing 4 stent pieces (186 edges in total) obtained from 4 different patients, and a test set containing 8 stent pieces (481 edges in total) obtained from 6 different patients. The training set was used for parameter tuning and the test set for measuring the accuracy of the algorithm. The final results are shown in Figure 6.7. Example results for both methods are shown in Figure 6.8.

6.4 Discussion

6.4.1 Training

For the results of tuning the parameter values of the MCP-based method we refer to [66]. For the centerline method, $d_{\text{point}} = 2.0$ was found to provide the best results. The value of d_{ring} was found to have little influence on the result; as long as the algorithm lands somewhere inside the next ring it is able to process the ring.

6.4.2 Accuracy

The results in Figure 6.7 clearly show that the MCP-based method outperforms the centerline-based method. This can also be seen from visual inspection of the results

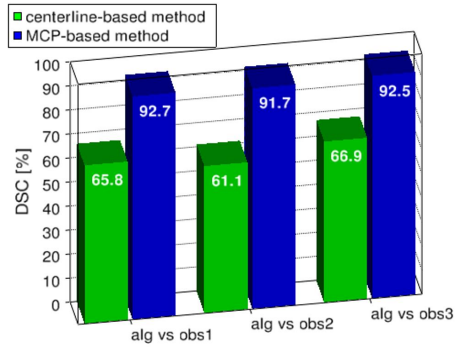

Figure 6.7: Illustration of the accuracy of the two algorithms, measured using the 8 stent pieces from the test set. ("alg" stands for the proposed algorithm, "obs" stands for observer.)

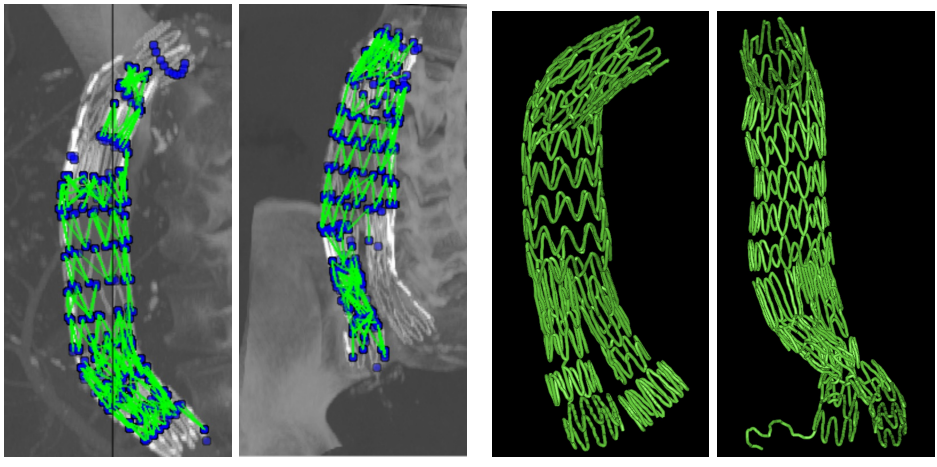

(a) Centerline-based method (b) MCP-based method

Figure 6.8: Illustration showing two example results for both methods.

in Figure 6.8.

6.4.3 Advantageous and limitations of the centerline method

An advantage of the centerline-based approach is that part of the algorithm is 2D, which makes visualization and algorithm design easier (for that part of the algorithm). However, the model of the stent as a series of stacked contours does not always hold, which makes further processing much harder, and is the root of several fundamental limitations.

For example, in handling bifurcations, it was found that although the expected location of both branches is accurate, the tracking of the branches failed in many cases, because the first part of the branch is usually embedded inside the trunk of the stent, causing problems with detecting the contour points. Similarly, when two rings

of the frame of the stent overlap, these rings are not segmented correctly.

Due to the inaccurate nature of the center estimate, the weighting factor to update the global centerline tensor should be very low. This causes the centerline to be so rigid that the algorithm cannot track sharp bends in the stent, particularly in the bifurcated branches. Furthermore, the center location estimate can sometimes be incorrect, particularly when the slice is sampled somewhere in between two rings. Although absurd results can be detected and ignored, this does sometimes cause the centerline to be updated in the wrong direction.

6.4.4 Advantageous and limitations of the MCP-based method

The MCP-based method avoids these problems because it works directly in 3D. It makes no assumption about the stent shape, not even that it is tubular. Therefore the bifurcations should in principal not cause any problems, although this cannot be verified with the current experimental method. The generic nature of the algorithm also means that the method can be applied to other stent types [66].

A downside of the fact that this method makes little assumptions about the topology of the stent is that it can also cause it to produce incorrect wires that for example run via calcifications. Another downside of this method is that it will detect structures on other high intensity data, such as the spinal cord or coils. Fortunately, these structures can easily be removed by user interaction.

6.5 Conclusions

Two methods were compared for the segmentation of stent grafts in AAA. The quantitative experiments show that the centerline-based method and the MCP-based method have an accuracy of approximately 65% and 92%, respectively. The difference in accuracy can be explained by the fact that the centerline method makes assumptions about the topology of the stent which do not always hold in practice. This causes difficulties that are hard and sometimes impossible to overcome. In contrast, the MCP-based method works directly in 3D and is capable of segmenting a large variety of stent shapes.

7

Initial steps towards automatic segmentation of the wire frame of stent grafts in CT data

This chapter is published in slightly modified form at the EMBS Benelux Chapter conference in Enschede, The Netherlands (2009) [60].

Abstract

This chapter presents an approach to segment the stent in 3D by tracking along the wires of its frame. Three tracking methods are proposed and compared in an experiment. A 2D test image was created by obtaining a projection of a 3D volume containing a stent. Ten versions of this image were obtained by adding different noise realizations. Each algorithm was started at the start of each of the ten images, after which the traveled paths were compared to the known correct path to determine the performance. Additionally, the algorithms were applied to 3D CT data and visually inspected.

The results show that the minimum cost path (MCP) is very suitable for tracking short pieces of wire of the frame of the stent. However, determining crossings and corners in the wire is difficult with this tracking approach. Therefore we decided to apply the MCP method on a set of seed points to segment the wire pieces, which resulted in the approach presented in Chapter 8.

7.1 Introduction

Endovascular aortic replacement (EVAR) is an established technique, which uses stent grafts to treat aortic aneurysms in patients at risk of aneurysm rupture [12]. The long-term durability of these stent grafts is affected by the stresses and hemodynamic forces applied to them, and may be reflected by the movements of the stent graft itself during the cardiac cycle. Applying ECG-gated CTA [45, 91, 65] provides three dimensional datasets at different phases of the cardiac cycle, and thus provides insight in the stent's motion. These data can be used to calculate the forces caused by the hemodynamics [62]. However, this requires obtaining the geometric model of the stent's frame in the form of connected landmarks placed at corners and crossings.

In literature, several studies have been published on the segmentation of blood vessels. However, two-step approaches [123], which first segment the vessel using a vessel measure [43] followed by centerline extraction, are not suitable because of the gaps in the wire (caused by artifacts) and its sharp corners compared to blood vessels. Methods that fit spheres/ellipsoids to the vessel (e.g. [8]), and methods that segment the contour in slices perpendicular to the vessel centerline (e.g. [76]), will not work due to the small diameter of the wire (1 to 3 voxels) and its sharp corners. Region growing methods [13] have problems with leaks and gaps and need a second stage to find the geometry from the segmented voxels. Methods that use the minimum cost path method [33] are robust against gaps and artifacts (e.g. [30, 122]). However, these method cannot be applied in their default form since it requires the selection of a start point (and often an end point).

We propose three tracking algorithms that finds the optimal path along the stent's wire, and compare them in an experiment. The proposed methods are applicable to both 2D and 3D (anisotropic) data.

7.2 Methods

The basic assumption of the proposed methods is that once a position on the wire of the stent is known, a next position along the wire can be calculated with great confidence. Given a 3 dimensional image $V(\mathbf{x})$ with $\mathbf{x} = (x, y, z)$, the algorithm places multiple objects inside the volume, which we refer to as *walkers*. These walkers are initially positioned in a regular grid with a certain spacing in mm. Next, each walker moves to the voxel in its local environment that has the highest intensity, after which each walker will proceed autonomously through the data, while tracking the wire frame of the stent.

7.2.1 Method A: 1D growing

In each iteration, this method compares its second order neighbors (i.e. its 8 or 26 neighbors for 2D and 3D, respectively) and selects the neighbor with the highest intensity. If the intensity of the selected pixel/voxel is above a certain threshold, it will be the current position for the next iteration, otherwise the walker is discarded.

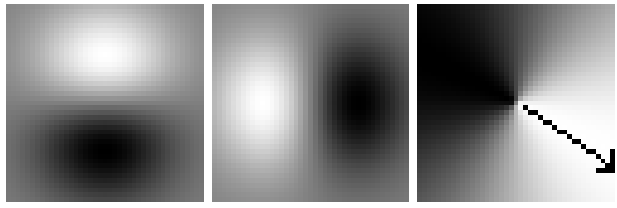

Figure 7.1: Illustration of the two 2D smoothed COM kernels and the direction-dependent weighting kernel.

7.2.2 Method B: directional center of mass

The second proposed method uses a modified version of the center-of-mass (COM) method, in which the kernels are Gaussian smoothed and multiplied by a direction-dependent weighting kernel. The Gaussian smoothing is applied to suppress the influence of voxels further away from the kernel center, and the weighting kernel is applied to maintain motion in the direction of the previous step (Figure 7.1). Since the smoothed COM kernel can be represented by Gaussian derivative kernels, this operator can be considered a gradient operator biased for a certain direction.

We describe the elements of the direction dependent weighting kernel as $w_{\mathbf{k}}$, with $\mathbf{k} = (k_x, k_y, k_z)$ the kernel coordinate relative to the center of the kernel. Given the direction of the previous step $\mathbf{d}$, the angle for each kernel element is $\alpha_{\mathbf{k}} = \arccos(\mathbf{k}/|\mathbf{k}| \cdot \mathbf{d}/|\mathbf{d}|)$. The weights of the kernel are calculated with:

$$w_{\mathbf{k}}(\alpha) = \cos(\alpha_{\mathbf{k}}) + 0.5. \tag{7.1}$$

Replacing $\alpha_{\mathbf{k}}$ gives:

$$w_{\mathbf{k}}(\mathbf{x},\mathbf{d}) = \frac{\mathbf{x}}{|\mathbf{x}|} \cdot \frac{\mathbf{d}}{|\mathbf{d}|} + 0.5. \tag{7.2}$$

Each iteration, a patch of data p is sampled, centered around the current position. For each dimension, this patch is multiplied with the corresponding Gaussian derivative kernel and with the weighting kernel to yield the new (unnormalized) direction:

$$\mathbf{d}_{new} = \left(\sum_{\forall \mathbf{k}} p_{\mathbf{k}} g_{\mathbf{k},x} w_{\mathbf{k}}, \sum_{\forall \mathbf{k}} p_{\mathbf{k}} g_{\mathbf{k},y} w_{\mathbf{k}}, \sum_{\forall \mathbf{k}} p_{\mathbf{k}} g_{\mathbf{k},z} w_{\mathbf{k}} \right), \tag{7.3}$$

where $g_{\mathbf{k},x}$, $g_{\mathbf{k},y}$ and $g_{\mathbf{k},z}$ represent the Gaussian derivative kernels for the x, y and z direction, respectively. Next, $\mathbf{d}_{new}$ is combined with the old value of $\mathbf{d}$ such that the angle between the old and the new direction is limited by a predefined value. This angle threshold gives control over the walkers' affinity to proceed in the same direction, and thus prevents the walker from "making a u-turn".

The new position is determined by selecting the voxels which are in the found direction, and choosing the one with the highest intensity.

7.2.3 Method C: minimum cost path

The third proposed algorithm is based on the minimum cost path algorithm [33, 122], which is applied in a step-by-step tracking approach. Given the speed (representing

Algorithm 7.1 The minimum cost path algorithm

```
01 init the time map t_x = ∞ ∀x
02 init the frozen map f_x = 0 ∀x
03 init t_x = 0 for the start nodes
04 insert the start nodes in list L
05 while not isempty(L):
06   take node m from L with min.  t_m
07   set f_m = 1
08   for n in neighborsOf(m):
09     if f_n == 0:
11       d = eucledianDist(n,m)
12       t = t_m + d/s_n
13       if t < t_n:
14         t_n = t
15         d_n = d_m + d
16         if n not in L:
17           insert n in L
```

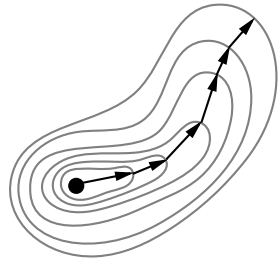

Figure 7.2: Illustration of the proposed minimum cost path approach. The front (the voxels in L) moves outward until the cumulative distance of the path is larger than d_{min}.

the inverse cost function) in each voxel $s_{\mathbf{x}}$, the algorithm iteratively calculates the time $t_{\mathbf{x}}$ it takes to travel from the start node to any other node in the image (see algorithm 7.1). The fastest (minimum cost) path can be found to any point $\mathbf{x}$ by backtracking in the resulting time map.

We modified the algorithm such that it stops as soon as a distance $d_{\mathbf{x}}$ larger than d_{min} is encountered (Figure 7.2). Since the algorithm selects the node with the minimum cumulative cost ($t_{\mathbf{x}}$) at each iteration, the first node $\mathbf{x}$ for which $d_{\mathbf{x}} > d_{min}$ is the node that is "traveled easiest to", and therefore has the highest probability of being positioned on the same wire as the previous node.

To be robust against artifacts, only the first half of the path from the start node to $\mathbf{x}$ is used. This path is appended to the total walked path and its last node will be the start node for the next iteration. At the start of the next iteration the last 20 nodes in the total walked path are set frozen to prevent backtracking.

The speed $s_{\mathbf{x}}$ is calculated with:

$$s_{\mathbf{x}} = 2^{V(\mathbf{x})/\delta}, \tag{7.4}$$

where δ is a scalar number. Using this speed map has the effect that the path

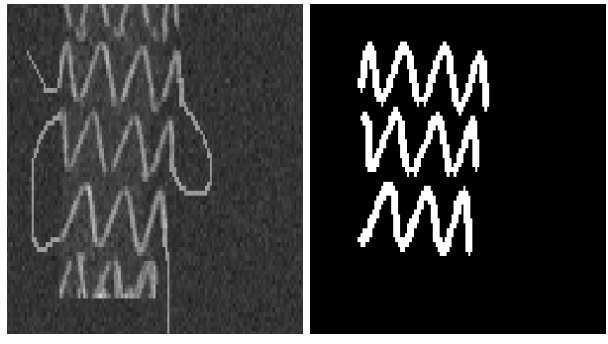

(a) (b)

Figure 7.3: Illustration of the 2D test image (a) and the used mask to calculate performance(b).

through a single voxel with a CT-value of a, is equally fast as another path through two voxels with CT-values of $a + \delta$. In other words, the optimal path has a strong affinity for high intensity voxels, rather than taking the shortest route. This prevents the algorithm to cut the sharp corners present in our data. In our method we used a value of $\delta = 100$.

7.3 Experiments

A 2D test image (Figure 7.3a) was created by obtaining a projection of a 3D volume containing a stent. The image was modified to connect the three parts of the stents' frame and thus create a single path. Ten versions of this image were obtained by adding different noise realizations (Gaussian noise with a standard deviation of 75 CT-values). A walker was placed to walk down the path of the ten images, after which the traveled paths were compared to the known correct path (Figure 7.1b) to determine the performance of each method. Beforehand, the same setup was used (but with different noise realizations as in the experiment) to choice the optimal parameters for each method.

Additionally, the algorithms were applied to 3D clinical data using multiple walkers, and the result was visually inspected.

7.4 Results and Discussion

From the average correspondence of the walked paths with the correct path (Table 7.1) it can be seen that the method A performed poorly, while method C performed excellent. The traveled paths for each method (Figure 7.4) confirm these findings. Visual inspection of the method applied to 3D data (Figure 7.5) shows the same trend.

During the process of the experiments, it was found that the performance of method B is sensitive to the chosen parameters (such as the scale of the kernels and the angle limit). In contrast, method C was very robust for the choice of its parameter

Method:	A	B	C
Performance:	0.32	0.88	1.00

Table 7.1: The quantitative results of the experiments.

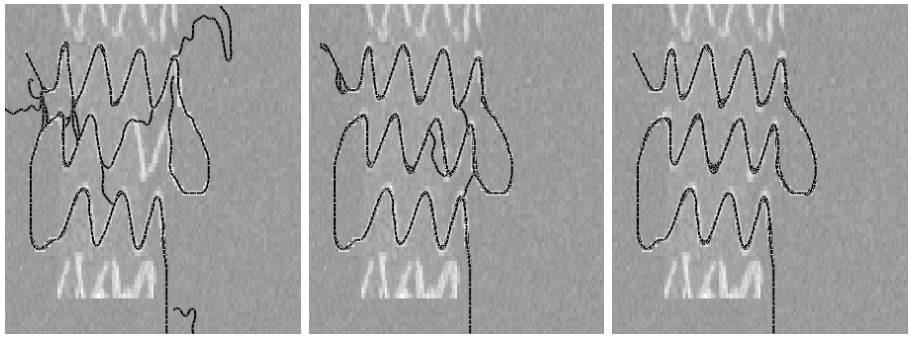

(a) (b) (c)

Figure 7.4: Illustration of the walked paths of the three methods for the 10 noise realizations.

values (d_{min} and δ).

We note that method C can be seen as a generalization of method A; instead of selecting the next position from the direct neighbors, the most probable voxel at a certain distance from the current position is selected. The concept of incorporating a larger neighborhood in determining the most probable next position is also used in method B. However, method B is sensitive for nearby structures (e.g. other wires) and has problems with sharp corners.

In the current method, the walkers are unaware of each other. By connecting the walked paths when two walkers meet, the walkers should be able to work together to perform the segmentation task.

7.5 Conclusions and Future Work

We have described and compared three methods to track the wire frame of the stent graft in 3D CT data and 2D projections thereof. The best performing method is based on the minimum cost path method and is capable of obtaining near-perfect segmentation of the stent. The method is robust to unforeseen situations and variation in its parameters.

Future research will focus on establishing a geometrical model by determining the corner points and the crossings from the results of the proposed method.

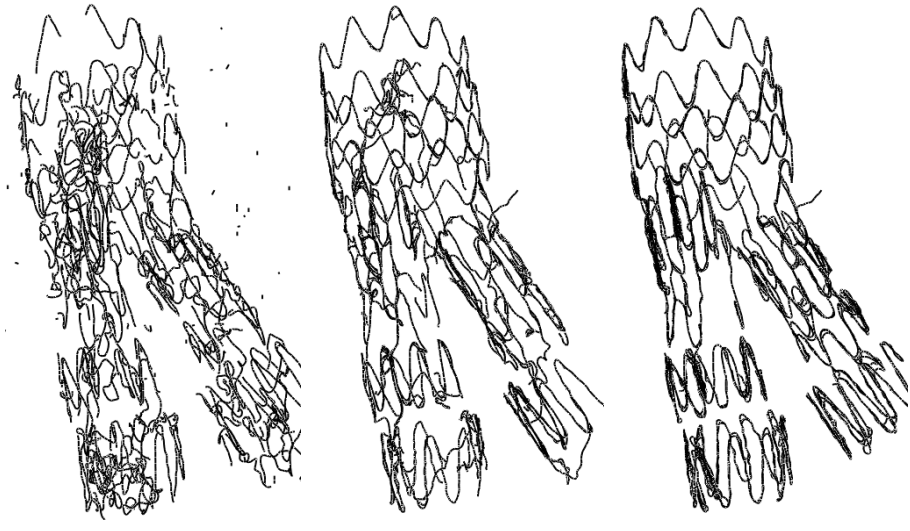

(a) (b) (c)

Figure 7.5: Illustration of the found paths of multiple walkers in a 3D volume.

8

Automatic segmentation of the wire frame of stent grafts from CT data

This chapter is published in slightly modified form in Medical Image Analysis (2011) [66].

Abstract

This chapter presents a stent segmentation algorithm that is based on the minimum cost path (MCP) method and works directly in 3D. It consists of three steps: the detection of seed points, finding the connections between these points to produce a graph, and graph processing to obtain the final geometric model in the form of an undirected graph.

Using annotated reference data, the method was optimized and its accuracy was evaluated. The experiments were performed using data containing the AneuRx and Zenith stent grafts. The algorithm is robust for noise and small variations in the used parameter values, does not require much memory according to modern standards, and is fast enough to be used in a clinical setting (65 and 30 seconds for the two stent types, respectively). Further, it is shown that the resulting graphs have a 95% (AneuRx) and 92% (Zenith) correspondence with the annotated data. The geometric model produced by the algorithm allows incorporation of high level information and material properties. This enables us to study the in vivo motions and forces that act on the frame of the stent.

8.1 Introduction

8.1.1 Clinical context

Endovascular aortic replacement (EVAR) is an established technique, which uses stent grafts to treat aortic aneurysms in patients at risk of aneurysm rupture [27, 53, 125]. The long-term durability of these stent grafts is affected by the stresses and hemodynamic forces applied to them, and may be reflected by the movements of the stent graft itself during the cardiac cycle. Late stent graft failure is a serious complication in endovascular repair of aortic aneurysms [16, 28, 54, 80, 88, 100], such as the formation of endoleaks (blood flow into the aneurysm sac) that can result in aneurysm expansion and rupture [81, 82, 106]. Better understanding of the motion characteristics of stent grafts will likely be beneficial for designing future devices. In addition, these data can be valuable in predicting stent graft failure in individual patients [74, 75].

Applying ECG-gated CTA [45] provides three dimensional datasets at different phases of the cardiac cycle. This allows 4D visualization of the scanned object and enables the investigation of its temporal behavior. ECG-gated CT has been used to study the motions of aneurysms [49, 91, 119] and stent grafts [74, 108]. In recent work [62] it was found that ECG gating is a suitable technique for studying the expected motions in the stent graft and vessel wall in abdominal aortic aneurysm (AAA).

Most studies on the motion of stent grafts focus on measuring the stent's diameter changes [51] or finding the motion for a sparse set of points on the stent [74]. A model that enables capturing material properties and high level knowledge about the stent graft would be a valuable tool to gain more insight into the stent's in vivo behavior [74]. Such a model can also help in performing more reliable (fluid dynamics) simulations, which is important for improving current stent designs [17, 68]. To the best of our knowledge such a model has not yet been proposed in literature.

One of the major downsides of CT in general, is the exposure of the patient to ionizing radiation, which can have negative effects on the long term health of the patient [98, 39]. The dose should therefore be kept as low as possible. However, this results in higher noise levels and more image artifacts, which can be a problem for automatic image analysis algorithms that often need high quality data to operate. Algorithms that can perform their task on low dose data can therefore contribute to better patient safety.

8.1.2 Contributions

In this work we take an important step towards evaluating the motion of stent grafts in vivo: the segmentation of the stent graft and its representation using a geometric model.

We propose a geometric model that represents the wire frame of the stent as an undirected graph, with nodes placed at the corners and crossings of the frame, and the edges between the nodes representing the wires (Figure 8.1). This model represents the topology of the stent's frame in a concise and natural manner.

The main goal of the present work is to present a method to obtain the geometric model of the stent's wire frame from volumetric CT data. The method is fast and robust for high levels of noise, such that it is applicable to CT data obtained with relatively low dose. To evaluate the quality of the produced geometric model, we

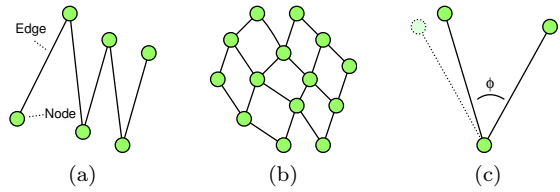

Figure 8.1: Example graphs that describe a geometric model of the stent's frame. The edges between the nodes represent the physical wire frame of the stent. Nodes are placed at corners (a) and crossings (b), which makes it possible to model different stent types. The (change of) angle ϕ when motion is applied to the model can be used to estimate the force present in the node (c).

performed experiments in which the produced model was compared with a reference model annotated by three experts.

8.1.3 Outlook

The geometric models produced by the proposed algorithm enable future studies on the motion of stent grafts in vivo; the motion field between the cardiac phases (which can be obtained by image registration) capture the motion of the stent and can be used to apply motion to the model. This will give quantitative information about the motion of stent grafts. Subsequently, by incorporating material properties such as stiffness, and by calculating the change of the angle between two edges, the forces present at the stent's frame can be estimated (Figure 8.1c). Both parameters will add valuable information to the analysis of stent grafts in vivo.

To be useful in future motion analysis, it is of importance that the model precisely represents the topology of the wire frame of the stent graft; if the model contains erroneous or missing nodes or edges, the force calculations can produce invalid results.

While the proposed method can be applied to non-gated CT data, in the present work ECG-gated data was used because of our intentions to study stent graft motion in the future.

8.1.4 Previous work

Few results have been published for the segmentation of the wire frame of stent grafts from CT data, despite the need for such a technique to evaluate large 3D datasets. In preliminary work we segmented the wire frame of the stent in 2D slices [65]. In a recent work by [75], the stent's frame (in the thoracic aorta) is found using a statistical model. They report that a threshold at 2000 Hounsfield Units is sufficient to obtain a connected set of points belonging to the stent. Unfortunately, our data contains too much noise (compared to the stent's intensities) for such a method. This can be due to differences in the applied dose (which is not not reported in [75]) or in the material of the stent's frame.

Several studies have been published on the segmentation of blood vessels in 3D, which have similarities with the wires of the frame of the stent (Figure 8.2) and may therefore be of interest (see [79] and [59] for an overview of vessel segmentation methods). One common method is the two-step approach [44, 57, 84, 123], which first

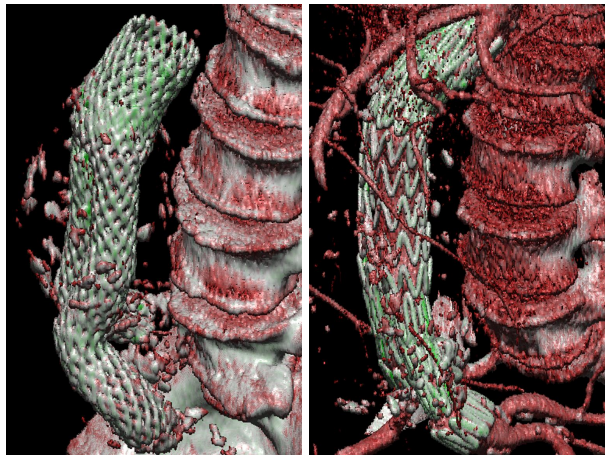

(a) AneuRx (b) Zenith

Figure 8.2: Illustration of iso-surface renderings from CT data of two different types of stent grafts.

segments the vessel using a vessel measure [43] followed by centerline extraction. Unfortunately, this method is not suitable for our research question because of gaps that can occur in the wire due to artifacts, and the sharp corners of the wire which blood vessels do not posses (Figure 8.2). Methods that fit a series of spheres or ellipsoids to the vessel [8, 126], and methods that segment the contour in slices perpendicular to the vessel centerline [50, 76] assume a solid vessel with a diameter of several voxels. Due to the small diameter of the stent's wire (1 to 3 voxels) and its sharp corners these methods are also not suitable. Region growing methods [13] have problems with leaks and gaps and need a second stage to find the geometry from the segmented voxels.

In recent work [60] we have proposed three tracking algorithms to find the optimal path along the wire of the stent, and compared them in an experiment. It was found that the method based on the Minimum Cost Path (MCP) method produced the best results. In the current work, we use the MCP method to segment the stent in a direct manner rather than applying a tracking approach. Additionally, the current work describes the method to create a geometrical model from the results obtained by the MCP method.

8.1.5 Outline

First the algorithm to calculate the geometric model from the CT data is discussed in three steps in section 8.2. Next, in section 8.3, we discuss the experiments that were performed to determine the optimal values for the algorithm's parameters, and the experiments to evaluate the accuracy of the algorithm. Finally, the results are discussed in section 8.4.

In this work we consider two types of stent grafts which are commonly used in our institution to treat patients with AAA: the AneuRx (Medtronic, Minneapolis, USA) and the Zenith (Cook, Bloomington, USA) stent graft (Figure 8.2). The proposed

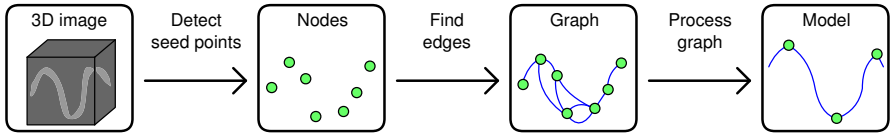

Figure 8.3: Flowchart illustrating the required processing steps to extract a geometrical model from the CT data.

method should be general enough to be capable of segmenting the wire frame of both types of stent grafts.

8.2 Methods

Given are multiple volumetric images $V(\mathbf{x}) \rightarrow \mathbb{R}$ with $\mathbf{x} = (x, y, z)$, acquired from a CT scanner using ECG-gating; each image represents a phase in the cardiac cycle. In a pre-processing step of the proposed segmentation algorithm, n_{phases} volumes are averaged to reduce the noise at the cost of increased motion artifacts. The averaged phases are taken from the end-diastole, which suffers the least from motion artifacts [31]. In section 8.3 the optimal number of phases to average is determined in an experiment.

The proposed method can be divided into three steps, which are illustrated in the flow chart in Figure 8.3. First, seed points are detected that are likely to be on the wire of the stent, resulting in a set of nodes $\mathbf{p}_i = (x, y, z)$ with $i \in [0 \ldots N-1]$ and N the number of nodes. These nodes represent a graph without edges. In the next step, the edges between the nodes are found using a modified version of the MCP method, resulting for each node $\mathbf{p}_i$ in a set of edges e_{ij}, each connecting it to another node $\mathbf{p}_j$. Finally, the resulting graph is processed by introducing new nodes and removing unwanted nodes and edges. These three steps are discussed in the following subsections.

8.2.1 Detecting the seed points

A set of seed points is found by searching V for voxels subject to three criteria:

- The voxel intensity must be a local maximum.
- The voxel intensity must be higher than a predefined threshold value v_{seed}.
- The voxel must have a direct neighbor with an intensity also above this threshold value.

Because the frame of the stent has an intensity that is (locally) constant, local maxima—and thus the resulting seed points—are distributed homogeneously along the wire frame of the stent. Since the diameter of the wire of the stent graft is much smaller than the voxels size, the partial volume effect causes the center of wire crossings (in the AneuRx stent graft) to always be a local maximum.

8.2.2 Connecting the seed points

The seed points are connected using a modified version of the minimum cost path (MCP) algorithm. The MCP method can be used for segmentation of vessels and other structures (e.g. [24, 30, 38, 47, 55, 99, 121]). It is a level set method in which a front is propagated monotonically following a (non-negative) cost function. The advantage of this method is that it can be implemented in a computationally efficient way. It is a form of Dijkstra's algorithm [33] that finds the shortest path in a graph of connected nodes. In the MCP method the nodes are represented by the pixels/voxels, and the edges are implicitly defined by the grid; each node is connected to its direct neighbors.

8.2.2.1 Theory

Given a set of seed points $P = \{\mathbf{p}_1, \mathbf{p}_2, \ldots, \mathbf{p}_n\}$ (with $\mathbf{p}_i = (x, y, z)$) and the costs in each voxel $C(\mathbf{x}) \rightarrow \mathbb{R}$, the algorithm iteratively calculates the cumulative costs $D(\mathbf{x}) \rightarrow \mathbb{R}$ it takes to travel from the seed points to any other point in the image (Algorithm 8.1). The minimum cost path from any voxel $\mathbf{x} = (x, y, z)$ to the seed point can be found by backtracking using D.

A major difference with respect to Dijkstra's algorithm is that the costs are defined at the voxels (i.e. nodes) rather than at the edges. Therefore, voxels $\mathbf{n}$ and $\mathbf{m}$ should contribute equally to the transition cost. This fact is sometimes neglected in literature, in which case line 9 in Algorithm 8.1 is replaced by $d = D(\mathbf{m}) + C(\mathbf{n})$, which resembles the update rule of Dijkstra's algorithm.

By weighting with the Euclidean distance between the voxels, the method corrects for the larger distance for diagonal voxels, and can account for the anisotropicity of the data. In applications where the cost and cumulative cost represent a distance measure, it is advantageous to consider the 8 (2D) or 26 (3D) nearest neighbors rather than only the 4 (2D) or 6 (3D) neighbors; the shape of the front is in that case octagonal rather than square, which results in smaller discretization errors.

A related method is the Fast Marching method [105] which replaces line 9 in Algorithm 8.1 with a formula to solve the Eikonal equation. It therefore remains (more) consistent with the continuous formulation of the problem, which significantly reduces the discretization errors. However, these errors are not significant in our situation since the paths are relatively short and the intensity difference between the wire and the background is relatively large. Additionally, the backtracking algorithm in the resulting cumulative cost map—which is applied relatively often in the proposed method—is considerably more easy and faster for the MCP method.

8.2.2.2 Cost function

Since the stent is made of metal, it is represented in the data $V(\mathbf{x})$ with relatively high CT-values compared to the soft tissue in the human body. Therefore we define our cost function such that an increase in CT-value results in a decrease in cost:

$$C(\mathbf{x}) = 2^{-V(\mathbf{x})/v_{aff}}, \tag{8.1}$$

where v_{aff} is a predefined value representing the algorithm's affinity for high intensities. Consider two hypothetical paths, of which one is exactly twice as long as the other but running through voxels of higher intensities. The above formula implies

Algorithm 8.1 Pseudo code for the normal Minimum Cost Path method.

```
01 function MCP(C):
02   init cumulative cost map D(x) = ∞ ∀x
03   init D(p_i) = 0 for all seed points
04   insert the seed points in list L
05   while not isempty(L):
06     take voxel m from L with minimum D
07     for n in neighborsOf(m):
08       δ = euclideanDist(n,m)
09       d = D(m) + 1/2 · δ · (C(n) + C(m))
10       if d < D(n):
11         D(n) = d
12         if n not in L:
13           insert n in L
14   return D
```

that the longer path is preferred if the intensities over the path are at least v_{aff} Hounsfield Units higher. This increased affinity for high values over a shorter path prevents the algorithm from cutting (sharp) corners.

8.2.2.3 Implementation

The MCP method can easily be modified to make it more suitable for a specific problem. For example, in [60] the algorithm executes until a certain distance from the seed point is reached, and in [24] the traveled distance of the front compared to the largest travelled distance is used as a measure to freeze the front and thus prevent leaking outside the vessel.

In the present work the MCP algorithm is modified (Algorithm 8.2) such that the fronts evolve from all the seed points $\mathbf{p}_i$ found in the seed point detection step (section 8.2.1). Connections between the nodes are detected when two fronts collide. As the fronts propagate, two maps are maintained (both maps are initialized with all values to -1). The first is the backtrack map B that specifies for each point what the previous point was. The second is the identification map I, a Voronoi diagram [5] that specifies for each point $\mathbf{x}$ from which seed point it originates: $I(\mathbf{x}) \rightarrow i$, maps point $\mathbf{x}$ to seed point $\mathbf{p}_i$. Consequently, the collision of two fronts can be detected by testing $I(\mathbf{m}) \neq I(\mathbf{n})$ and $I(\mathbf{n}) \neq -1$ (line 10 in Algorithm 8.2). If this test is positive, an edge e_{ij} between nodes $\mathbf{p}_i$ and $\mathbf{p}_j$ is created (with $i = I(\mathbf{n})$ and $j = I(\mathbf{m})$), and the path between $\mathbf{p}_i$ and $\mathbf{p}_j$ is found by backtracing the route using B (Figure 8.4b). Since it is possible that another (lower cost) path between i and j is found later (Figure 8.4c), the above check should be performed at each iteration and for each neighbor. Figure 8.4d shows how the front originating from one seed point partly replaces another front, thereby changing the identification map I and the backtrack map B. This illustrates that the paths need to be traced at the moment an edge is found.

As a stop criterion the value of $D(\mathbf{m})$ is compared to a predefined threshold d_{max}. If the threshold is reached, the cumulative cost is at such a level that we can assume all significant edges have been found.

Algorithm 8.2 Pseudo code for the proposed (modified) version of the Minimum Cost Path method.

```
01 function MCP_connect_seed_points(C):
02   init D(x) = ∞, I(x) = −1, B(x) = −1 ∀x
03   init D(p_i) = 0 for all seed points p_i
04   init B(p_i) = i for all seed points p_i
05   init I(p_i) = i for all seed points p_i
06   insert the seed points in list L
07   while not isempty(L):
08     take voxel m from L with minimum D
09     for n in neighborsOf(m):
10       if I(n) ≠ −1 and I(n) ≠ I(m):
11         makeEdge(n,m)
12       δ = EuclideanDist(n,m)
13       d = D(m) + ½ · δ · (C(n) + C(m))
14       if d < D(n):
15         D(n) = d
16         B(n) = m
17         I(n) = I(m)
18         if n not in L:
19           insert n in L
20   return edges
```

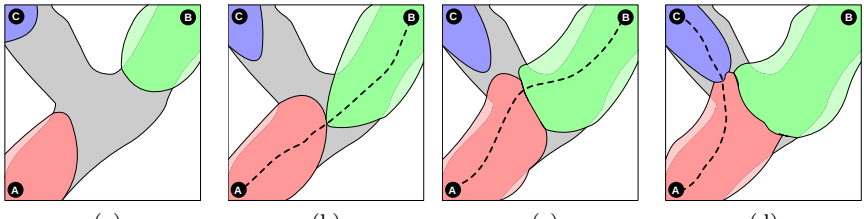

(a) (b) (c) (d)

Figure 8.4: Illustration of three meeting fronts. The black circular shapes indicate seed points A, B and C. In (a) the fronts do not yet meet. In (b) front A meets front B, and the path is traced. A few iterations later, in (c) a better path (lower D) is found between A and B. The path is traced again and updated. In (d) a third front meets with the first, connection seed points A and C.

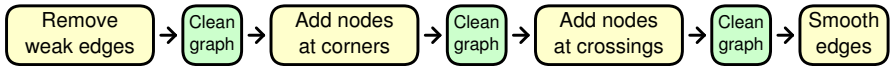

Figure 8.5: Flow chart illustrating the process of graph processing.

8.2.3 Graph processing

The result of the MCP algorithm is a graph consisting of nodes (the seed points) connected by edges. Each edge is associated with a path $P_{ij} = [\mathbf{p}_i, \ldots, \mathbf{p}_j]$ of voxels connecting one node to another. Although the paths themselves are optimal (in terms of the cost C), many of these edges are false edges, and must be removed. For this purpose, two scalar values are associated with each edge. The first is α_{ij}, the maximum cumulative cost $D(\mathbf{x})$ on the path. It represents the *weakness* (i.e. inverse strength) of the edge. This value is used to establish the order of the edges; a stronger connection (lower α) is preferred over a weaker one. The second scalar value is β_{ij}, the minimum intensity (the CT-value in Hounsfield Units) on the path. Due to the definition of CT-values (-1000 representing air and 0 representing water) this value has a physical meaning and represents the *quality* of the edge; it is used to determine whether an edge should be removed or not. The processing of the graph occurs in multiple different passes which are illustrated in Figure 8.5 and described below.

1. **Removing weak edges.** Weak edges are removed based on the expected number of edges for each node n_e, which is two for the Zenith and four for the AneuRx stent graft. The algorithm iterates over all nodes $\mathbf{p}_i$, and evaluates their edges. If a node has more than n_e edges, the weakest surplus edges are selected (using α_{ij}) and considered for removal: the edge e_{ij} is removed if β_{ij} is smaller than the predefined threshold $v_{quality}$, and if the same criteria are also met when considering the edge from the node at the other end ($\mathbf{p}_j$). If $\beta_{ij} > v_{quality}$ there is enough evidence to assume the path P_{ij} corresponds to a wire and the edge e_{ij} is considered valid (Figure 8.6a).

2. **Clean-up.** Next, redundant edges are removed: all connections s_{ij} are considered in the order of decreasing α_{ij}, starting with the weakest edge. An edge is found redundant if there is a path of one or two stronger (i.e. lower α) edges that connect the same nodes. Next, each cluster (i.e. group of interconnected nodes that is not connected to any other nodes) is removed if it consist of less than $n_{cluster}$ nodes, and each tail (i.e. string of nodes) is trimmed if it consists of less than n_{tail} nodes (Figure 8.6b).

3. **Adding nodes at corners.** The third pass starts with removing nodes from the graph that have exactly two edges. The two edges are combined to form a new edge. Next, all edges are evaluated and new nodes are inserted at the corners of the wire by searching for points on the path P_{ij} that have maximum curvature (Figure 8.6c): for each point $\mathbf{x}$ on the path, two vectors are calculated, which each span five positions on the path (corresponding to maximal 5 mm, which is less than the length of the wires). The angle ϕ between the two vectors is calculated (Figure 8.7), and the resulting angle values for all the points in the edge are smoothed with a Gaussian filter (with a standard deviation of 3 voxels) to account for the discretization. By detecting local minima in the resulting list of angles, the point of maximum curvature can be calculated.

4. Repeat pass 2. Since the graph has changed, new edges can now be marked redundant. These should be removed before proceeding to the next pass. Furthermore, pass 5 gets rid of many nodes in invalid clusters, which therefore are now eligible for removal.

5. **Adding nodes at crossings.** By tracing the paths of two edges from the same node, the algorithm determines whether the paths are partly the same. In these cases a new node is inserted at the point where the paths diverge, and the overlapping parts of the paths are replaced by a single edge (Figure 8.6e) [24]. This pass is not applied if $n_e = 4$, because in that case the nodes are already correctly positioned at the crossings and this pass would only introduce false crossings.

6. Repeat pass 2. Since the graph has changed, new edges can now be marked redundant (Figure 8.6f).

7. **Smoothing.** Finally, the paths are smoothed using a moving average filter (with a range of three voxels).

8.3 Experiments

8.3.1 Materials

ECG-gated CT data was obtained for 10 different patients (4 with AneuRx, 6 with Zenith) using a Siemens Somatom 64-slice CT scanner (Siemens Medical Solutions, Erlangen, Germany). A rotation time of 0.37 s, a pitch of 0.34, and $2 \times 32 \times 0.6$ mm^2 collimation were used. The effective tube current (per rotation) time product was 180 mAs and the tube voltage was 120 kVp. Each volume was reconstructed using the B36f reconstruction filter and resulted in approximately 300 slices of 512×512 voxels. The slices (with a thickness of 2 mm) were spaced 1 mm apart, and the resolution in the xy plane was approximately 0.5 mm. Retrospective gating was applied to obtain 10 (equal distant) cardiac phases. The resulting data were manually cropped to $256 \times 256 \times 256$ voxels to reduce the memory requirements.

For the purpose of creating a reference graph, two dimensional images of unfolded pieces of stent were obtained via radial raycasting. Virtual cylinders were manually aligned with the stent's centerline, and given a diameter slightly larger than that of the stent. Rays were projected (using maximum intensity projection) from the axis onto the face of the cylinder to obtain an unfolded image of the stent (Figure 8.8a). Since a bifurcation cannot be unfolded in this manner, only tubular sections without bifurcations were used. This resulted in 18 unfolded stent images: 6 of the AneuRx stent graft and 12 of the Zenith stent graft. In total there were approximately 700 edges for each stent type.

An annotation tool was developed which could be used to annotate a graph in the 2D image of the unfolded stent. Using the mouse, nodes could be created, repositioned and removed, and edges between the nodes could be defined. The annotator could initialize the graph with the (un-tuned) algorithm to alleviate the required work for the annotator in regions which can be considered "easy".

Three trained observers (an experienced radiologist, an experienced vascular surgeon and a computer scientist) annotated the stent's frame using the annotation tool.

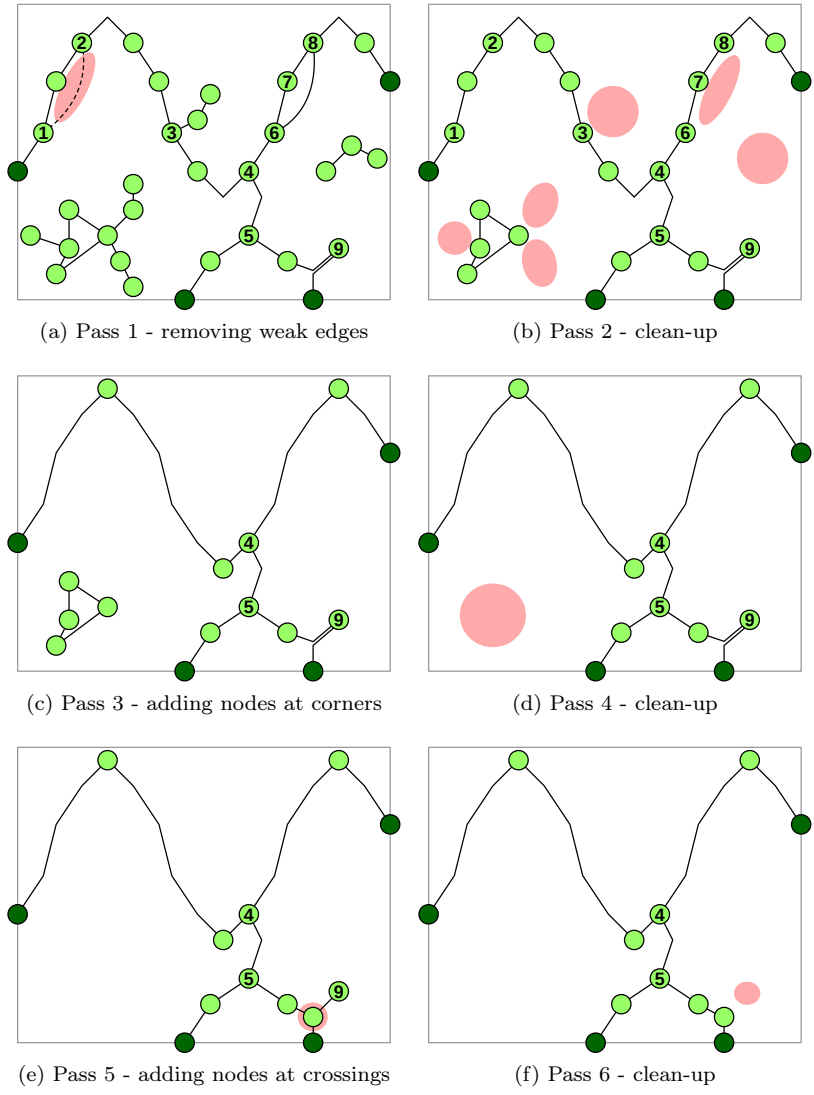

(a) Pass 1 - removing weak edges

(b) Pass 2 - clean-up

(c) Pass 3 - adding nodes at corners

(d) Pass 4 - clean-up

(e) Pass 5 - adding nodes at crossings

(f) Pass 6 - clean-up

Figure 8.6: Illustration of an example of the different steps in the graph processing. For the sake of the example, the dark nodes at the image boundaries are fixed. Changes are highlighted with red ellipsoids. In (a) the weak ($\beta < v_{quality}$) edge e_{12} is removed because it is found not-expected by both end-nodes. Edge e_{68} is also not expected, but its β is high enough to support evidence of a wire, and the edge is thus not (yet) removed. In (b) e_{68} is removed because it is redundant and e_{67} and e_{78} are apparently stronger. Further the tail at node 3 is removed, as well as a few tails in the bottom-left cluster. The bottom-right cluster is removed because it is too small. In (c) new nodes are placed at corners, and nodes with only two edges are "dissolved". In (d) the bottom-left cluster is removed. In (e) a node is placed at a crossing. In (f) node 9 is removed because it is now on a tail.

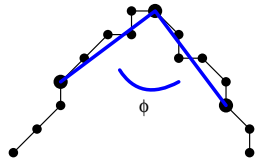

Figure 8.7: Illustration of how the path's curvature is calculated to be able to insert nodes where the curvature is maximal. The black dots represent voxel positions. The algorithm selects the voxel for which the angle ϕ is the smallest.

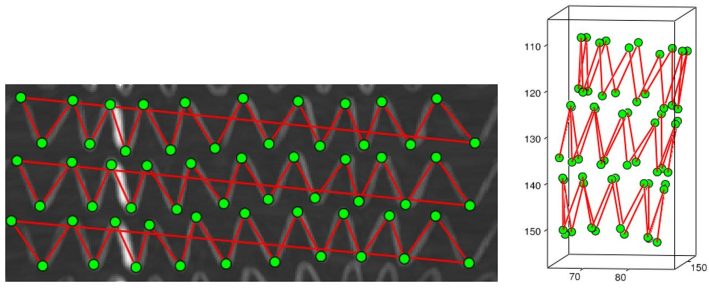

(a) 2D image and annotations (b) Reference graph

Figure 8.8: Illustration of an example image used for the manual annotations (a), obtained using radial raycasting, and the resulting 3D reference graph (b).

Because the (x, y, z) location in the original volume was known for each pixel in the unfolded images, a three dimensional graph could be obtained from the expert's 2D annotations (Figure 8.8b), which served as a reference to evaluate the proposed algorithm.

8.3.2 Comparison Method

To compare the geometric model produced by the algorithm with the annotated reference, a method to compare two graphs is required. In graph theory the *edit distance* [15] is a common method to compare two graphs. The method considers the list of operations on one graph to transform it into the other. A final distance measure is obtained by adding the costs for each edit operation. Since for the geometric model the edges represent the physical wires, while the nodes are only a way to connect these, we use a cost of 1 to add or remove an edge, and a cost of 0 to add or remove a node. Consequently, the edit distance becomes $N_{miss} + N_{wrong}$, representing the number of false negatives plus the number of false positives. To produce a score expressed as a percentage, we apply the Dice Similarity Coefficient (DSC) [26]:

$$DSC = \frac{2 \cdot N_{match}}{(N_{match} + N_{miss}) + (N_{match} + N_{wrong})} \cdot 100\% \tag{8.2}$$

where N_{match} represents the number of true positives.

8.3.3 Results

Multiple experiments were performed using the annotated reference data. The reference data was divided in a training set containing 6 stent pieces (400 edges in total) obtained from 6 different patients, and a test set containing 12 stent pieces (946 edges in total) obtained from 8 different patients. The sets were divided such that in each set the amount of edges for each stent type is approximately the same. The training set was used for parameter tuning and the test set for measuring the accuracy of the algorithm. The experiments and their results are described in the next paragraphs.

8.3.3.1 Parameter tuning

To choose suitable values for the parameters of the algorithm and to study the algorithm's sensitivity to the chosen parameter values, tuning experiments were performed on the training set. To tune each parameter (n_{phases}, v_{seed}, d_{max}, v_{aff} and $v_{quality}$), the algorithm was applied multiple times using a range of different values for the parameter. Accuracy plots were generated by comparing the resulting models against the annotated reference data, resulting in one plot per annotator. The resulting accuracy plots are shown in Figure 8.9. The plots show the optimal value for each parameter, and also give an idea of how sensitive the algorithm is for the parameter value. The parameters (and their found optimal values) are listed in Table 8.1. Although the removal of tails and small clusters has a positive effect on the resulting geometric model, this cannot be measured in the experiments; because the nodes/edges of tails and erroneous clusters are in general not inside the cylinder that is compared with the reference data, removing these edges has no effect on the accuracy measure.

8.3.3.2 Accuracy

An experiment to evaluate the accuracy of the algorithm was performed with the test set. Using the optimal parameter values found in the tuning experiments, the algorithm's results were compared with the results of the three annotators. The resulting values for N_{match},N_{wrong} and N_{miss} are given in Table 8.2, and the matching scores are illustrated in Figure 8.10a. It can be seen that the accuracy is approximately 95% and 92% for the AneuRx and Zenith respectively. In order to assess the quality of the used reference data, the graphs produced by the annotators were also compared with each-other (Figure 8.10b).

Figure 8.11 and Figure 8.12 give an impression of the different steps of the algorithm and provide a way to visually assess the quality of the algorithm for the AneuRx and Zenith stent graft, respectively. In Figure 8.13 and Figure 8.14 surface rendered images of the geometric models of both stent types is shown. [1] For the purpose of visualization, clusters of edges not attached to the stent were manually removed; because the topology of the graph is known, clusters of interconnected nodes are easily identified and can be removed using an interactive process with just a few mouse clicks.

[1] Interactive flash movies of two examples are available at: http://www.sas.el.utwente.nl/home/almar/stents.html

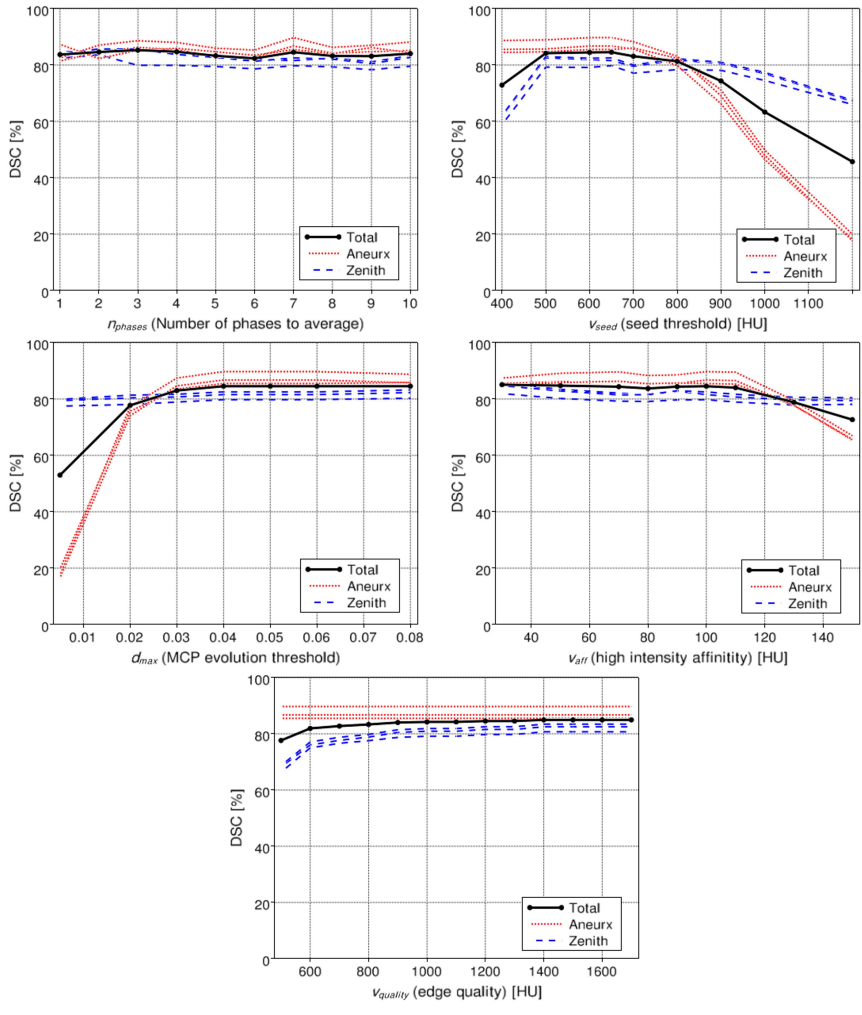

Figure 8.9: Results for the parameter tuning experiments. Shown are three lines (corresponding to the three annotators) for each type of stent graft, and a thick line indicating the total accuracy.

Parameter	Optimal value	Description
n_{phases}	7	Number of phases to average
v_{seed}	650	Threshold to detect seed points
d_{max}	0.06	MCP evolution threshold
v_{aff}	100	High intensity affinity
$v_{quality}$	1200	Threshold for edge quality
$n_{cluster}$	8	Minimum cluster size
n_{tail}	3	Maximum tail length

Table 8.1: List of the algorithm's parameters and the resulting optimal values.

Stent type	Annotator	N_{match}	N_{wrong}	N_{miss}	DSC
AneuRx	1	448	16	19	96.2%
	2	448	16	26	95.5%
	3	442	22	19	95.6%
Zenith	1	470	53	21	92.7%
	2	466	57	27	91.7%
	3	470	53	23	92.5%

Table 8.2: List of the number of true positives, false positives, false negatives and the resulting scores in the accuracy experiment.

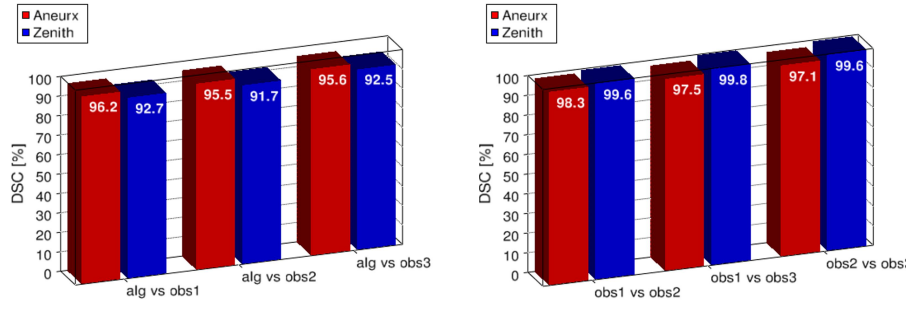

(a) (b)

Figure 8.10: Illustration of the accuracy of the proposed algorithm (a) and the inter-observer agreement of the annotators (b), measured using the 12 stent pieces from the test set. ("alg" stands for the proposed algorithm, "obs" stands for observer.)

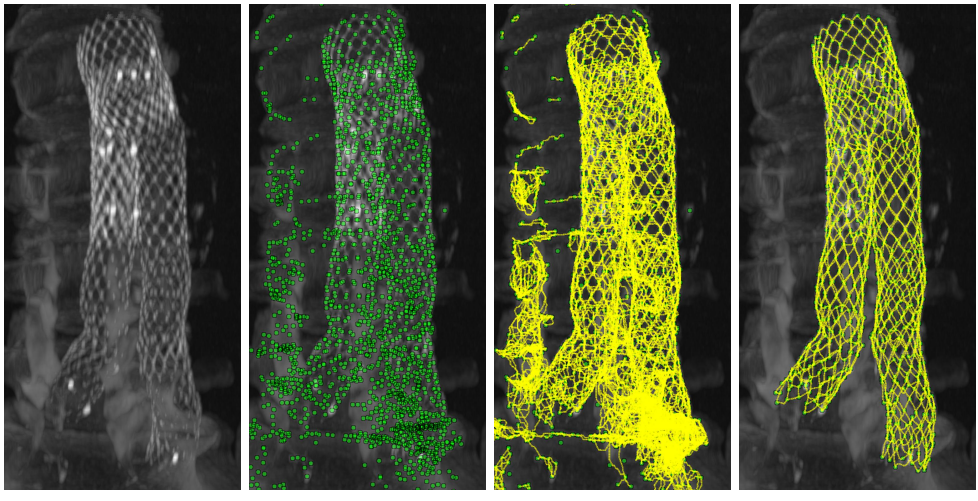

(a) MIP of the CT data (b) 2873 seed points (c) 9630 initial edges (d) 1496 final edges

Figure 8.11: Illustration of the different algorithm steps for an AneuRx stent graft. Shown are a Maximum Intensity Projection (MIP) of the original data (a), the detected seed points (b), the found edges (c), the result after processing the graph (d).

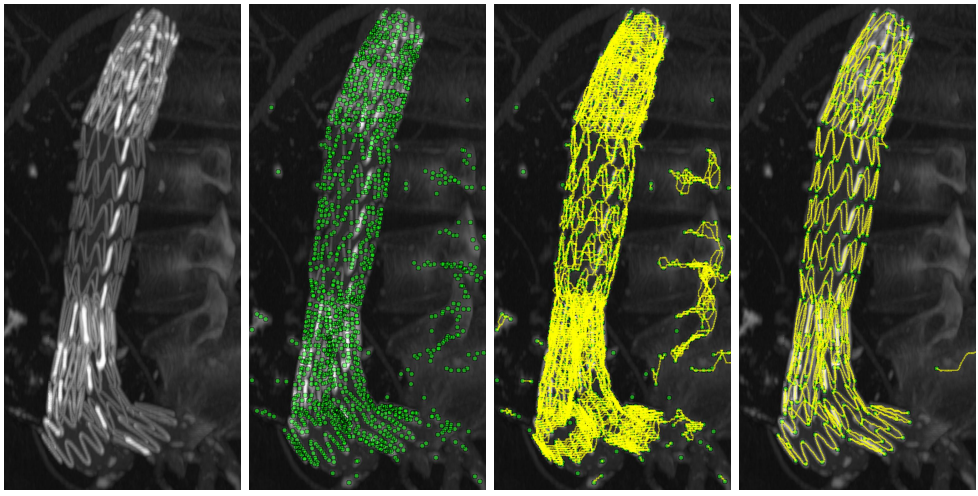

(a) MIP of the CT data (b) 1732 seed points (c) 4963 initial edges (d) 531 final edges

Figure 8.12: Illustration of the different algorithm steps for a Zenith stent graft. Shown are a Maximum Intensity Projection (MIP) of the original data (a), the detected seed points (b), the found edges (c), the result after processing the graph (d).

8.4 Discussion

8.4.1 Parameters

From Figure 8.9 it can be seen that the plots for most parameters are relatively smooth, which implies that small variations in the parameter values do not affect the accuracy of the algorithm. This means that finding the (global) optimal parameter values is relatively easy, as there are no local maxima and the parameters have little effect on each-other. Additionally, this suggests that the algorithm performs consistently on different data using the same parameter values. However, the optimal parameter values can differ if a different scanner type or different acquisition settings are used. All parameters—except for d_{max}—have a direct relation to a physical quantity and are therefore relatively easy to choose.

The plot for n_{phases} (Figure 8.9) shows that this parameter has little effect on the accuracy of the algorithm. This suggests that the algorithm is relatively robust for noise. There is a slight decrease after $n_{phases} = 3$, which can probably be attributed by the increase in motion artifacts. Further, the plot shows a small peak at $n_{phases} = 7$. Whether this peak is significant remains uncertain, but it can be explained as follows. The pressure present in the abdominal aorta first rises quickly after which it slowly declines [49, 80, 107]. The phases to average are selected primarily from the diastolic phase. As more phases are selected for averaging, more phases are included that are not in the diastolic phase, and the increased motion artifacts will negate the positive effect of the reduced noise. The peak at the seventh phase can be explained by the fact that the seventh included phase is usually at the maximum of the pressure pulse, where the amount of motion is low, and thus causes a relatively large accuracy increase. Naturally, when a different number of phases are reconstructed, the optimal

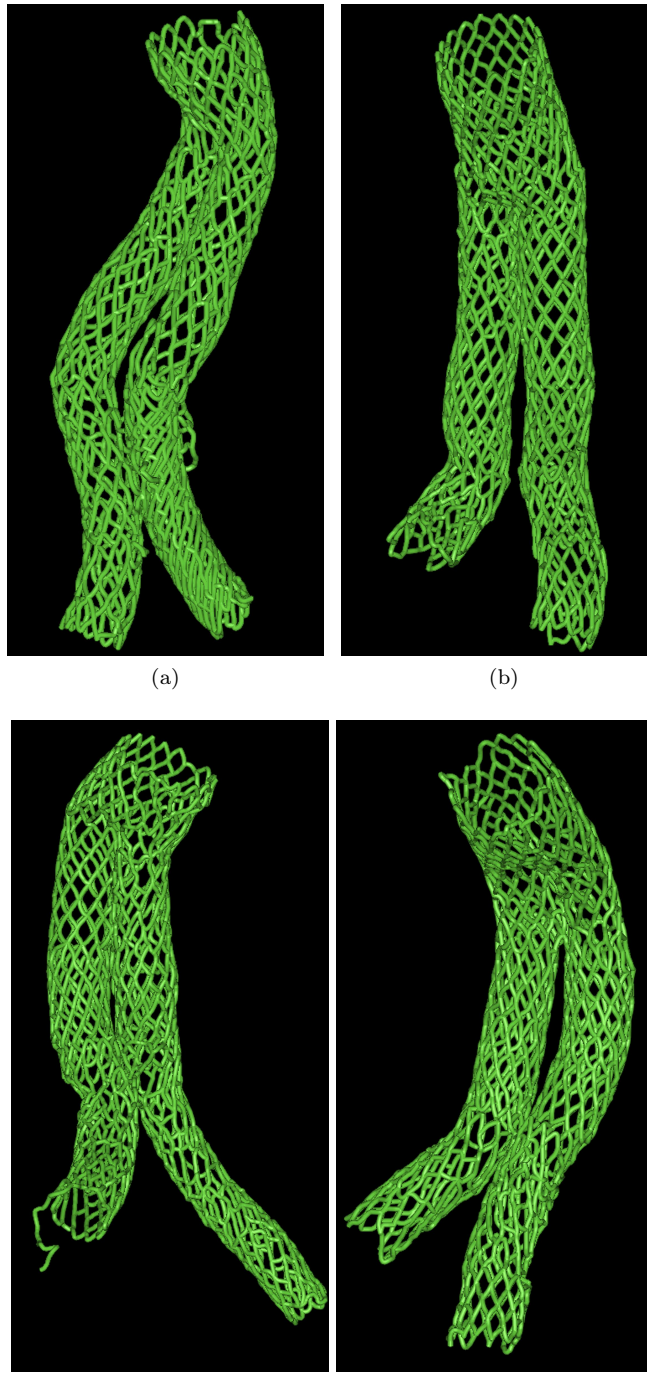

(a) (b)

(c) (d)

Figure 8.13: Illustration of lit surface renders of the resulting models for the 4 datasets containing the AneuRx stent.

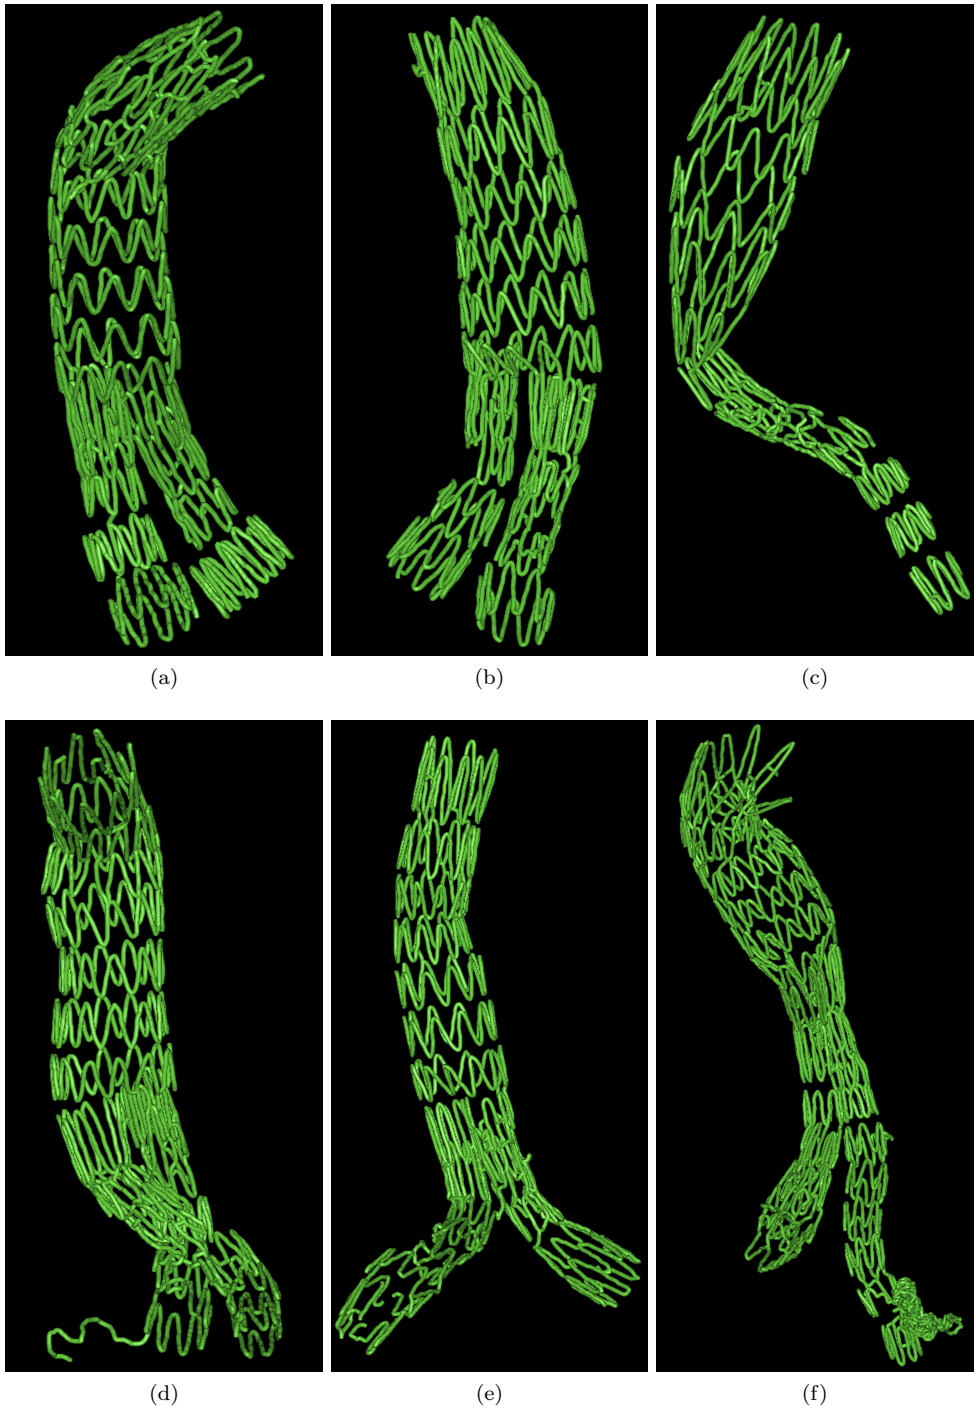

(a) (b) (c)

(d) (e) (f)

Figure 8.14: Illustration of lit surface renders of the resulting models for the 6 datasets containing the Zenith stent.

value of n_{phases} will also change. The relative flatness of this plot is an indication that the algorithm is relatively invariant to noise; the algorithm performs well even when applied to a single phase.

The plot for v_{seed} illustrates a clear maximum accuracy around 650 HU (Hounsfield units). While the materials of which stents are made (stainless steel and nitinol) are usually expressed in much larger CT-values, the partial volume effect causes the parts of the stent where the wire is thin to be represented by CT-values around 650 HU. The accuracy for the AneuRx decays faster because its contrast (in the CT data) is lower than for the Zenith stent graft. This difference can be explained by the difference in material or wire diameter. If the value of v_{seed} is chosen too high, too few seed points are detected in order to segment the stent completely. If the value is chosen too low, many false seed points are introduced that cannot be cleaned up by the graph processing algorithm, particularly when contrast fluid was applied during scanning.

The plot for the MCP evolution threshold d_{max} shows the minimal value to use in order not to miss any connections between nodes. While higher values will result in more edges being produced, these are successfully cleaned up by the graph processing algorithm. The plot for the high intensity affinity parameter v_{aff} shows a broad optimum around 100 HU.

The plot for $v_{quality}$ illustrates that this parameter has no effect for the AneuRx stent graft. This suggests that there are usually no extra edges found between nodes. Although the plot suggests that (for the Zenith stent graft) there is no maximum for this parameter value, we chose a value of 1200; the situations in which this parameter is applied (when there are more genuine edges than expected) are relatively rare and are probably not represented sufficiently in the training data. Fortunately, the value can be chosen by considering what it physically represents; an edge with a minimum CT-value of 1200 can only represent bone or a metallic structure. The plot does show that using values lower than 1200 HU results in erroneous edges.

The values of $n_{cluster}$ and n_{tail} could not be chosen based on experimental data (see section 8.3.3.1), and their values might therefore not be optimal. Visual inspection of the results showed that most tails are successfully removed, but some invalid clusters remain, specifically on the spinal cord (which were removed interactively). Although increasing the value of $n_{cluster}$ might remove these false clusters of nodes, this will also result in parts of the stent being removed (particularly for the Zenith stent graft, for which the graph consists of several non-connected clusters).

In the brief discussion on the algorithm to calculate the curvature on the path P_{ij}, two parameter values were mentioned. These were not tested in the experiments because it was found that the corners are placed in the correct location even when the parameters were varied significantly.

8.4.2 Accuracy

Figure 8.10a show that the correspondence of the algorithm with the experienced human annotators is 95% and 92% for the AneuRx and Zenith, respectively. This is currently (to the best of our knowledge) the only published algorithm to segment stent grafts is such a precise way. The fact that the accuracy is slightly lower for the Zenith can be explained by the occasional sharp corners in its wire frame, which are sometimes cut in the MCP method.

Figure 8.10b shows that the consensus between the different annotators is high.

This means that the process of annotating provides consistent reference data and suggests that the reference data is a good approximation of the real topology of the stent. In the future the reference data can be used to perform similar experiments to compare different algorithms that perform the same task.

From Figure 8.13 and Figure 8.14 it can be verified that the overall accuracy is good, but some errors are clearly visible. Most errors are located at "difficult situations" such as the parts where the wire frame has relatively sharp corners (Zenith), small mesh size (AneuRx), or bifurcations (AneuRx and Zenith). Further it can be seen that for the AneuRx the frame contains holes, and that nearby structures are sometimes attached to the stent (such as the coil in Figure 8.14f).

Figure 8.13 and Figure 8.14 show that there is still room for improvement. We expect to gain the most by improving the graph processing, for example by more explicitly incorporating information about the known topology of the stent. Nevertheless, we believe that the current results are good enough to enable studying the motion of stent grafts in vivo. However, it should be noted that small errors in the topology of the found model can lead to significant errors when calculating the forces present in the stent. The extent to which the errors affect such calculations should be investigated.

Currently the algorithm requires approximately 65 seconds for the AneuRx and 30 second for the Zenith on a Pentium 4 with 2.4 GHz (using a single thread). This difference can be explained by the different number of edges between the stent types.

8.4.3 Limitations

In this study we focused on two different stent types—the AneuRx and Zenith—which are made of nitinol (a nickel titanium alloy) and stainless steel, respectively. We expect that the proposed algorithm can be used to segment other stent types as well, provided that the contrast is high enough. This should be the case if such a device is made of a similar material and has the same (or larger) wire diameter. In initial tests (not shown in this work) we have seen positive results for the Talent stent graft (Medtronic, Minneapolis, USA), while the Excluder stent graft (Gore, Flagstaff, USA) seemed too ill-defined in the resulting CT data when scanned with our scanner settings. More research is needed to determine the cause for these differences.

While the optimal parameter values are the same for both stent types under consideration (except for the expected number of edges per node), this is not necessarily true for other stent types.

The annotated reference data does not contains parts of the stent with bifurcations. These are hard to annotate for three reasons: firstly, a bifurcation cannot be unfolded in the way as proposed in this work. Secondly, the topology of the frame can be relatively complex in these regions. Lastly, such regions can contain wires that touch each-other, which are difficult to distinguish from wires that are physically attached. Since the algorithm does not suffer from the first two limitations, we can expect the algorithm to produce accurate results in these regions. This can, however, not be verified with the presented experimental method.

8.4.4 Implementation details

Most of the algorithm was implemented in the Python programming language[2]. A time-critical part of the algorithm is the list L that stores the voxels of the "front" in the MCP method (Algorithm 8.1). Because at each iteration the voxel with the minimum D must be selected, usually some form of a binary heap is employed [30], which enables obtaining the minimum value in $O(log2(N))$. For increased speed, the detection of seed points, the binary heap, and a large part of the MCP method were implemented in Cython [104], a dialect of the Python programming language that is compiled to the C language.

The memory requirements of the algorithm depend mostly on the MCP method. Considering a dataset of $512 \times 512 \times 512$ voxels, the memory requirements can be shown to be 2944 MB to store the volumetric maps required by the MCP method. By cropping the data and selecting the region containing the stent, the data could in our case be reduced to $256 \times 256 \times 256$ voxels, implying a memory requirement of just 368 MB. The latter requirement should not be a problem even for older (32 bit) personal computers.

Accuracy and speed improvements may be achieved by applying a method that detects the seed points in a more reliable way, while reducing the amount of seed points found on bone. Further speed improvements may be realized by using a different criterion to stop the evolution of the MCP algorithm, and by optimizing the graph processing algorithms. Nevertheless, we believe that the algorithm in its current form is fast enough to be used in the clinic.

8.5 Conclusions

We propose a method to segment stent grafts from CT data, consisting of three steps: the detection of seed points, finding the connections between these points to produce a graph, and graph processing to obtain the final geometric model. Using annotated reference data, the method was optimized and its accuracy evaluated. It was found that the algorithm performs good for both types of stent grafts under consideration and is robust for noise and small variations in the used parameter values. Additionally, the algorithm does not consume much memory for modern standards and is fast enough to be used in a clinical setting.

The proposed algorithm produces an accurate geometric model of the stent in the form of an undirected graph. This model allows incorporation of high level information and material properties. We can now use this model to study the in vivo motions and forces that act on the frame of the stent. We believe that such studies will provide new insights in the behavior of the stent in vivo, enables the detection and prediction of stent failure, and by better understanding the reasons for stent failure can play a role in designing better stent grafts in the future.

Acknowledgment

The authors wish to thank Zachary Pincus from Yale University for the fruitful discussions on the MCP method and its implementation in Cython.

[2] http://www.python.org

Part C

Registration

image registration: *the process of transforming different sets of data into one coordinate system. The goal of registration is to enable comparison or to find the transformation between the different sets of data.*

9

A generic approach for groupwise diffeomorphic image registration

A paper based on this chapter is in preparation.

Abstract

This chapter presents the technical details of the groupwise registration algorithm that is used in Chapter 10 to estimate the motion of stent grafts. It is essential that the transformations are diffeomorphic, such that the transformations have a unique inverse and the motions can be reconstructed from the registration result. Transformations can be made diffeomorphic using a simple method based on B-spline grids. However, this requires that the total deformation is represented as a composition of many small deformations.

We propose a groupwise diffeomorphic registration method in which deformations are represented using discrete displacements fields, and that uses B-spline grids as a tool to regularize the delta deformations. The proposed generic method enables the use of different types of forces to drive the registration. Using Demons forces we obtain a groupwise diffeomorphic variant of the Demons algorithm. We also propose a "gravity" force based on the attraction between edges in the images.

The proposed methods were quantitatively evaluated, and compared against the classic Demons algorithm and an algorithm based on the optimization of mutual information. The results show that the proposed Demons algorithm has a high accuracy if the intensities between the images are similar. The proposed gravity algorithm is shown to be robust for distortions such as different lighting and bias fields.

9.1 Introduction

Groupwise registration is the process of registering multiple images simultaneously—in contrast to pairwise registration, in which just two images are aligned. The application for this approach ranges from the creation of atlases (e.g. [7, 103, 56, 10]), as well as the formation of active shape models [25] and calculating the motions in dynamic data [89, 35].

In the latter two applications it is essential that the transformations are diffeomorphic, since such transformations have a unique inverse [35, 4]. Further, a diffeomorphic deformation preserves topology, which is important in medical image registration to ensures that the anatomy does not change in an unnatural manner. A transformation that it is diffeomorphic (i.e. a diffeomorphism) is continuous, one-to-one (injective), onto (surjective), and differentiable [20]. In the context of deformations, this means that there are no folds or gaps.

When a deformation is described with a B-spline grid, continuity and differentiability are implicitly guaranteed. Injectivity for B-spline grids can be guaranteed using a simple method proposed by Choi et al.[19]. Since constraining a grid in this manner allows just very small deformations, this approach is only applicable by maintaining a total transformation and combining it with a delta deformations at each iteration.

9.1.1 Contribution

We propose a groupwise diffeomorphic registration method in which the total transformation is composed of many small transformations. The total transformation for each image is represented using discrete displacement fields, and a B-spline grid is used as a tool to regularize the delta transformation at each iteration. Consequetly, deformations can be constrained to be diffeomorphic in a simple way. Since the resulting transformations can easily be inverted and combined they are suitable for various applications, such as atlas building and calculating changes or motions in time series.

Further, this generic approach allows the use of different types of "forces" to drive the registration. We demonstrate this with Demon forces to create a variant of the Demons algorithm that is groupwise and diffeomorphic. Additionally we propose a new image derived *gravity* force, which is robust for changes in the image intensities, and can also be used for multi modal image registration.

9.1.2 Overview

In order to explain the proposed method we shall first discuss the theory and state of the art related to groupwise diffeomorphic image registration (section 9.2). Next, we will explain the proposed generic registration method and two implementations thereof (section 9.3). In section 9.4, we describe the experiments that we performed in order to validate the proposed method, and in section 9.5 we discuss the results.

9.2 Theory

Let $I_i(\mathbf{x}) : \mathbb{R}^d \to \mathbb{R}$ be an image of dimension d, with $i \in [1 \dots n]$ and n the number of images to register. Let $T_{ij}(\mathbf{x}) : \mathbb{R}^d \to \mathbb{R}^d$ be the transformation function that maps I_i to match with I_j. We assume backward mapping transformations: $I_{ij} = I_i(T_{ij}(\mathbf{x}))$.

In a pairwise registration setting ($n = 2$), the goal is to find T_{12}. In a groupwise registration setting, the goal can be to find T_i for all i such that all images can be transformed to a common reference frame. Alternatively, the goal can be to find T_{ij} for all i to transform all images to match with image I_j. For simplicity, we assume that all images have the same size.

9.2.1 Representation

Without loss of generality, we assume that transformations are represented using a displacement vector field, which describes the relative displacement for each location $\mathbf{x}$ in the image. In practice, a *discrete displacement field* is often used: an image with the same size as I_i, where each pixel represents the relative deformation at that pixel. With Γ such an image we obtain for the corresponding T:

$$T(\mathbf{x}) = \mathbf{x} + \Gamma(\mathbf{x}) \tag{9.1}$$

Another common approach is to use a *B-spline grid* to describe the deformation [101, 89]. It consists of a grid of knots that is overlaid on the image. To calculate the deformation of a pixel, the contributions of the four nearest knots in each dimension are combined:

$$T_{ij}(\mathbf{x}) = \mathbf{x} + \sum_{\mathbf{x}_k \in \Pi_{\mathbf{x}}} B_k(\mathbf{x} - \mathbf{x}_k)\phi_k \tag{9.2}$$

where $\mathbf{x}_k$ are the control points, $B(\mathbf{x})$ the multidimensional B-spline *basis functions* [113], ϕ_k the B-spline knot values, and $\Pi_{\mathbf{x}}$ the set of all control points within the compact support of $\mathbf{x}$. The advantage of using a B-spline grid is that it describes a continuous field in an efficient manner, and its inherent smoothness leads to realistic deformations.

A transformation represented as a B-spline grid can easily be converted into a discrete displacement field. The other way around, however, is more difficult. Lee et al. [78] have proposed a solution that can be shown to minimize the sum of squared differences. We rewrite Equation 5 from [78] to apply it to image data:

$$\phi_k = \frac{\sum_{\mathbf{x} \in \Pi_k} \omega_{\mathbf{x}}^2 \cdot \phi_{\mathbf{x}}}{\sum_{\mathbf{x} \in \Pi_k} \omega_{\mathbf{x}}^2}, \tag{9.3}$$

where $\omega_{\mathbf{x}} = B_k(\mathbf{x} - \mathbf{x}_k)$, the set Π_k contains the pixel locations within the compact support of k, and $\phi_{\mathbf{x}}$ is the deformation to apply at $\mathbf{x}$ weighted by the B-spline basis functions (see Equation 4 in [78]). This formula is applied in an iterative fashion by going from a high/coarse scale to the finest required scale in steps of two [78]. Therefore this method is relatively computationally expensive. Further, it should be noted that a B-spline grid can only model a subset of the deformations that can be modeled with a discrete displacement field. For these reasons, if both representations are used together, it is more straightforward to transform a B-spline grid into a discrete displacement field than the other way around.

Although there is no technical reason to avoid using these two representations together, there seems to be a tendency to use either a discrete displacement field or a B-spline grid—not both.

9.2.2 Two classes of algorithms

The current range of common region based image registration algorithms can be divided into two classes. Both classes usually adopt a multiscale approach to prevent convergence to local minima, and to speed up the registration process. In the first class the total transformation consists of a combination of small transformations. At each iteration a *delta transformation* is calculated, which is regularized and then combined with the total transformation. The delta deformations are often calculated using pixel-based image forces. The most notable method is the popular Demons algorithm [109].

In the second class only a single total transformation is maintained, which is updated at each iteration using an optimization procedure that minimizes/maximizes a certain similarity metric. Often a B-spline grid is used to describe the deformation field (e.g. [101, 118]), which has the advantage that the deformations are described with few parameters and are physically realistic [93].

Regularization in the second class of registration algorithms is by definition applied to the total transformation. Although this allows for a high degree of control by for example penalizing small Jacobians [22, 116], some deformations (e.g. rotations) are difficult to describe using a B-spline grid [72]. In algorithms belonging to the first class usually only the delta deformations are regularized, which causes a fluid-like effect and enables modeling a larger set of transformations. However, this can also lead to unrealistic transformations if no special care is taken.

9.2.3 Combining deformations

In literature deformations are sometimes reported to be combined by addition (e.g. [109]). Even though this can produce satisfactory results in an iterative setting, this operation is an approximation. Furthermore, adding two diffeomorphic transformations does not guarantee that the result is diffeomorphic too.

A better approach is to combine deformations using *composition*, which is an operation with the property that if both deformations are diffeomorphic, the result is also diffeomorphic [77, 4]. Given a deformation T_1 and a delta deformation T_2 that should be combined with it, we obtain the total deform T_3. With $\circ$ the mathematical composition operator:

$$T_3 = T_1(T_2) = T_1 \circ T_2 \tag{9.4}$$

Unfortunately, the composition of two transformations can not be described using a B-spline grid, even if both transformations are B-spline grids [116]. Therefore, either a cascade of transformations is employed [102], which significantly reduces the speed of the algorithm, or one must use a discrete displacement field to describe the total deformation.

9.2.4 Diffeomorphic transformations

A transformation that it is diffeomorphic (i.e. a diffeomorphism) is continuous, one-to-one (injective), onto (surjective), and differentiable [20]. If a transformation is continuous, injectivity and surjectivity are analogous (Figure 9.1).

In the context of image deformations, we can distinguish two types of injectivity. We call a transformation *normally injective* if the domains of I_i and I_j are considered

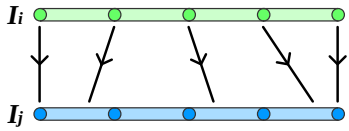

(a) Injective: every pixel in I_i maps to a unique point in I_j.

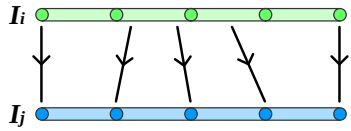

(b) Surjective: every pixel in I_j is accounted for.

Figure 9.1: Illustration of an injective map (a) and a surjective map (b). It can be seen that if an injective mapping is also continuous, it is also surjective.

(a) Not injective (b) Normally injective (c) Fully injective

Figure 9.2: Illustration of an image with an imposed deformation. The result is shown for different injectivity constraints.

to expand beyond the boundaries of the images. In this case, some pixels may be mapped outside the image boundaries, and some pixels from outside the boundaries may be mapped inside the image (Figure 9.2b). This can be a problem in case multiple deformations are combined or when the inverse deformation is calculated, because then the result may not be diffeomorphic.

We call a transformation *fully injective* if the domains of I_i and I_j are bounded by the borders of the image. With such a mapping all pixels remain inside the image and no pixels from outside the image are introduced (Figure 9.2c).

9.2.4.1 Penalizing small Jacobians

One approach to constrain a transformation to be diffeomorphic is to prevent folding by ensuring that the Jacobians of the total deformation do not go below zero [102, 22, 116]. This approach has the advantage that it is relatively easy to include in algorithms that optimize a similarity metric. However, it has been shown that the constraint imposed on the transformation is sometimes so large that it can have a significant negative effect on the registration result [102].

9.2.4.2 Diffeomorphic constraints based on continuum mechanics

Inspired by the transport equations from continuum mechanics, early approaches for diffeomorphic image registration use a dynamic vector field to model one image

as it flows to match the shape of another image (e.g. [22, 20, 9, 56]). Since this approach is very time and memory consuming and therefore not suitable for large volumetric datasets [35], algorithms based on a stationary vector field have gained popularity. While this approach has fewer degrees of freedom and can only model a subgroup of diffeomorphisms, the versatility of this stationary approach has been shown to be sufficient in several applications [35, 116, 4]. Most of these algorithms use a variant of the Log-Euclidean framework originally proposed by Arsigny et al. [3], which also enables calculating statistics on diffeomorphisms. A disadvantage is that the logarithm may become ill-defined in the presence of large deformations [3]. This approach has been combined with the popular Demons algorithm [109] to obtain a version of the algorithm that is diffeomorphic [116].

9.2.4.3 Diffeomorphic constraint using a B-spline grid

When a deformation is described with a B-spline grid, continuity and differentiability are implicitly guaranteed. Injectivity for B-spline grids can be guaranteed using a simple method proposed by Choi et al.[19]. They propose and proof simple conditions for the knot values of the B-spline grid: if the values of all knots in the grid are less then the critical value ϕ_d/K the grid is guaranteed to be injective. In this expression ϕ_d is the distance between the knots and the value of K is 2.04639 and 2.47947 for 2D and 3D, respectively.

One advantage of this method is that it enables averaging multiple deformations to obtain a mean deformation that is also diffeomorphic. While this might seem obvious, this is not possible for diffeomorphisms in general [3]. To show that averaging indeed yields a diffeomorphism, consider averaging a set of transformations that are described using B-spline grids subject to Choi's constraint. Clearly, the values of the knots in the resulting average grid will also be below the critical value.

Since constraining the grid in this manner allows just very small deformations, this approach is only applicable to the class of algorithms that update the total transformation by composing it with many delta deformations. Because the composition of two deformations cannot be described using a B-spline grid, the final deformation should be described using a cascade of transformations or using a discrete displacement field.

We believe that this is the main reason why this relatively simple method is currently not widely applied, particularly in algorithms that describe the deformation using a B-spline grid and minimize a certain cost function [101, 67].

9.2.4.4 Constraining a transformation to be fully injective

A transformation can be constrained to be fully injective if everywhere in the image the deformation is not larger than the distance to the nearest edge, i.e. if no point on the inside of the image is transformed outside of the image and vice versa. This constraint can be attained by—in addition to the normal diffeomorphic constraint—constraining the B-spline grid in such a way that the transformation is zero at the edges.

Without loss of generality, assume a 1D B-spline grid that has exactly two knots beyond the edge of the image, such that the edge of the image is between knots ϕ_2 and ϕ_3. The value can be made to go to zero at the edges in a smooth way by setting

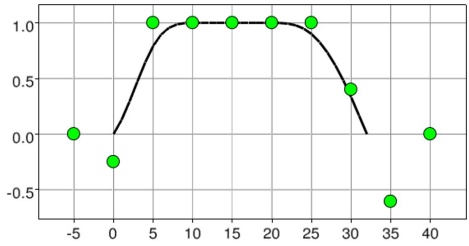

Figure 9.3: Illustration of a B-spline grid that is modified such that the field that it describes is zero at the edges. Shown are the modified knots of a 1D B-spine grid (that were initially all 1). The field smoothly goes to zero at the edges in about two knot-spacings.

the first three knot values (Figure 9.3):

$$\phi_3 \leftarrow (1-t)\cdot\phi_3 \tag{9.5}$$
$$\phi_2 \leftarrow -(\phi_3\cdot B_3(t)+\phi_4\cdot B_4(t))/B_2(t) \tag{9.6}$$
$$\phi_1 \leftarrow 0 \tag{9.7}$$

where B_1 through B_4 are the B-spline basis functions [113], and t is a factor between 0 and 1, corresponding to the image edge being between ϕ_2 and ϕ_3, respectively. The value of ϕ_3 in the second line is the new value as determined by the first line. If the grid is aligned with the image ($t = 0$) the equations reduce to $\phi_1 \leftarrow 0$ and $\phi_2 = -\frac{1}{4}\phi_3$.

9.2.5 Groupwise registration

One groupwise registration approach is to apply pairwise matching repeatedly; each image is transformed to all the other images, and the final transformation is the average of these transformations. By doing this for all images, they are all mapped to a reference frame which represents the mean shape of the population [103]. Another approach is to deform all images simultaneously while constraining the mean transformation to be the identity transform. This has been demonstrated using Mutual Information as a similarity metric [11], and using a metric based on the sum of entropies along pixel stacks [7]. In a recent approach the deformation is described using a 4D B-spline grid (or 3D for 2D data), such that the deformation between the images is constrained to be smooth [89]. Groupwise registration has also been performed in a diffeomorphic framework by [25] and [21].

Under certain conditions, groupwise registration is *symmetric*: T_{ij} and T_{ji} are calculated simultaneously, while ensuring that they are each-other's inverse [20]. Even if the inverse transformation is not of interest, the constraint itself will be beneficial for the registration result (except in cases that uniquely determine the correspondence between the images, i.e. landmark based registration) [20].

9.3 Methods

Based on the theory described in the previous section, we propose a groupwise diffeomorphic image registration algorithm that belongs to the class of algorithms that

uses delta transformations to update the total transformation at each iteration.

9.3.1 Generic registration method

The algorithm guarantees fully diffeomorphic transforms by applying Choi's constraint [19] to the delta transformations, and by applying the method discussed in section 9.2.4.4. Composition is used to combine the delta deformations with the total deformation. Consequently, the total deformation is described using a discrete displacement field. The registration is performed iteratively while the scale s is reduced from coarse to fine. At each iteration, I_i^s is the deformed image I_i at scale s.

9.3.1.1 Groupwise approach

Groupwise registration is obtained by simultaneously registering the images towards a virtual reference frame. A total transformation T_i is maintained for each of the images being registered. The delta deformation ΔT_i is calculated by taking the mean of the delta deformation between the images:

$$\Delta T_i = \frac{1}{n} \sum_{j=1}^{n} \Delta T_{ij}, \tag{9.8}$$

This method is similar to [103], but the proposed method applies the averaging of the transformations at each iteration rather than at the end.

9.3.1.2 Delta deformations

The delta transform ΔT_{ij} is calculated from the deformation force F_{ij}, which is derived from the images I_i^s and I_j^s, and can be obtained in any way. Such a wide variety of forces is possible because of the regularization method that is applied to F_{ij}.

To convert F_{ij} to ΔT_{ij} the transformation is made diffeomorphic. To convert F_{ij} to ΔT_{ij} it is first converted to a B-spline grid (if necessary) using Equation 9.3, with a grid spacing that is proportional to the scale s. A multiscale approach is not required in this case, since the forces already match with the scale. Next, the knot values are limited according to Choi's constraint [19]. Instead of clipping the values at the required limit, we limit the values in a smooth way:

$$\phi_k \leftarrow \begin{cases} \beta \cdot (1 - e^{-\phi_k/\beta}) & \text{where } \phi_k > 0 \\ -\beta \cdot (1 - e^{\phi_k/\beta}) & \text{otherwise} \end{cases} \tag{9.9}$$

This function has a slope of 1 at the origin and has an asymptote at $+\beta$ and $-\beta$. The value of β according to Choi's constraint is $\beta = \frac{\phi_d}{K}$, with ϕ_d the distance between the knots of the B-spline grid. We propose to tighten this constraint even further to limit the maximum deformation at each iteration to a value relative to the scale, thereby preventing abnormally large deformations and improving convergence: $\beta = \frac{s}{K}$. Next, the grid is constrained to be fully injective (section 9.2.4.4).

9.3.1.3 Scale space approach

In most publications on image registration, the scale s is reduced with factors of two. In contrast, we consider the scale as a continuous parameter. Therefore, we chose to

smoothly decrease the scale and define a fixed number of iterations p_{iters} for each factor-of-two reduction. This has the effect that the average deformation found at each iteration scales smoothly with scale, which might increase the probability of finding the global optimum registration. The parameter p_{iters} represents the scale sampling, and is comparable to the "number of iterations per scale" parameter of other registration algorithms.

The algorithm moves down in scale space until the final scale is reached, at which point the registration process stops. There is no need for a stop criterion. Instead, the assumption is made that the algorithm converges fast enough to “keep up". If this is not the case the value of p_{iters} should be increased.

9.3.1.4 Obtaining the final transformation

After the registration is finished, there is for each image I_i a transformation T_i that maps the image to the reference frame (Figure 9.4a). It is now possible to obtain any transformation T_{ij} by using the following equation (Figure 9.4b):

$$T_{ij} = C(T_i, T_j^{-1}). \tag{9.10}$$

It can be shown that T_{ij} and T_{ji} are each other inverse. With $(f \circ g)^{-1} = (g^{-1} \circ f^{-1})$ it follows that:

$$T_{ij}^{-1} = C(T_i, T_j^{-1})^{-1} = C(T_j, T_i^{-1}) = T_{ji} \tag{9.11}$$

For some applications, such as atlas building, the transformation to the mean shape is of interest: $\hat{T}_i = \frac{1}{n} \sum_{j=1}^{n} T_{ij}$. Provided that the delta transformations are symmetric ($\Delta T_{ij} = -\Delta T_{ji}$) the reference frame is also the mean shape: $T_i = \hat{T}_i$.

9.3.1.5 Obtaining transformations in between images

Any transformation to another image can be obtained by composing a transformation that maps to the reference frame with one that maps back to the "image space". In some situations (e.g. motion analysis) it may be of interest to map the images to a virtual image in between two images, i.e. to find T_{ij} where j is not integer. This is possible by using a transformation from reference frame to image space that is a linear combination of transformations (Figure 9.4b). This makes it possible to interpolate between different time frames in a motion sequence, using linear or even cubic interpolation.

It should be noted that in this case the linear combination of transformations is not mathematically guaranteed to be diffeomorphic [3]. However, due to the relation between the transformations that are combined, it is unlikely that such a transformation becomes non-diffeomorphic.

9.3.2 Diffeomorphic Demons registration

An image-based force that is known to provide good results if the images have similar intensities is the Demons algorithm [109]. By calculating F_{ij} using the Demon forces, we obtain a variant of the Demons algorithm that is diffeomorphic. We use the improved force equation for which the force is derived from both images, which is symmetric ($F_{ij} = -F_{ji}$) and more robust [117]. The force is calculated as follows:

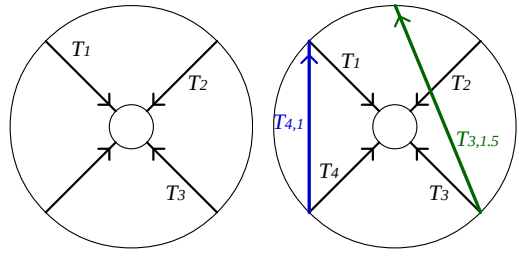

Figure 9.4: Illustration of the deformations that are obtained with the registration algorithm. The outer circle represents the space of original images. The small inner circle represents the reference frame. In (a) an example is shown for four images, using forward mapping. In (b) it is illustrated how one can obtain transformations from one image to another image, and even in between images: $T_{41} = C(T_4, T_1^{-1})$ and $T_{3,1.5} = C(T_3, 0.5{\cdot}(T_1{+}T_2)^{-1})$. Note that linearly combining multiple transformations is only valid for forward mapping transformations.

$$F_{ij} = \left(\frac{\nabla I_i^s}{|\nabla I_i^s|^2 + p_{\text{noise}} \cdot (I_i^s - I_j^s)^2} + \frac{\nabla I_j^s}{|\nabla I_j^s|^2 + p_{\text{noise}} \cdot (I_i^s - I_j^s)^2} \right) \cdot (I_i^s - I_j^s) \cdot p_{\text{speed}}, \tag{9.12}$$

where I_i^s is the deformed image I_i at scale s, and p_{noise} is a parameter that is a measure for the amount of noise in the image. We also introduce the parameter p_{speed} to influence the relative strength of the deformation. Instead of regularizing using Gaussian diffusion, the B-spline based regularization as described in the previous subsection is used.

9.3.3 Gravity Registration

In addition to the diffeomorphic Demons algorithm we also propose a registration algorithm based on attracting forces between objects in the images. A preliminary version of this algorithm was proposed in [61]. The first step of the algorithm is to obtain the mass-images M_i^s and M_j^s. These mass images should be created in such a way that they both contain a sparsely distributed set of structures (a group of high intensity pixels/voxels). The main idea of the Gravity algorithm is that the structures in both mass images attracts each-other; the larger a structure's weight (i.e. higher intensities), the larger the attractive force.

9.3.3.1 Creating the mass images

Because images can be considered similar if the intensity changes occur at the same location, we propose to use the gradient magnitude to detect intensity changes in

both images[1] (The index i is omitted in the next equations for simplicity.):

$$M^s = |\nabla I^s|. \tag{9.13}$$

Further, a series of operations is applied to normalize the mass image. First, the image is normalized such that its mean is zero and its standard deviation is 2. We now assume that the values below zero belong to the background noise, and that the important masses are between 0 and 1. To prevent pixels with high values in the mass image to result in extreme deformations, these pixels are tempered. All values below zero are made zero, and the values are smoothly limited to 1 using the same approach as proposed in Equation 9.9:

$$M^s \leftarrow \begin{cases} 0 & \text{where } M^s(\mathbf{x}) < 0 \\ 1 - e^{-M^s} & \text{otherwise} \end{cases} \tag{9.14}$$

Finally, linear diffusion is applied to smooth the mass image and increase the region of attraction: $M^s \leftarrow M^s \otimes g^s$, where g^s is a Gaussian kernel corresponding to scale s.

9.3.3.2 Calculating and regularizing the deformation

The total force is calculated by combining the obtained forces of the two images. First, the "gravity field" is obtained by calculating the gradient of the mass:

$$F_{ij} = (\nabla M_j^s - \nabla M_i^s) \cdot s \cdot p_{\text{speed}} \tag{9.15}$$

where p_{speed} is a parameter of the proposed algorithm that can be used to influence the relative strength of the deformation at each iteration. The vector field F_{ij} describes the deformation for each pixel. However, since in our situation the intensities (i.e. masses) are sparsely distributed over the image, normal regularization will lead to unsatisfactory results; pixels which have a high mass should contribute more than pixels which have a small or zero mass. In order to solve this problem, we modify Equation 9.3 by introducing a weight for each pixel:

$$\phi_k = \frac{\sum_{\mathbf{x} \in \Pi_k} \alpha_{\mathbf{x}} \cdot \omega_{\mathbf{x}}^2 \cdot \phi_{\mathbf{x}}}{\sum_{\mathbf{x} \in \Pi_k} \alpha_{\mathbf{x}} \cdot \omega_{\mathbf{x}}^2}. \tag{9.16}$$

The introduced weight factor is set to $\alpha_{\mathbf{x}} = M_i^s(\mathbf{x}) \cdot M_j^s(\mathbf{x})$. The underlying idea to this approach is that the deformation is calculated by evaluating the gravity field created by both images at the location of the mass. In addition to enabling the deformation to be described by a sparse set of pixels, the introduced weight factor also has a noise regulating effect when used in combination with the normalization step in the calculation of the mass images.

[1] Depending on the type of images, other operators such as the absolute of the Laplacian may be used as well.

9.4 Experiments and results

A series of experiments have been performed to demonstrate and evaluate the proposed algorithm. The first is a qualitative experiment to demonstrate the applicability of the proposed algorithm to a groupwise registration problem, and to show how the resulting transformations can be used further.

The next four experiments were performed to quantitatively evaluate the accuracy and robustness of the proposed algorithms. They were compared with an implementation of the Demons algorithm as proposed in [117], and with a registration algorithm that optimizes Mutual Information (MI). For the latter we have used the implementation that is freely available in the Elastix toolkit[67][2]. Each of the experiments was performed on different kind of data, and in each experiment a distortion was introduced to evaluate the robustness of the algorithms. For brevity, we denote the proposed algorithms "DD" for diffeomorphic Demons, "Grav" for Gravity, "CD" for classic Demons, and "MI" for Mutual Information.

9.4.1 Random transformations

In some of the following experiments, random transformations were used. These deformations were calculated by randomly sampling 50 points inside the image. From each point, a random vector was generated with a maximum norm of 20 pixels. This description was then converted to a fully diffeomorphic displacement field using a method similar to [77]. By setting the seed for the random number generator before each experiment, the deformations could be repeated, such that the deformation were the same for each algorithm.

9.4.2 Experiments on four images containing a square

In order to demonstrate the groupwise property of the proposed registration algorithm, the Gravity algorithm was applied to four images (100×100 pixels) each containing a 10 pixels wide square, which was for each image located in a different corner. Gaussian distributed noise was added to all images (Figure 9.5). The images were registered towards each other, resulting in transformations that each map the square to the center of the image. These transformation fields were the basis for further processing.

First, the transformation fields were calculated to map each image to the shape of image I_1.[3] For each image i, $T_{i1} = C(T_i, T_1^{-1})$. All images were mapped to this shape and then averaged. The result in Figure 9.6 shows that the noise is reduced and the square still has sharp edges. It can also be seen that the deformation fields are fully diffeomorphic.

Next, the same method was repeatedly applied, but all images were mapped to a shape in between the original images as explained in section 9.3.1.5. In this way, a series of images was obtained that show the square moving from one corner to the other. This was done using linear and cubic interpolation of the transformations. To illustrate this in a single image, the original squares are shown together with the

[2] http://elastix.isi.uu.nl/

[3] In practice it would likely be more convenient to transform all images to the mean reference frame. For the sake of the example we chose to map all images to match with the first image.

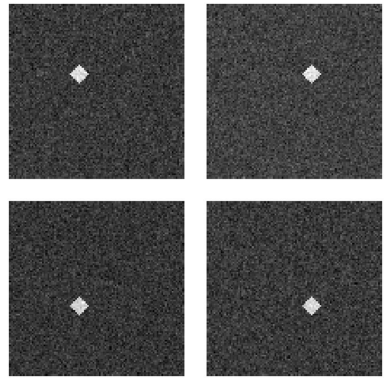

Figure 9.5: Illustration of the four images used to demonstrate the groupwise registration. Each image contains a 10×10 pixel square and additive Gaussian noise with a standard deviation of 10%. The distance between the squares is 10 pixels.

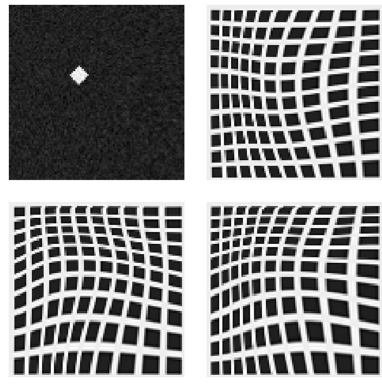

Figure 9.6: Illustration of the deformation fields to transform the images to the shape of the top left image. The top left itself shows the average of all four images. It can be seen how the noise is reduced, while the square has sharp edges.

squares exactly in between (Figure 9.7). It can be seen how the "deformation path" can be made smooth by using cubic interpolation.

9.4.3 Experiments with the Lena image

Experiments were performed on two-dimensional images (256×256 pixels) based on the well-known Lena image (e.g. Figure 9.2).

Random but known deformations were applied to the images and the algorithms were applied to reconstruct this deformation. The image and its deformed counterpart were distorted with 10% Gaussian distributed noise (i.e. the sigma was 10% of the maximum value of the image). A series of 10 image-pairs were used as a training set, and a series of 30 image-pairs was used as a test set.

All algorithms were first tuned using the training set to find the optimal param-

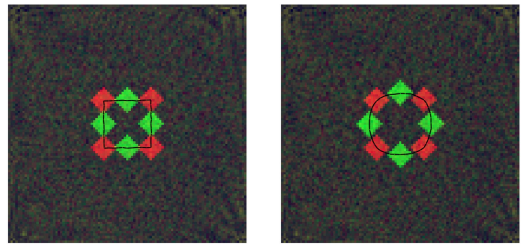

(a) Linear (b) Cubic

Figure 9.7: Illustration of the creation of new images (green) in between the original images (red), by linear and cubic interpolation of the transformations. The black line shows the approximate location of the square's center for 40 locations.

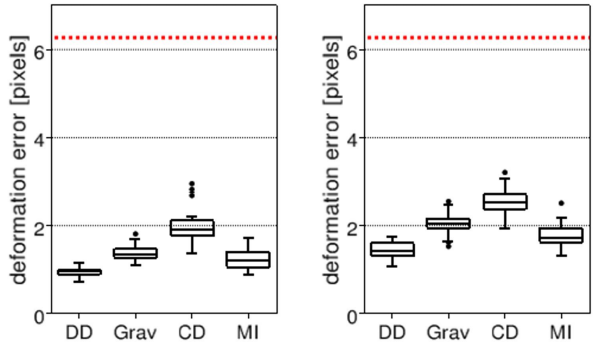

(a) 10% noise (b) 20% noise

Figure 9.8: Illustration of the mean deformation error for registering the Lena image. The dotted horizontal line represents the error when the images are not registered.

eters. The resulting values are shown in Table 9.1. Next, the experiments were performed on the test set. To test the algorithms robustness for noise, the experiments were repeated with 20% Gaussian noise. The results of the experiments are shown in Figure 9.8. Registering a single image took on average 5.3 s and 3.6 s for the diffeomorphic Demons and Gravity algorithms, and 1.7 s and 21 s for the classic Demons and MI-based algorithms.

To evaluate the effect of the differences between the proposed and the classic Demons algorithm, the experiments were repeated (with 10% noise) for four variants of the Demons algorithm: the proposed algorithm (DDemons), the classic algorithm (CDemons), a variant of the proposed algorithm that uses a classic (non-smooth) scale sampling approach (DD_{ns}), and a variant of the classic algorithm that uses the proposed smooth scale sampling approach (CD_s). The results are shown in Figure 9.9.

	#iters per scale level p_{iters}	Grid sampling (smoothing) p_{gs}	Speed factor p_{speed}	Noise factor p_{noise}
Lena image				
DD	16	15	1.5	(1.0)
Grav	16	15	2.0	-
CD	16	4.0	-	1.25
MI	250	15	-	-
Toys image				
DD	16	40	5.0	(1.0)
Grav	16	40	2.0	-
CD	16	5.0	-	1.5
MI	250	40	-	-
T1-T1				
DD	16	15	0.5	(1.0)
Grav	16	15	2.0	-
CD	16	5.0	-	2.5
MI	250	15	-	-
T1-T2				
DD	16	30	0.5	(1.0)
Grav	16	30	2.0	-
CD	16	5.0	-	2.5
MI	250	30	-	-

Table 9.1: List of the parameters used in the experiments to register the different images. For CD (classic Demons) p_{gs} is the sigma if the smoothing kernel.

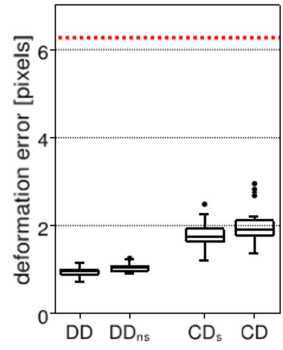

Figure 9.9: Illustration of the mean deformation error (on the Lena image) for different variants of the Demons algorithm. DD_{ns} represents the proposed diffeomorphic Demons algorithm but not with a smooth scale space. CD_{s} represents the classic Demons algorithm with a smooth scale space.

(a) With flash (b) Without flash

Figure 9.10: Illustration of the two photos of a scene taken under different lighting conditions.

9.4.4 Experiments with images of a scene under different lighting conditions

Experiments were performed on photos of a scene under different lighting conditions. The images (400×300 pixels) contain a table on top of which are placed a few toys, and there is a couch in the background. This scene was shot using a regular consumer camera (Canon Powershot) with and without the use of a flash (Figure 9.10).

First, only the image with flash was used. Random but known deformations were applied to this image and the algorithms were applied to reconstruct this deformation. The image and its deformed counterpart were distorted with 10% Gaussian distributed noise. A series of 10 image-pairs were used as a training set, and a series of 30 image-pairs was used as a test set.

All algorithms were first tuned using the training set to find the optimal parameters (Table 9.1). Next, the experiments were performed on the test set. To test the algorithms robustness for lighting conditions, the experiments were repeated, but in this case the image with flash was registered to the deformed image without flash. The results of the experiments are shown in Figure 9.11. Registering a single image took on average 5.3 s and 5.7 s for the proposed Demons and Gravity algorithms, and 2.8 s and 22 s for the classic Demons and MI-based algorithms.

9.4.5 Experiment with single modal MRI data

Experiments were performed on two-dimensional slices extracted from simulated MRI images ($181 \times 217 \times 181$ voxels, spaced 1 mm in all dimensions, with a slice thickness of 1 mm). These images were obtained from the Brainweb database[73][4]. We obtained T1 weighted images with 3% noise (as defined by Brainweb), and an intensity non-uniformity of 0%. For practical reasons, the slices were padded to 256×256 pixels (Figure 9.13a).

Random but known deformations were applied to the images and the algorithms were applied to reconstruct this deformation. A series of 14 slices was used as a training set, and a series of 28 slices (not overlapping with the training set) was used as a test set.

[4] http://www.bic.mni.mcgill.ca/brainweb/

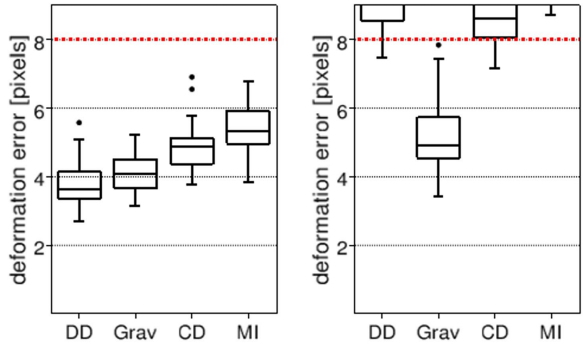

(a) Same lighting (b) Different lighting

Figure 9.11: Illustration of the mean deformation error for registering images of a scene containing toys. The dotted horizontal line represents the error when the images are not registered.

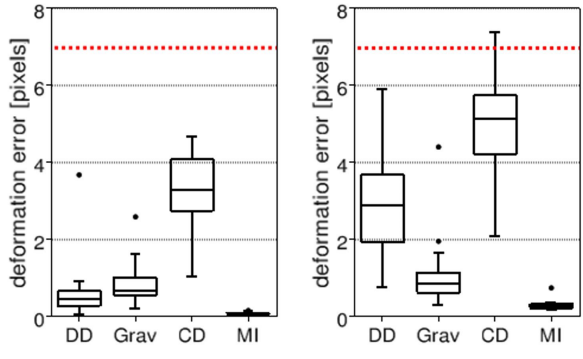

(a) 0% bias field (b) 40% bias field

Figure 9.12: Illustration of the mean deformation error for registering T1 to T1 images. The dotted horizontal line represents the error when the images are not registered.

All algorithms were first tuned using the training set to find the optimal parameters (Table 9.1). Next, the experiments were performed on the test set. To test the algorithms robustness for the bias field, the experiments were repeated, but in this case using a bias field of 40% for the deformed image. The results of the experiments are shown in Figure 9.12. Registering a single image took on average 4.4 s and 3 s for the proposed Demons and Gravity algorithms, and 1.2 s and 20 s for the classical Demons and MI-based algorithms.

9.4.6 Experiment with multi modal MRI data

Experiments were performed on multi modal simulated MRI data. The same data was used as in the previous experiment, but now the task of the algorithms was to

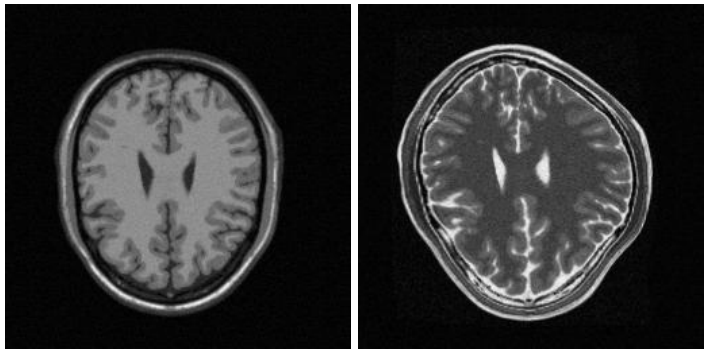

(a) T1 (b) T2

Figure 9.13: Illustration of the a T1 image and the corresponding (but deformed) T2 image.

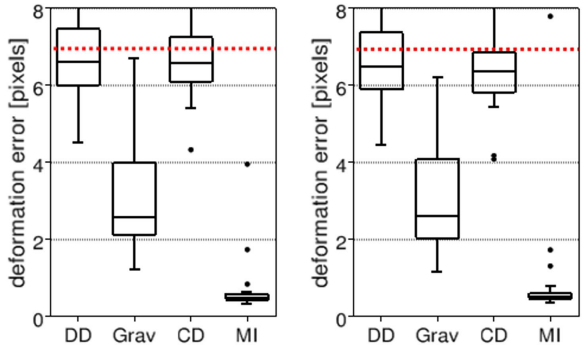

(a) 0% bias field (b) 40% bias field

Figure 9.14: Illustration of the mean deformation error for registering T1 to T2 images. The dotted horizontal line represents the error when the images are not registered.

register a T1 image to a deformed T2 image (Figure 9.13).

All algorithms were first tuned using the training set to find the optimal parameters. The resulting values are shown in Table 9.1. Next the experiments were performed using a bias field of 0% for the T1 and T2 images. Again, the experiments were repeated using a bias field of 0% and 40%, respectively. The results of the experiments are shown in Figure 9.14.

9.5 Discussion

9.5.1 Experiments

Using a set of simulated images containing a white square, we demonstrated the capability of the proposed algorithm to perform groupwise registration. We also showed

how the resulting transformations can be used to map all images to the reference frame, and how a motion sequence can be temporally interpolated.

The results for registering the Lena images (Figure 9.8) show that the proposed variant of the Demons algorithm has the best performance. MI and Gravity registration perform slightly worse, and classic Demons performs the worst. As the noise is increased, this trend is maintained.

The comparison of different variants of the Demons algorithm (on the Lena image) shows that using smooth scale sampling produces slightly better results compared to using factor-of-two scale sampling. This is shown for both the proposed ($p < 0.0002$) and the classic ($p < 0.09$) Demons algorithm.

The results for the experiments with images containing toys (Figure 9.11) show that the proposed algorithms perform comparably well, but that classic Demons and MI seems less capable of accurately aligning this type of images. When the lighting conditions for the two images to register are different, all algorithms suffer greatly, except Gravity registration, which only has a slightly increased deformation error.

The results for the single modal MRI data clearly show that the diffeomorphic Demons and MI algorithms perform the best. Although the Gravity algorithm performs considerably worse, its average error is still smaller than the pixel size. When a bias field is introduced, it can be seen that both Demons-based algorithms perform much worse. This is because Demon forces rely on the similarity between pixel intensities.

The results for the multi modal MRI data show that the accuracy of both Demons algorithms is lower than the other two algorithms. This is as expected, because the intensities differ greatly between the images. Although Gravity registration has an error slightly larger around 3 pixels (which can be considered sufficient for many applications), it is clearly outperformed by the MI-based algorithm. Since the pixel intensities between the images are already different, no significant changes are observed when a bias field is introduced.

Although the experiments were all performed on 2D data, the proposed algorithms can be applied to 3D data as well.

9.5.2 Algorithms

9.5.2.1 Demons

For the proposed Demons algorithm, the parameter p_{noise} was always optimal when set to 1.0, which effectively disables this parameter. This is probably because it is not necessary in combination with the B-spline based regularization; the strength of the delta deformations is regulated by the proposed constraint that keeps the B-spline knot values smaller than $\frac{s}{K}$.

Although we have shown that the usage of smooth scale sampling improves the performance of the Demons algorithm, the major part of the performance improvement of the diffeomorphic Demons algorithm can be attributed to the proposed regularization method based on B-splines. Since both images have to be deformed at each iteration, and because the proposed regularization is more costly than simple Gaussian diffusion, the proposed method is roughly a factor 3 slower than the classic Demons algorithm. This decrease in speed is justified given the large performance increase compared to the classic Demons algorithm.

The experiments show that the diffeomorphic Demons algorithm produces excellent results when the image intensities between the images are similar. Unfortunately, the Demons algorithm is inherently sensitive to distortions such as lighting conditions and bias fields.

9.5.2.2 Gravity

Since Gravity registration is based on attraction forces, it is decoupled from the absolute image intensities, which enables (and requires) normalization of the mass image. Therefore the algorithm is robust for differences in image intensities such as bias fields and image artifacts, and operates well on a wide range of images with little need to change the parameter values (as can be seen from Table 9.1). Its performance in situations with extreme differences in the appearance between the images was already demonstrated in [61]. This also enables using Gravity registration on multi modal images. Unfortunately, in some situations this feature makes the algorithm less accurate than for example the proposed Demons algorithm, which uses the image intensities more directly to calculate the force F_{ij}.

The experiments show that although Gravity registration is often not the most accurate registration algorithm, it has a good overall performance and is robust for different image artifacts, causing it to outperform the other registration algorithms in extreme situations.

9.5.2.3 Mutual Information

In most experiments involving the MI-based algorithm, the performance of this algorithm could be improved by increasing the number of iterations. However, this would make the processing time even longer. In the presented experiments the proposed algorithms are more than five times as fast as the MI-based algorithm, even though they run in a single thread. The number of iterations could probably be reduced for the MI-based algorithm to trade performance for speed.

Except for the experiments with the images containing toys, MI performs very well compared to the other algorithms, and clearly performs best on the multi-modal image data. Unfortunately, the used MI registration algorithm is not diffeomorphic nor groupwise. It would therefore be interesting for future studies to incorporate a force in the proposed algorithm that is based on Mutual Information.

9.6 Conclusions

We propose a registration algorithm that can be applied to two or more images, and which produces transformations that are symmetric and diffeomorphic. This is achieved by representing the transformation as a discrete displacement field and using a B-spline grid to constrain the delta transformations at each iteration to be diffeomorphic.

The algorithm has a generic design; different forces can be used to drive the registration process. In the current work we propose two forces. The first are Demons forces, resulting in a variant of the Demons algorithm that is diffeomorphic and can be used in a groupwise setting. The second are gravity forces, resulting in an algorithm

that is robust for differences in the intensities between the images, making it suitable also for multi modal registration problems.

In quantitative experiments we evaluated both proposed methods against the classic Demons algorithm and an algorithm based on the optimization of mutual information. The results show that the proposed diffeomorphic Demons algorithm has a high accuracy if the pixel values between the images are similar. When distortions are introduced that affect the relation of the pixel values between the images (such as different lighting or bias fields) the performance of both Demons algorithms decreases. The proposed Gravity algorithm is shown to be robust for such distortions.

9.7 Outlook

A major problem with image registration is that the selection of the best registration algorithm (and what parameters should be used) depends greatly on the problem and the type of images. We believe that the proposed generic registration algorithm can be valuable for handling this problem. Firstly, it's regularization method leads to good convergence (as is demonstrated by how much the proposed Demons algorithm has improved compared to classic Demons). Secondly, it allows using a large variety of forces to drive the registration. The force could, for example, be tailored to fit a particular problem, without the need to implement a whole new registration algorithm.

10

A tool for studying the motion of stent grafts in AAA

A paper based on this chapter is in preparation.

Abstract

In this chapter we propose a method to quantitatively measure the stent graft motion in a robust manner. Groupwise image registration is used to find the deformations between the multiple three-dimensional images (Chapter 9). These deformations are then used to incorporate motion to the geometric model of the stent that is obtained by our segmentation algorithm (Chapter 8). The result is a quantitative dynamic model of the stent. Since the geometric model describes the topology of the stent, the forces acting on the stent can be estimated as well.

We performed experiments to compare four different registration algorithms on data containing stent grafts. Further, we calculated the motions in data acquired from 13 clinical cases. The motions show mainly distal motions and expanding motions in the transversal plane. Further, we show that depending on the geometry of the stent, similar motion patterns can result in different force distributions.

These results give a glimpse of the possibilities that this new source of information can provide. The proposed method should therefore be seen as a tool that enables further research to the behavior of stent grafts in vivo.

10.1 Introduction

10.1.1 Clinical context

Endovascular aneurysm repair (EVAR) is an established technique, which uses stent grafts to treat aortic aneurysms in patients at risk of aneurysm rupture [125]. However, due the need for reintervention it does not have a significant advantage over open repair on the long term [53, 27]. Late stent graft failure is therefore a serious complication in endovascular repair of aortic aneurysms [16, 28, 54, 80, 88, 100]. Examples are metal fatigue, stent graft migration [80, 70], and the formation of endoleaks [82, 106].

The long-term durability of stent grafts is affected by the stresses and hemodynamic forces applied to them, which may be reflected by the movements of the stent graft itself during the cardiac cycle. Studying the dynamic behavior of stent grafts can therefore give a better understanding of their motion characteristics, and can give insights into how these motion characteristics relate to certain stent-related problems. This information will be beneficial for designing future devices and can be valuable in predicting stent graft failure in individual patients [74].

Applying ECG-gated CTA [45] provides three-dimensional datasets at different phases of the cardiac cycle. This allows 4D visualization of the scanned object and enables the investigation of its temporal behavior. ECG-gated CT has been used to study the motions of aneurysms [49, 91, 119] and stent grafts [74, 108]. In [62] we demonstrated that ECG gating is a suitable technique for studying the expected motions in the stent graft and vessel wall in abdominal aortic aneurysm (AAA).

10.1.2 Previous work

Most studies on the motion of stent grafts focus on measuring the stent's diameter changes [51] or finding the motion for a sparse set of points on the stent [74]. A model that enables capturing material properties and high level knowledge about the stent graft would be a valuable tool to gain more insight into the stent's in vivo behavior [74]. Such a model can also help in performing more reliable (fluid dynamics) simulations, which is important for improving current stent designs [17, 68].

In [66] we proposed a method to segment the wire frame of the stent. This algorithm can operate on low-dose (i.e. high noise) data, which is important because this reduces the exposure of the patient to ionizing radiation, thereby increasing patient safety on the long term [98, 39]. The algorithm produces a geometric model of the stent, consisting of a graph, where the nodes represent the corners and crossings of the stents frame and the edges represent the actual wires (Figure 10.1). For each edge, the path through the voxels is known as well, allowing for a more precise representation of the stent. This particular graph representation is also known as a spatial graph [6].

One method to apply motion to the obtained geometric model would be to segment the model in each of the different phases. However, this could lead to inconsistencies between the models, which is difficult to deal with. Furthermore, the accuracy of the measured motions would then depend greatly on the accuracy with which the nodes are localized in the individual volumes. This is particularly problematic for motions in AAA, since the data is noisy and the motions of interest are small [62].

A better method is to find the deformation fields between the different phases in an accurate manner, and use these to apply motion to the geometric model. The

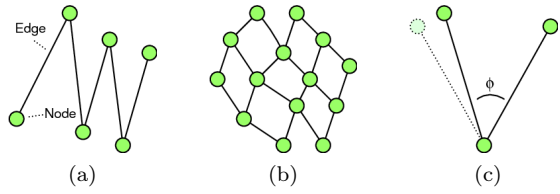

Figure 10.1: Example graphs that describe a geometric model of the stent's frame. The edges between the nodes represent the physical wire frame of the stent. Nodes are placed at corners (a) and crossings (b), which makes it possible to model different stent types. The (change of) angle ϕ when motion is applied to the model can be used to estimate the force present in the node (c).

geometric model can then also be calculated from the averaged phases, which results in a more accurate segmentation[66]. These deformation fields can be found through image registration techniques. Because all the phases must be registered simultaneously, a groupwise registration approach is required. This demands specific features, such as the guarantee for a unique inverse [35, 4](Chapter 9). In previous work we have proposed a generic groupwise diffeomorphic image registration algorithm, which allows using different forces to drive the registration process (Chapter 9).

10.1.3 Contribution

In this work we propose a method to study the motions of stent grafts in AAA, which consists of two parts: segmentation and registration. First, segmentation is used to obtain a geometric model of the stent. Next, groupwise image registration is used to obtain the deformation fields which are used to apply motion to the geometric model. The initial concepts of this idea were briefly discussed in [64].

We compare different image registration algorithms to determine which one is the most suitable for our purpose. Additionally, we show the resulting motions of stent grafts obtained from the data of 13 different patients.

10.2 Algorithms

The segmentation algorithm and the registration algorithm have been discussed in detail in [66] and (Chapter 9)respectively. For completeness, we will now briefly discuss both algorithms, and shall describe any changes with respect to the original methods.

10.2.1 Segmentation

The segmentation algorithm consists of three steps (Figure 10.2). First, seed points are detected by finding the voxels that are subject to a few simple criteria: it must be a local maximum, its intensity must be above a certain threshold value, and it must have a direct neighbor that also has an intensity above the threshold value.

Next, a modified version of the minimum cost path (MCP) method [122] is applied to find connections between the seed points. A front is grown from each seed point,

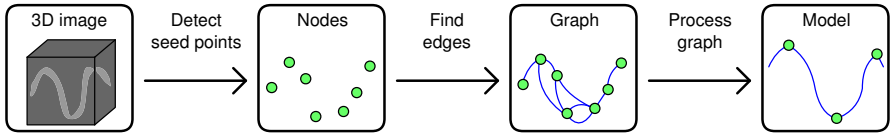

Figure 10.2: Flow chart of the segmentation algorithm.

and two fronts meet, a connection is made and the optimal path between the seed points is found. By choosing an appropriate cost function, the fronts can be made to have a high affinity for the wire of the stent, such that the paths between the points follow the wire frame of the stent graft. The result is a graph (in which the nodes are the seed points) with a lot of connections.

In the third step of a series of graph processing operations are performed to clean up the graph. Redundant edges are removed, and new nodes are placed at corners and crossings. The final result is a graph that describes the wire frame of the stent in a concise manner .

10.2.2 Registration

Groupwise image registration is the process of aligning multiple image with each-other. A large amount of research is published on image registration, and it has been applied to many practical problems. A major difficulty with image registration is that the specific application determines to a large extent what algorithm and parameters work best. Different algorithms can also behave differently when exposed to the same image artifacts.

In the current work, four image registration algorithms are considered. The first two are image driven algorithms based on the generic registration framework proposed in Chapter 9:a diffeomorphic variant of the Demons algorithm (DDemons), and the Gravity registration algorithm. The third and fourth algorithm are algorithms that optimize the mean square error (MSE) (as used in [89]) and mutual information (MI), respectively.[1] Although the latter two algorithms do not produce diffeomorphic deformations (i.e. do not guarantee a unique inverse), this in generally not to be a problem in practice [89]. We chose to incorporate the MI algorithm because of its promising results in a wide range of fields, even though it can currently not be used for groupwise registration.

The Gravity algorithm is based on the idea of masses attracting one another. For this purpose, a mass image is created, which has a sparsely distributed set of "objects" in it. In the original approach this mass image was created by taking the gradient magnitude, such that the algorithm is effectively based on the attraction of edges. However, because the wire of the stent's frame is relatively thin, we propose to use a mass images based on the absolute of the Laplacian instead. This results in masses on the wire instead of around them, and therefore causes less "cross-talk". The mass is calculated as follows:

$$L_{\text{mass}} = |L_{xx} + L_{yy} + L_{zz}|, \tag{10.1}$$

where L_{xx}, L_{yy} and L_{zz} are the second order derivatives for x, y and z, respectively.

[1] For the MSE and MI algorithms we used the implementation that is freely available as part of the Elastix toolkit: http://elastix.isi.uu.nl/

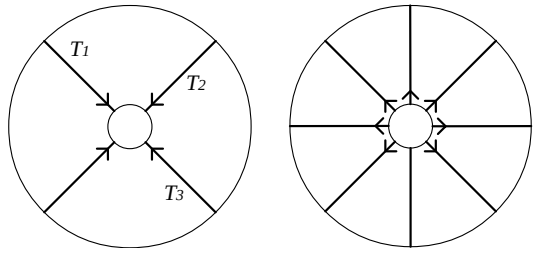

(a) From image space to reference frame

(b) From reference frame to image space

Figure 10.3: Schematic representation of the deformations in groupwise image registration. The outer circle represents the image space, an the inner circle the reference frame. The deformations produced by a registration algorithm represent the mapping from image space to reference frame (a), and can be used to describe a (continuous) mapping to the image space (b).

10.2.3 Applying motion to the model

The deformation fields found by a registration algorithm describe how each image should be deformed to make it match the shape of the reference shape, which is in general the average shape of the original images (Figure 10.3a). Similarly, the deformations can be used to describe how the reference frame should be deformed to match the shape of any of the original images. By interpolating the deformation fields (Chapter 9) it is also possible to map the reference frame in between the original images (Figure 10.3b). In other words, the deformation fields describe the motion of the reference frame as a continuous function. This property can be of value in applying calculations to the motion field, and is allows the motion to be visualized in a smooth fashion.

In order to use this information for quantitatively studying stent graft behavior, the motion at the position of the stent (at the nodes of the geometric model) is calculated from the deformation fields. This makes it possible to visualize the geometric model in motion, and enables performing calculations on the now dynamic geometric model. As the basis for estimating the forces that act on the frame of the stent, the change of angle between two edges in the model can be calculate (Figure 10.1c).

10.3 Experimental methods and materials

In any experiment there is a trade-off between realism and availability of a (reliable) ground truth. Therefore we performed a series of experiments on different kinds of data: 1) artificial 2D images with artificial deformations; 2) data obtained by scanning a phantom; 3) clinical data. In the first two experiments we tested the performance of four different registration algorithms, and distortions were introduced to test the robustness of these algorithms. In the third experiment we tested the proposed algorithms by applying it to patient data.

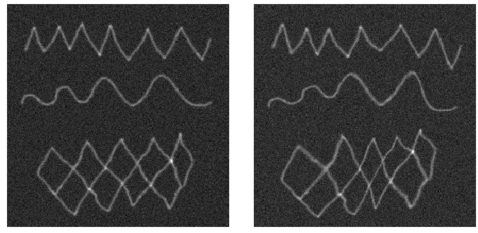

Figure 10.4: Example of the image used in the 2D experiments. The shown image has 10% noise and the right image is deformed ($d_{\text{amp}} = 20$).

10.3.1 Experiments on 2D simulated data

An artificial image (256×256 pixels) was created containing lines in zigzag and diamond shapes similar to some stent types (Figure 10.4). Random but known deformations were applied to the images and the algorithms were applied to reconstruct this deformation. The image and its deformed counterpart were distorted with 10% zero-mean Gaussian distributed noise (i.e. the sigma of the Gaussian kernel was 10% of the maximum value of the image). A series of 10 image-pairs was used as a training set, and a series of 30 image-pairs was used as a test set.

The random deformations were calculated by randomly sampling 50 points inside the image. From each point, a random vector was generated with a maximum length of d_{amp} pixels. The value d_{amp} is a measure of deformation amplitude, and was varied in the experiments. This description was then converted to a (diffeomorphic) displacement field using a method similar to [77]. By setting the seed for the random number generator before each experiment, the deformations could be repeated, such that they were the same for each algorithm.

All algorithms were first tuned using the training set (with $d_{\text{amp}} = 10$) to find the optimal parameters. Next, the experiments were performed on the test set, with d_{amp} set to 5, 10 and 20. These experiment were repeated three times with different distortions. For the first distortion we doubled the amount of Gaussian noise. For the second distortion the intensities of the wire in the deformed image were decreased with a factor $\frac{1}{2}$. For the third distortion a grid of horizontal and vertical lines (1 pixel in width and positioned at every 6th pixel) was added to the deformed image. These distortions have similarities with CT image artifacts such as noise, the partial volume effect and streak artifacts, and were introduced to test the robustness of the algorithms.

10.3.2 Experiments on 3D phantom data

In order to obtain more realistic data with some form of a ground truth, an aorta phantom was developed, consisting of a cylindrical perspex container, inside which a rubber tube is fixated (Figure 10.5a). A stent was placed inside the rubber tube, which has a larger diameter than that of the tube, such that in its natural state, the tube compressed the stent. The container and the tube can both be filled with water, but water cannot escape from the inside of the tube. This enables injecting extra water inside the tube such that it inflates, allowing the stent diameter to increase. Further, a perspex bar could be inserted via a hole at the top of the container, such

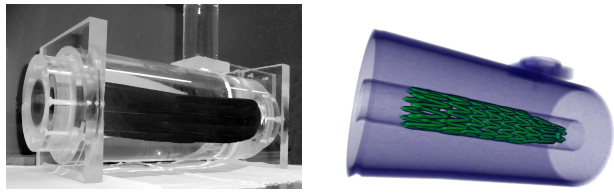

(a) (b)

Figure 10.5: Example of the phantom used in the 3D experiments. Shown are a photograph of the phantom (a) and a volume rendering of a scan containing the phantom (b).

that the tube could be displaced (and thus deformed). The phantom was emerged in a water bath with a width of approximately 50 cm.

The phantom was scanned using a Siemens Somatom 64-slice CT scanner (Siemens Medical Solutions, Erlangen, Germany). A rotation time of 0.37 s, a pitch of 0.34, and $2 \times 32 \times 0.6$ mm 2 collimation were used. The tube voltage was 120 kVp. Each volume was reconstructed using the B36f 245 reconstruction filter and resulted in approximately 220 slices of 512×512 voxels. The slices (with a thickness of 2 mm) were spaced 1 mm apart, and the resolution in the xy plane was approximately 0.5 mm. The resulting data were manually cropped to $200 \times 256 \times 256$ voxels to reduce the memory requirements.

The setup was scanned in nine different configurations, by scanning all possibilities of three different pressures on the tube and three different bar displacements. For each configuration, two scan were made with a tube current of 40 mAs and 600 mAs, respectively. Special care was taken to prevent any displacement of the phantom in between the different scans. The low exposure scan functions as a realistic scan to perform the algorithms on. The high exposure scan functions as a reference to evaluate the performance. Another three configurations were scanned with the introduction of a metal bar in different positions to introduce metal artifacts.

The registration algorithms were used to deform one configuration (no inflation, no bar) to the other configurations, using the scans of 40 mAs.

10.3.3 Experiments on 3D patient data

10.3.3.1 Materials

For the final experiment we used data obtained from 13 clinical cases (of 13 different patients). The patients were threated with either the AneuRx (Medtronic, Minneapolis, USA) or the Zenith (Cook, Bloomington, USA) stent graft. ECG-gated CT data was used to obtain ten volumetric images, using the same scanner and settings as in the phantom experiment. The effective tube current (per rotation) time product was 180 mAs. Each volume resulted in approximately 300 slices of 512×512 voxels. Retrospective gating was applied to obtain ten (equal distant) cardiac phases. The resulting data were manually cropped to $256 \times 256 \times 256$ voxels to reduce the memory requirements.

Next, the Gravity algorithm was used to register the different phases towards eachother. This resulted for each dataset in 10 deformation fields that describe for each

phase how it should be deformed to the reference frame. The deformation fields were applied to the geometric models obtained with the segmentation algorithm. Next, the maximum change of the angle between the edges was calculated for all nodes in the model.

The dynamic model of the stent can be visualized interactively; the user can inspect the model while it moves, using the mouse to zoom and view the model from different directions. This allows the viewer to get a good understanding of the geometry of the stent and its motion. The maximum angle change is displayed using color.

10.4 Results

The tuning of the algorithm, as performed using the artificial 2D imaged and the 3D phantom data, resulted in the set of parameters shown in Table 10.1.

For the 2D images registering one image pair took on average 2.1 s for the DDemons and Gravity algorithms, and respectively 7.4 s and 12.0 s for the MSE and MI algorithms. For the 3D images, registration took 20 min for DDemons, 7.5 min for Gravity, 3.1 min for MSE, and 3.6 min for MI. For the clinical data, groupwise registration took around 5 hours per case.

10.4.1 2D artificial data

To quantitatively evaluate the results of the experiments on the 2D data, the error between the applied deformation and the found deformation was calculated. Because we are only interested in the deformation at the lines, the error was evaluated at the lines by masking it with the noiseless image. The results are shown in Figure 10.6.

10.4.2 3D phantom data

For the phantom experiments the quality of the deformation was measured by applying the found transformations to the corresponding scans of 600 mAs, and subtracting the results. The sum of the absolute intensity error was divided by the sum of the absolute intensity errors obtained by subtracting the un-registered scans. The results are shown in Figure 10.7.

Because the metal bar (and the subsequent artifacts) causes large intensity differences between the images even if they would be perfectly aligned, these registration result cannot be quantitatively evaluated with the above approach. Instead, the results were evaluated qualitatively; Figure 10.8 shows the maximum intensity projections of the average of the two registered images. A good registration results in a sharp image of the stent, while registration errors causes the two (misaligned) stents to be visible.

10.4.3 Motion in 3D patient data

Since there is no ground truth available for the clinical data, the registration results can only be visually inspected. In Figure 10.9 a slice is shown from a volume that was obtained by registering the 10 phases of a dataset and then averaging them. By

comparing the result to a slice from the volume obtained by only averaging the phases, it can be seen that the motion artifacts are removed.

To show the motion in a 2D image, we divided the motion cycle in four phases and show the motion vectors using lines overlaid on the geometric model. [2] Additionally, we show an image of the geometric model with the maximum angle change displayed using color. The results for eight datasets are shown in Figure 10.10 and Figure 10.11.

For reasons of space we show a selection of the datasets in the current work. Of the cases not shown, there was one with no motion at all (case 4), two cases (case 6 and 8) that showed mainly distal motion, and two cases (case 7 and 9) that showed an expanding motion.

In case 1 (Figure 10.10a) the body of the Zenith stent first expands, which seems to pull the top of the stent in the second phase. The stent relaxes to its resting state in the last two (diastolic) phases. Case 2 also shows a downward and expanding motion. In fact, case 1 and case 2 are the same patient scanned at different times scanned 7 months apart. It can be seen that the motion pattern has changed slightly. More interesting is that the forces acting on the stent seem to have reduced considerably. The motion for case 3 is mainly in the lateral direction, resulting in the generation of only small forces as the stent seems to move as a whole. For case 5 the small motions result in a slight expanding of the stent (the movie shows this more clearly than the 2D motion image).

For the AneuRx stent graft in case 10 (Figure 10.11a) we can observe a relatively complex motion consisting of motions in the transverse direction as well as in the distal direction. This causes the stent to bend, resulting in local forces on the stent. The motion for case 11 is mainly in the lateral direction, but there are local variations, which seem to cause large forces at the upper part of the stent. In case 12 the motion is mainly in the distal direction, but the body of the stent has a lateral motion as well. The motion causes forces in the bend part of the stent. Case 13 shows mainly a distal motion, but also contains large irregular motions, causing large angular changes in different regions of the stent.

It can be seen that the motion of case 13 shows extreme dynamics compared to the other datasets. This result is caused by noise bands that propagate through the data (Figure 10.12), caused by the heart beat of this patient being too low during the scan [62].

10.5 Discussion

10.5.1 Registration accuracy

From the results of Figure 10.6 it can be seen that for the synthetic data all algorithms are capable of reconstructing the deformation with a mean error well below one pixel. When the noise is increased, the error increases, but the relationship between the algorithms remains approximately the same. When the intensity artifact was introduced, this affected mainly the accuracy of the DDemons and MSE algorithms. The lines artifact negatively affects all algorithms, although MI is clearly least affected and DDemons the most. This can be explained by the fact that both the DDemons and MSE algorithm operate on the intensities directly, making them more vulnerable to

[2] Movies that show the dynamics of the stent grafts in a more intuitive fashion are available at http://www.sas.el.utwente.nl/home/almar/stents.html.

	#iters per scale level p_{iters}	Grid sampling (smoothing) p_{gs}	Speed factor p_{speed}	Number of samples p_{samples}
2D images				
DDemons	16	20	2.0	-
Gravity	16	20	2.0	-
MSE	200	20	-	2^{11} (2048)
MI	200	20	-	2^{11} (2048)
3D images				
DDemons	16	20	2.0	-
Gravity	16	20	2.0	-
MSE	200	20	-	2^{13}(8192)
MI	200	20	-	2^{13}(8192)

Table 10.1: List of the parameters used in the experiments to register the different types of images. See (Chapter 9)and [67] for an explanation of the parameters.

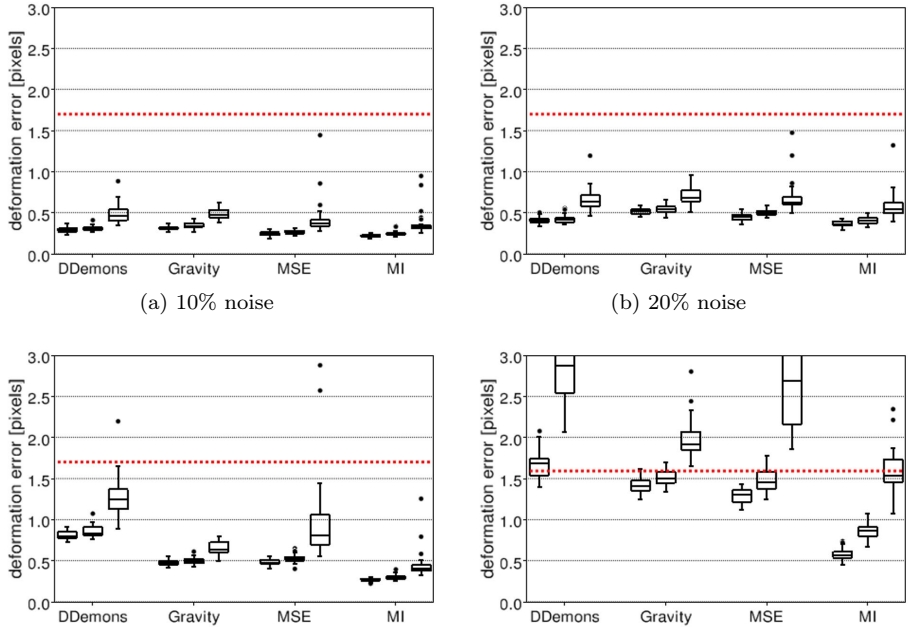

(a) 10% noise

(b) 20% noise

(c) intensity artifact

(d) grid artifact

Figure 10.6: Illustration of the mean deformation error for registering the artificial 2D images containing lines. The results are shown for a d_{amp} of 5, 10 and 20 pixels. The dotted horizontal lines represent the error when the images are not registered (for $d_{\text{amp}} = 5$).

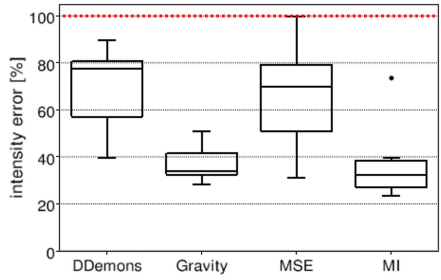

Figure 10.7: Illustration of the intensity error for registering the3D images with the aorta phantom.

distortions. The Gravity and MI algorithms use the image intensities in a less direct manner, which makes them more robust for variation in the intensities between the images (Chapter 9). The MI algorithm is particularly robust, probably because it uses an optimization strategy rather than using the pixel data to drive the registration.

The reliability of the ground truth for the phantom experiments is subject to differences in the reference images caused by image artifacts. By scanning the reference images with a high dose, this effect was minimized. From Figure 10.7 it can be seen that on the phantom data the DDemons and MSE algorithms perform significantly worse than the Gravity and MI algorithms. Similarly, Figure 10.8 shows that all algorithms suffer from the introduction of a metal bar, but that the Gravity and MI algorithm are much less affected than the other algorithms. Both results show that MI performs better than the Gravity algorithm.

We conclude that of the four registration algorithms tested the Gravity and MI algorithm perform sufficiently for our needs. The results show that MI performs better than the Gravity algorithm, but unfortunately a groupwise registration algorithm based on MI is currently not available. The reason for this might be that MI is fundamentally pairwise in nature [87]. With the groupwise framework proposed in this work, however, this should not be a problem. It would therefore be interesting for future research to realize a groupwise MI-based algorithm. We expect that other registration approaches, such as mutual information combined with gradient data (e.g. [95]) may produce good results as well.

10.5.2 Speed

We found that the MSE and MI algorithm perform slower on 2D data, but faster on 3D data. This is because these algorithms use a sampling strategy to avoid calculating the similarity measure for every element (i.e. pixel/voxel). The number of samples had to be increased in the 3D case to get reasonable results (last row of Table 10.1), but this increase was much less than the increase in the amount of elements.

Since the DDemons and Gravity algorithms can be parallelized well, we expect that their speed can be improved significantly by using a multi-threaded approach or using GPU-based techniques. This is subject for future research.

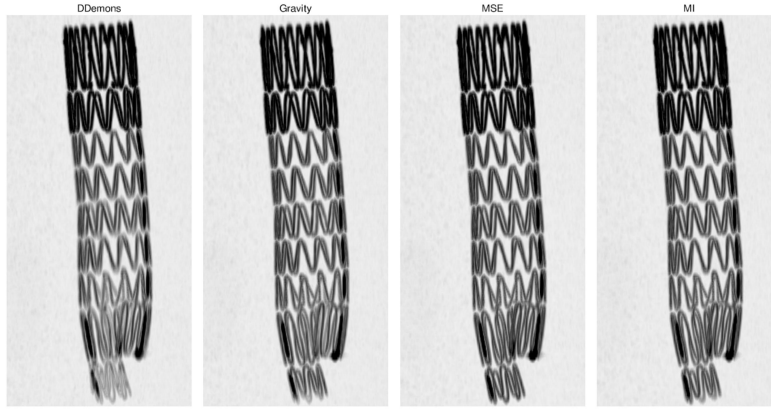

(a) No metal bar

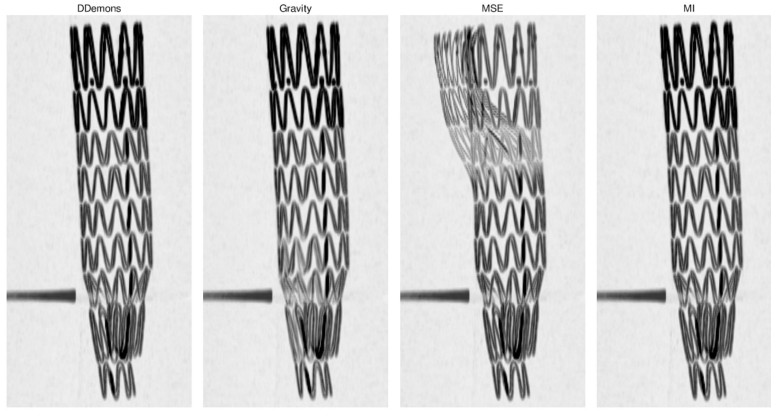

(b) Metal bar perpendicular to the stent

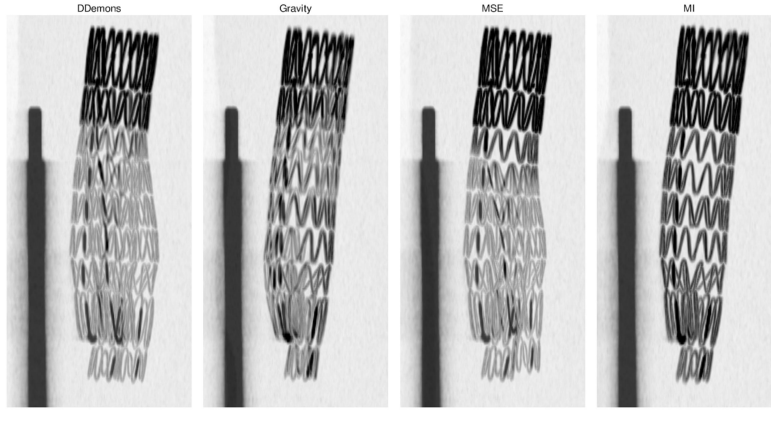

(c) Metal bar parallel to the stent

Figure 10.8: Maximum intensity projections of registered data under the influence of a strong metal artifact.

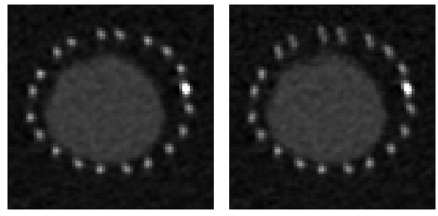

Figure 10.9: Comparison of averaging the ten volumes using registration (left) and without (right). Motion artifacts are clearly visible in the right image.

10.5.3 Deformation amplitude and grid spacing

Every image registration algorithm uses a form of regularization, and consequently, there is always a trade-off between this amount of regularization and the complexity of the deformations that can be modeled. For the four registration algorithms described in this work, the regularization is controlled by the grid sampling parameter p_{gs}, which specifies the spacing between the knots of the final B-spline grid [113, 101]. A smaller value allows more variation in the deformation, at the cost of decreased regularization. It can be seen from the 2D experiment results (Figure 10.6) that the error increases as the deformation amplitude is increased. The error for $d_{\text{amp}} = 20$ is particularly large, which can be attributed to the fact that the algorithms were trained with $d_{\text{amp}} = 10$; the deformations were too complex to be appropriately modeled with the used grid sampling. To reduce this error, the grid sampling will have to be decreased, but due to the reduced regularization, this would cause the error for the smaller amplitudes to increase.

This illustrates the importance of choosing the right amount of regularization; it should be chosen as high as possible, but not too high, since this would constrain the possible transformation too much. To chose a suitable value, one can estimate it from the expected range of deformations. Unfortunately these are often unknown in practice. In the phantom experiments we could quantitatively determine that the grid spacing value was optimal at 20 mm, and this worked for the patient data as well. When we tried to reduce the value to 10 mm we found that the deformations became inaccurate because they were not sufficiently constrained.

The effects of the amount of regularization with respect to the kind of deformations found in patient data should be investigated further.

10.5.4 Volume gaps

The erroneous extreme dynamics seen in case 14 are caused by "volume gaps" where no data is defined for a particular space and time [62]. The scanner attempts to restore the data via a form of interpolation, which causes the stent to be represented in a wrong way. Because the registration algorithm tries to match the shape of this wrongly represented data, the deformation fields contain locally extreme deformations.

Therefore, for the proposed method to work well it is crucial that these volume gaps are avoided by preventing the patient's heart beat to fall below the crucial value (see [62] how this can can be calculated).

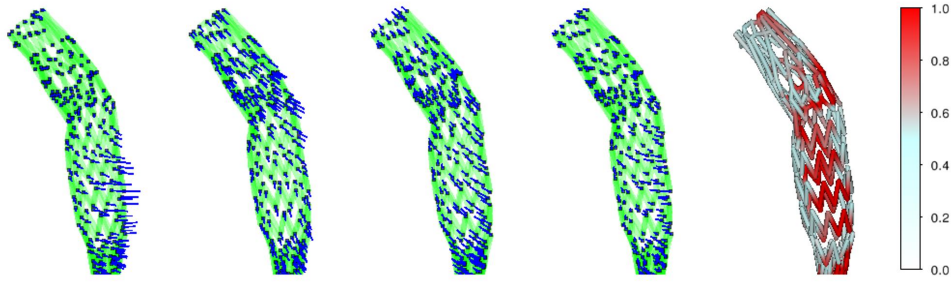

(a) case 1

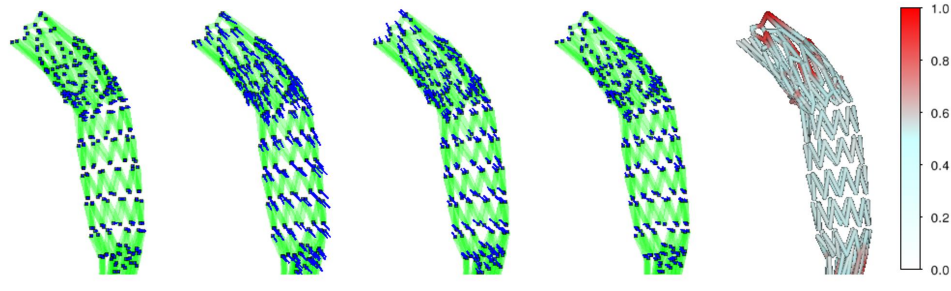

(b) case 2

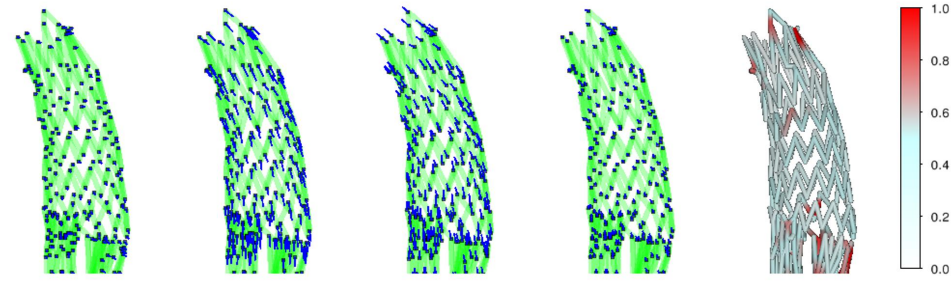

(c) case 3

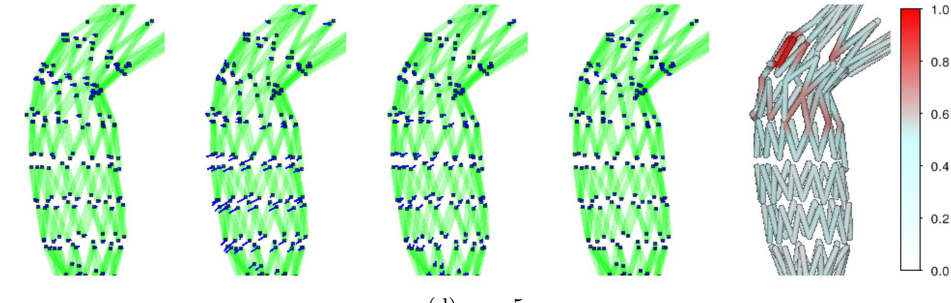

(d) case 5

Figure 10.10: Illustration of stent motion measured for 4 cases treated with the Zenith stent graft. The motion is represented using vectors (amplified five times) in four phases. The image on the right shows the change of angle between the edges (in degrees).

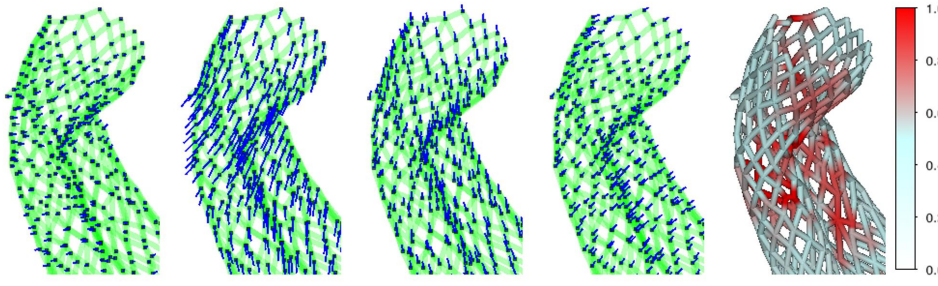

(a) case 10

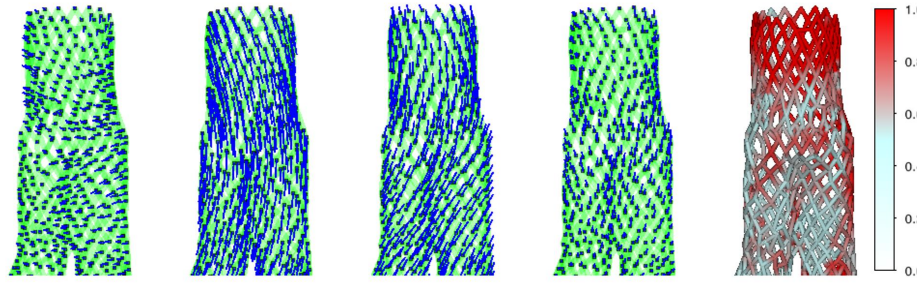

(b) case 11

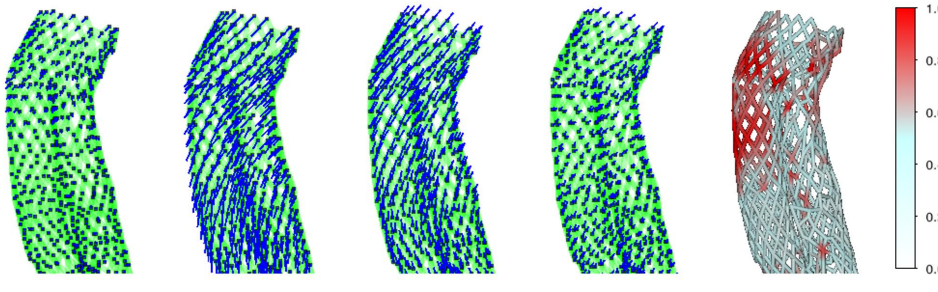

(c) case 12

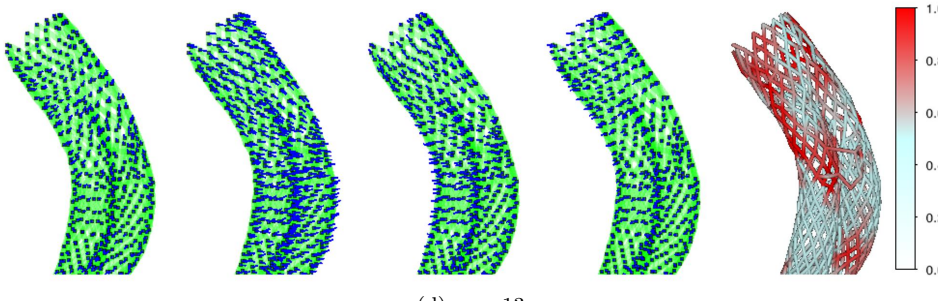

(d) case 13

Figure 10.11: Illustration of stent motion measured for 4 cases treated with the AneuRx stent graft. The motion is represented using vectors (amplified five times) in four phases. The image on the right shows the change of angle between the edges (in degrees).

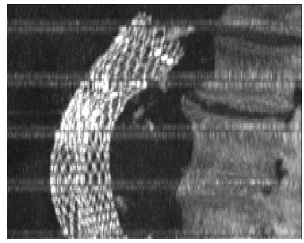

Figure 10.12: Illustration of the noise bands present in the data for case 13.

10.5.5 Patient data

The general tendency observed for the motions in the different patients is a distal and or lateral motion, sometimes accompanied by an expanding body, followed by a slow returning to the resting state (diastolic phase). Depending on the geometry of the stent, the motion can result in forces, mainly at bends and the bifurcation (e.g. case 2, 5, 10, 12). Nevertheless, there are also rather large differences between the observed motions; some datasets have very little or no motion, some are mainly in the distal direction, and some show a more expanding motion in the transversal plane. To what extent these differences are significant for the prediction stent related problem remains to investigated.

The calculated angle changes are an indication of the force present at that location. For genuine force calculations, one would have to take into account the length and diameter of the wires, as well as the material properties. We also observed that in cases where the segmentation contained an error (e.g. an extra node) the angular change was sometimes very high. Because this occurs very locally (see for example the top left of case 5) these errors can usually be recognized. Nevertheless, the effect of errors in the segmentation on the force calculations should be investigated as well.

The used registration algorithm was carefully selected and evaluated for data containing stent grafts. Nevertheless, the resulting deformation fields are an approximation and may contain errors. Further research is required to study (and possibly improve) the robustness of the registration algorithm.

10.6 Conclusions and future work

We performed experiments to compare four different registration algorithms on data containing stent grafts. Further, we used the best method to find the motions in data acquired from 13 different cases. By applying these motions to the geometric model of the stent, a dynamic model was obtained which enables visualization and calculations of the motions of the stent. The motions show mainly distal motions and expanding motions in the transversal plane. Additionally, we showed the change of the angles between the edges in the model, which are an indication of the forces acting on the stent. These show that depending on the geometry of the stent, similar motion patterns can result in different force distributions.

These results give a glimpse of the possibilities that this new source of information can provide. The proposed method should therefore be seen as a tool that enables further research to the behavior of stent grafts in vivo. It could, for instance, enable

studying the motion patterns of individual patients, relate them to data of a previous date, or relate them to the motion patterns of other patients. It would also be interesting to study the range of motion patterns of stent grafts in patients without problems, and compare them to the motions in patients who do have problems. Such studies would, however, require large datasets to incorporate all the variabilities in motion patterns, particularly because problems with stent grafts are relatively rare. Nevertheless, we believe that such studies can help our understanding of the dynamics and failure of stent grafts, and can thereby help in designing better stent grafts in the future. Further, we hope that we are able to correlate certain distinct motion patterns to specific stent-related problems, so that this technique can be used for diagnostic purposes and prediction of stent failure.

11 Conclusions

This thesis presents a method that enables quantitative measurements of the motions of stent grafts from ECG-gated CT data. The proposed method consists of two parts. In the segmentation part the stent is detected from the data and a geometric model is produced that explicitly describes the topology of the stent. In the registration part the deformation field of the data is calculated and used to apply motion to the geometric model. Further, the applicability of ECG-gated CT for measuring the motions of stent grafts in AAA was studied.

11.1 Conclusions

We shall now answer the research questions that were posed in the introduction.

Is the data obtained with ECG-gated CT suitable for measuring the motions expected for stent grafts in AAA?

In Chapter 3 we performed several experiments to investigate the various aspects of this question. We have found that we can measure motions with frequency components up to 2.7 Hz, and with amplitudes from 0.4 mm and 0.7 mm in x/y-direction and z-direction, respectively. This is (just) sufficient for the intended purpose. It was also found that volume gaps occur if the patient's heart rate is below 55 BPM. The data inside these volume gaps is unreliable, and can cause severe misregistration (Chapter 10). This problem should therefore be avoided by preventing the patient's heartbeat to go below 55 BPM if possible.

Can ECG-gated CT replace the regular CT study that patients currently have, without any negative effects on the clinical procedure?

Averaging of the multiple volumes obtained by ECG-gated CT results in a volume with frequency and noise properties that are very similar to those of the single volume

obtained by a regular CT scan (Chapter 4). It was also shown that the artifacts due to motion appear as blurring.

It should be noted that an ECG-gated CT scan requires the patient to hold his/her breath longer than for a regular CT scan. If the patient is unable to hold his/her breath that long, the exhaling might cause motion artifacts. By scanning in distal direction, this effect can be minimized.

Can we segment the stent graft from these (noisy) data, and with what accuracy?

Yes, although this problem proofed to be harder than was anticipated. Our first approach was to segment the stent in 2D slices while tracking along the centerline of the stent. Although segmentation of the stent in the 2D slices was quite successful (Chapter 5), the subsequent tracking turned out to be difficult due to the fundamental limitations of the approach (Chapter 6). To overcome these problems we have worked on a method that works directly in 3D. The second approach consists of a method in which many wire-parts are tracked simultaneously (Chapter 7). This approach was found to work the best using the MCP method for the tracking. Because we anticipated problems with connecting the different wire-parts being tracked, the third (final) approach uses the MCP method to connect a set of seed points (Chapter 8).

The result of the segmentation algorithm is a geometric model in the form of a graph consisting of nodes connected by edges. The nodes are placed at the corners and crossings of the stent's frame and the edges represent the wires.

We have demonstrated (in Chapter 8) that the final method has a 95% and 92% correspondence with expert annotations for the AneuRx and Zenith stent type, respectively. It should be noted, however, that because annotating bifurcations is very difficult, the performance was only measured at the tubular parts of the stents.

Is it possible to measure the motion of the stent graft from these data, and with what accuracy?

Yes, by applying image registration the deformation field between the images can be found. This can then be used to obtain a dynamic stent model by incorporating the motion in the geometric model. Since we are dealing with multiple images, it is of importance to use an algorithm that registers the images simultaneously (i.e. groupwise). Therefore it is crucial that the deformations are invertible, which can be guaranteed by ensuring that the deformations are diffeomorphic. We have designed a registration algorithm that has these properties, and which is generic enough to allow different "forces" to drive the registration progress (Chapter 9). This allows us to select a force which is robust for the image artifacts present in CT data, and also produces accurate results—which is important because the motions of interest are relatively small.

A ground truth for the motion of the stent grafts is not available, making quantitative analysis of the accuracy difficult. Nevertheless, we were able to compare different registration algorithms and select the one with the best accuracy and robustness. The experiments suggest that the proposed Gravity algorithm is sufficient for our needs.

What kind of motions do stent grafts make inside the human body?

For the patients that participate in this study we are able to quantitatively measure the motions of the stent graft (Chapter 10). We have found that the general tendency is a distal motion and/or expanding motion, followed by a slow return to the resting state (diastolic phase). Additionally, we show the change of the angles between the edges in the stent model, which are an indication of the forces acting on the stent. These show that depending on the geometry of the stent, similar motion patterns can result in different force distributions. There are also rather large differences between the observed motions; some datasets have very little or no motion, some are mainly in the distal direction, and some show a more expanding motion in the transverse plane.

To be able to get real insight in these motions, and to draw relations with stent-related problems, much more patients should be studied. We now have the tools to do this in a highly automated manner.

11.2 Recommendations

Using the registered volumes to create an average volume

A single volume is currently generated by averaging the multiple 3D images. This volume is then used in the clinical procedure. Similarly, a single volume is created to which the segmentation algorithm is applied. In the latter case, only a subset of the volumes is used, because that proofed to be a better compromise between noise and motion artifacts (Chapter 7).

By registering the different volumes to a mean shape before averaging (as discussed in Chapter 9), we expect to remove the motion artifacts almost completely. This is beneficial for the clinic, and might also improve the segmentation results.

Improvements for the segmentation algorithm

Although the segmentation algorithm produces relatively accurate results, the errors that are made can have a significant effect on the estimation of forces. To what extent this is a problem should be investigated.

We believe that the most significant improvements for the current segmentation algorithm can be achieved by improving the graph processing. Although the rules for cleaning up the graph are relatively clear, the order in which these are applied can affect the outcome. Instead of doing the processing in a fixed amount of steps, it might be better to adopt a more gradual approach by repeatedly applying a set of operations, where each operation has only an infinitesimally chance of introducing an error.

Improvements for the registration algorithm

In Chapter 10 we also show that an algorithm that uses mutual information might provide better results than the Gravity algorithm, but unfortunately this algorithm can currently not be used in a groupwise setting. It is therefore interesting to implement a registration-force based on mutual information that can be used with the generic registration algorithm proposed in Chapter 9.

Although the parameters of the registration algorithm can be tuned in a controlled experiment, this is often not the case for real-world data. In Chapter 10 we have tried to solve this by tuning the parameters in a series of experiments of increasing realism. It would be interesting to see the effect of the parameters on the estimated motion in patient data, particularly for the "grid spacing" parameter, as it limits the complexity of the motions that can be modeled.

Further, the speed of the registration algorithm can be improved by implementing it on the GPU. The Demons and Gravity algorithms are expected to be particularly easy to parallelize.

Understanding stent dynamics

The proposed method should be seen as a tool that enables further research to the behavior of stent grafts in vivo. It can, for instance, enable studying the motion patterns of individual patients, relate them to data of a previous date, or relate them to the motion patterns of other patients. It is also interesting to study the range of motion of failing stent grafts, and compare them to the motions in patients without problems. Such studies, however, should be performed on large patient populations in order to incorporate all the variabilities in motion patterns, particularly because problems with stent grafts are relatively rare. Nevertheless, we believe that such studies can help our understanding of the dynamics and failure of stent grafts, and can thereby help in designing better stent grafts in the future. Further, we hope that we are able to correlate certain distinct motion patterns to specific stent-related problems, so that this technique can be used for diagnostic purposes and prediction of stent failure.

Bibliography

[1] Stephan Achenbach. Half-Scan vs. Multi-Segment reconstruction for computed tomography coronary angiography. Technical report, Siemens medical, December 2006.

[2] Jörg Albers, Jan M Boese, Christian F Vahl, and Siegfried Hagl. In vivo validation of cardiac spiral computed tomography using retrospective gating. *The Annals of Thoracic Surgery*, 75:885—889, 2003.

[3] Vincent Arsigny, Olivier Commowick, Xavier Pennec, and Nicholas Ayache. A Log-Euclidean framework for statistics on diffeomorphisms. In *Proceedings of MICCAI*, volume 4190, pages 924–931. Springer Berlin Heidelberg, Berlin, Heidelberg, 2006.

[4] John Ashburner. A fast diffeomorphic image registration algorithm. *NeuroImage*, 38(1):95–113, October 2007.

[5] Franz Aurenhammer, Rolf Klein, J.-R. Sack, and J. Urrutia. Voronoi diagrams. In *Handbook of Computational Geometry*, pages 201–290. North-Holland, Amsterdam, 2000.

[6] Stephen R. Aylward, Julien Jomier, Christelle Vivert, Vincent LeDigarcher, and Elizabeth Bullitt. Spatial graphs for intra-cranial vascular network characterization, generation, and discrimination. In James S. Duncan and Guido Gerig, editors, *Proceedings of MICCAI*, volume 3749, pages 59–66. Springer Berlin Heidelberg, Berlin, Heidelberg, 2005.

[7] Serdar K Balci, Polina Golland, Martha Shenton, and William M Wells. Free-Form b-spline deformation model for groupwise registration. *Proceedings of MICCAI*, 10(WS):23–30, 2007.

[8] Thomas Beck, Christina Biermann, Dominik Fritz, Rudiger Dillmann, Josien P. W. Pluim, and Benoit M. Dawant. Robust model-based centerline extraction of vessels in CTA data. In *Proc. of SPIE*, volume 7259, page 72593O, Lake Buena Vista, FL, USA, February 2009.

[9] M. Faisal Beg, Michael I. Miller, Alain Trouvé, and Laurent Younes. Computing large deformation metric mappings via geodesic flows of diffeomorphisms. *International Journal of Computer Vision*, 61(2):139–157, February 2005.

[10] Kanwal K. Bhatia, Paul Aljabar, James P. Boardman, Latha Srinivasan, Maria Murgasova, Serena J. Counsell, Mary A. Rutherford, Joseph V. Hajnal, A. David Edwards, and Daniel Rueckert. Groupwise combined segmentation and registration for atlas construction. In Nicholas Ayache, Sébastien Ourselin, and

Anthony Maeder, editors, *Medical Image Computing and Computer-Assisted Intervention – MICCAI 2007*, volume 4791, pages 532–540. Springer Berlin Heidelberg, Berlin, Heidelberg, 2007.

[11] K.K. Bhatia, J.V. Hajnal, B.K. Puri, A.D. Edwards, and D. Rueckert. Consistent groupwise non-rigid registration for atlas construction. In *Biomedical Imaging: Nano to Macro, 2004. IEEE International Symposium on*, pages 908–911 Vol. 1, 2004.

[12] Jan D. Blankensteijn, Sjors E.C.A. de Jong, Monique Prinssen, Arie C. van der Ham, Jaap Buth, Steven M.M. van Sterkenburg, Hence J.M. Verhagen, Erik Buskens, Diederick E. Grobbee, and the Dutch Randomized Endovascular Aneurysm Management (DREAM) Trial Group. Two-Year outcomes after conventional or endovascular repair of abdominal aortic aneurysms. *N Engl J Med*, 352(23):2398–2405, June 2005.

[13] Susanne Bock, Caroline Kuhhnel, Tobias Boskamp, and Heinz-Otto Peitgen. Robust vessel segmentation. In *Proc. of SPIE*, volume 6915, page 691539, San Diego, CA, USA, April 2008.

[14] K. L. Boedeker, V. N. Cooper, and M.F. McNitt-Gray. Application of the noise power spectrum in modern diagnostic MDCT: part i. measurement of noise power spectra and noise equivalent quanta. *Physics in Medicine and Biology*, 52:4027—4046, 2007.

[15] Horst Bunke. Error-Tolerant graph matching: A formal framework and algorithms. In *Proc. of the Joint IAPR International Workshops on Advances in Pattern Recognition*, pages 1–14. Springer-Verlag, 1998.

[16] Piergiorgio Cao, Paola De Rango, Gianbattista Parlani, and Fabio Verzini. Durability of abdominal aortic endograft with the talent unidoc stent graft in common practice: Core lab reanalysis from the TAURIS multicenter study. *J. Cardiovasc. Surg.*, 49(4):859–865, April 2009.

[17] Juan Cebral, Fernando Mut, Sunil Appanaboyina, Rainald Lohner, Carlos Miranda, Esteban Escrivano, Pedro Lylyk, Christopher Putman, Xiaoping P. Hu, and Anne V. Clough. Image-based analysis of blood flow modification in stented aneurysms. In *Proc. of SPIE*, volume 7262, page 72621G, Lake Buena Vista, FL, USA, February 2009.

[18] Carl Chartrand-Lefebvre, Alexandre Cadrin-Chenevert, Edith Bordeleau, Patricia Ugolini, Robert Ouellet, Jean-Louis Sablayrolles, and Julie Prenovault. Coronary computed tomography angiography: Overview of technical aspects, current concepts, and perspectives. *Canadian Association of Radiology*, 58:92—108, 2007.

[19] Yongchoel Choi and Seungyong Lee. Injectivity conditions of 2d and 3d uniform cubic b-spline functions. *Graphical Models*, 62(6):411–427, 2000.

[20] Gary Christensen. Consistent Linear-Elastic transformations for image matching. *Information Processing in Medical Imaging*, 1613:224–237, 1999.

[21] Gary E Christensen, Hans J Johnson, and Michael W Vannier. Synthesizing average 3D anatomical shapes. *NeuroImage*, 32(1):146–158, August 2006.

[22] GE Christensen, RD Rabitt, and MI Miller. Deformable Templates Using Large Deformation Kinematics. In *IEEE Transactions on Medical Imaging*, volume 5(10), pages 1447, 1435, 1996.

[23] Laurent D Cohen. On active contour models and balloons. *Computer Vision, Graphics, and Image Processing: Image Understanding*, 53(2):211—218, 1991.

[24] Laurent D. Cohen and Thomas Deschamps. Segmentation of 3D tubular objects with adaptive front propagation and minimal tree extraction for 3D medical imaging. *Comput Methods Biomech Biomed Engin*, 10(4):289, 2007.

[25] T. F Cootes, S. Marsland, C. J Twining, K. Smith, and C. J Taylor. Groupwise diffeomorphic non-rigid registration for automatic model building. *Proceedings of ECCV*, 4:316—327, 2004.

[26] W.R. Crum, O. Camara, and D.L.G. Hill. Generalized overlap measures for evaluation and validation in medical image analysis. *Medical Imaging, IEEE Transactions on*, 25(11):1451–1461, 2006.

[27] Jorg L. De Bruin, Annette F. Baas, Jaap Buth, Monique Prinssen, Eric L.G. Verhoeven, Philippe W.M. Cuypers, Marc R.H.M. van Sambeek, Ron Balm, Diederick E. Grobbee, Jan D. Blankensteijn, and the DREAM Study Group. Long-Term outcome of open or endovascular repair of abdominal aortic aneurysm. *N. Engl. J. Med.*, 362(20):1881–1889, May 2010.

[28] Stefanie Demirci, Frode Manstad-Hulaas, Nassir Navab, Michael I. Miga, and Kenneth H. Wong. Quantification of abdominal aortic deformation after EVAR. In *Proc. of SPIE*, volume 7261, page 72611U, Lake Buena Vista, FL, USA, February 2009.

[29] Catherine Desbleds-Mansard, Alfred Anwander, Linda Chaabane, Maciej Orkisz, Bruno Neyran, Philippe C Douek, and Isabelle E Magnin. Dynamic active contour model for size independent blood vessel lumen segmentation and quantification in High-Resolution magnetic resonance images. In *Computer Analysis of Images and Patterns: 9th International Conference, CAIP 2001 Warsaw, Poland, September 5-7, 2001, Proceedings*, 2001.

[30] Thomas Deschamps and Laurent D Cohen. Fast extraction of minimal paths in 3D images and applications to virtual endoscopy. *Med Image Anal*, 5:281—299, 2000.

[31] Benoit Desjardins and Ella A. Kazerooni. ECG-Gated cardiac CT. *AJR Am J Roentgenol*, 182(4):993–1010, April 2004.

[32] Marc Dewey, Florian Teige, Wolfgang Rutsch, Tania Schink, and Bernd Hamm. CT coronary angiography: Influence of different cardiac reconstruction intervals on image quality and diagnostic accuracy. *European Journal of Radiology*, 67:92—99, 2008.

[33] E.W. Dijkstra. A note on two problems in connexion with graphs. *Numerische Mathematik*, 1:269—271, 1959.

[34] David Eberly. *Distance from a Point to an Ellipse in 2D*. 2004.

[35] J. Ehrhardt, R. Werner, A. Schmidt-Richberg, and H. Handels. Statistical modeling of 4D respiratory lung motion using diffeomorphic image registration. *IEEE Transactions on Medical Imaging*, 30(2):251–265, 2011.

[36] Kari Ravn Eide, Asbjorn Odegard, Hans Olav Myhre, and Olav Haraldseth. Initial observations of endovasculair aneurysm repair using Dyna-CT. *Journal of Endovascular Therapy*, 14(1):50—53, 2007.

[37] Dirk Ertel, Evelyn Krober, Yiannis Kyriakou, Oliver Langner, and Willi A. Kalender. Modulation transfer function-based assessment of temporal resolution: Validation for single- and Dual-Source CT. *Radiology*, 248(3):1013–1017, September 2008.

[38] Rachid Fahmi, Anna Jerebko, Matthias Wolf, and Aly A. Farag. Robust segmentation of tubular structures in medical images. In *Proc. of SPIE*, page 691443, San Diego, CA, USA, 2008.

[39] Reza Fazel, Harlan M. Krumholz, Yongfei Wang, Joseph S. Ross, Jersey Chen, Henry H. Ting, Nilay D. Shah, Khurram Nasir, Andrew J. Einstein, and Brahmajee K. Nallamothu. Exposure to Low-Dose ionizing radiation from medical imaging procedures. *New England Journal of Medicine*, 361(9):849–857, 2009.

[40] Andrew W Fitzgibbon, Maurizio Pilu, and Robert B Fisher. Direct least squares fitting of ellipses. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 21(5):476—480, 1999.

[41] T. Flohr, B. Ohnesorge, H. Bruder, K. Stierstorfer, J. Simon, C. Suess, and S. Schaller. Image reconstruction and performance evaluation for ECG-gated spiral scanning with a 16-slice CT system. *Medical Physics*, 30:2650—2662, 2003.

[42] L. Florez-Valencia, J. Montagnat, and M. Orkisz. 3D graphical models for vascular-stent pose simulation. *International Journal of Machine Graphics & Vision*, 13(3):235—248, 2004.

[43] Alejandro F Frangi, Wiro J Niessen, Koen L Vincken, and Max A Viergever. Multiscale vessel enhancement filtering. *Lect Notes Comput Sci*, 1496:130—137, 1998.

[44] M. Freiman, L. Joskowicz, J. Sosna, Michael I. Miga, and Kenneth H. Wong. A variational method for vessels segmentation: algorithm and application to liver vessels visualization. In *Proc. of SPIE*, volume 7261, page 72610H, Lake Buena Vista, FL, USA, February 2009.

[45] T. O. Fuchs, M. Kachelriess, and W. A. Kalender. System performance of multislice spiral computed tomography. *IEEE Eng Med Biol Mag*, 19:63—70, 2000.

[46] Marcel J. W Greuter, Joost Dorgelo, Wim G. J Tukker, and Matthijs Oudkerk. Study on motion artifacts in coronary arteries with an anthropomorphic moving heart phantom on an ECG-gated multidetector computed tomography unit. *European Radiology*, 5:995—1007, 2005.

[47] M Akif Gülsün and Hüseyin Tek. Robust vessel tree modeling. *Proceedings of MICCAI*, 11:602–611, 2008.

[48] Radim Halir and Jan Flusser. Numerical stable direct least squares fitting of ellipses. In *Internatianal Conference in Central Europe on Computer Graphics, Visualization and Interactive Digital Media*, 1998.

[49] Dilana Hazer, Ender A. Finol, Michael Kostrzewa, Maria Kopaigorenko, Gotz-M. Richter, Rudiger Dillmann, Xiaoping P. Hu, and Anne V. Clough. Computational biomechanics and experimental validation of vessel deformation based on 4D-CT imaging of the porcine aorta. In *Proc. of SPIE*, volume 7262, page 72621F, Lake Buena Vista, FL, USA, February 2009.

[50] Marcela Hernandez-Hoyos, Maciej Orkisz, Philippe Puech, Catherine Mansard-Desbleds, Philippe Douek, and Isabelle E Magnin. Computer-assisted analysis of three-dimensional MR angiograms. *Radiographics*, 22(2):421—436, 2002.

[51] J. A. van Herwaarden, L. W Bartels, B. E Muhs, K. L Vincken, M. Y.A Lindeboom, A. Teutelink, F. L Moll, and H. J Verhagen. Dynamic magnetic resonance angiography of the aneurysm neck: Conformational changes during the cardiac cycle with possible consequences for endograft sizing and future design. *J. Vasc. Surg.*, 44(1):22—28, 2006.

[52] J. A. van Herwaarden, B. E Muhs, K. L Vincken, J. van Prehn, A. Teutelink, L. W Bartels, F. L Moll, and H. J Verhagen. Aortic compliance following EVAR and the influence of different endografts: determination using dynamic MRA. *Journal of Endovascular Therapy*, 13(3):406—414, 2006.

[53] The United Kingdom EVAR Trial Investigators. Endovascular versus open repair of abdominal aortic aneurysm. *N. Engl. J. Med.*, 362(20):1863–1871, May 2010.

[54] Tikva S. Jacobs, Jamie Won, Edwin C. Gravereaux, Peter L. Faries, Nicholas Morrissey, Victoria J. Teodorescu, Larry H. Hollier, and Michael L. Marin. Mechanical failure of prosthetic human implants: A 10-year experience with aortic stent graft devices. *J. Vasc. Surg.*, 37(1):16–26, January 2003.

[55] Uwe Jandt, Dirk Schäfer, Michael Grass, and Volker Rasche. Automatic generation of 3D coronary artery centerlines using rotational x-ray angiography. *Med Image Anal*, 13(6):846–858, December 2009.

[56] S Joshi, B Davis, M Jomier, and G Gerig. Unbiased diffeomorphic atlas construction for computational anatomy. *NeuroImage*, 23:S151–S160, 2004.

[57] Jens N. Kaftan, Huseyin Tek, Til Aach, Josien P. W. Pluim, and Benoit M. Dawant. A two-stage approach for fully automatic segmentation of venous vascular structures in liver CT images. In *Proc. of SPIE*, volume 7259, page 725911, Lake Buena Vista, FL, USA, February 2009.

[58] Willi A. Kalender. *Computed Tomography.* Publicis Corporate Publishing, Erlangen, 2005.

[59] Cemil Kirbas and Francis Quek. A review of vessel extraction techniques and algorithms. *ACM Comput. Surv.*, 36(2):81–121, 2004.

[60] A. Klein, W. K Renema, L. J Schultze Kool, and C. H Slump. Initial steps towards automatic segmentation of the wire frame of stent grafts in CT data. In *Proc. of IEEE-EMBS Benelux Chapter*, pages 116—119, Enschede, The Netherlands, November 2009.

[61] Almar Klein, Dirk-Jan Kroon, Yvonne Hoogeveen, Leo J. Schultze Kool, W. KlaasJan Renema, and Cornelis H. Slump. Multimodal image registration by edge attraction and regularization using a b-spline grid. In *Proc. of SPIE*, pages 7962–71, Orlando, USA, 2011.

[62] Almar Klein, Luuk J. Oostveen, Marcel J. W. Greuter, Yvonne Hoogeveen, Leo J. Schultze Kool, Cornelis H. Slump, and W. Klaas Jan Renema. Detectability of motions in AAA with ECG-gated CTA: a quantitative study. *Medical Physics*, 36(10):4616–4624, October 2009.

[63] Almar Klein, Luuk J. Oostveen, Marcel J. W. Greuter, Yvonne Hoogeveen, Leo J. Schultze Kool, Cornelis H. Slump, and W. Klaas Jan Renema. Diagnostic quality of time-averaged ECG-gated CT data. In *Proceedings of SPIE*, page 725836, Lake Buena Vista, FL, USA, 2009.

[64] Almar Klein, W. Klaas Jan Renema, J. Adam van der Vliet, Luuk J. Oostveen, Yvonne Hoogeveen, Leo J. Schultze Kool, and Cornelis H. Slump. Motion calculations on stent grafts in AAA. In *Diagnosis, Screening and Treatment of Abdominal, Thoracoabdominal and Thoracic Aortic Aneurysms*, pages 126–144. InTech, 2011.

[65] Almar Klein, W. KlaasJan Renema, Luuk J. Oostveen, Leo J. Schultze Kool, and Cornelis H. Slump. A segmentation method for stentgrafts in the abdominal aorta from ECG-gated CTA data. In *Proc. of SPIE*, page 69160R, San Diego, CA, USA, 2008.

[66] Almar Klein, J. Adam van der Vliet, Luuk J. Oostveen, Yvonne Hoogeveen, Leo J. Schultze Kool, W. Klaas Jan Renema, and Cornelis.H. Slump. Automatic segmentation of the wire frame of stent grafts from CT data. *Medical Image Analysis*, In Press, Accepted Manuscript, 2011.

[67] S. Klein, M. Staring, K. Murphy, M.A. Viergever, and J. Pluim. elastix: A toolbox for Intensity-Based medical image registration. *Medical Imaging, IEEE Transactions on*, 29(1):196–205, 2010.

[68] C. Kleinstreuer, Zhonghua Li, and M. A Farber. Fluid-Structure interaction analyses of stented abdominal aortic aneurysms. *Annu Rev Biomed Eng*, 9:169—204, 2007.

[69] Jan J Koenderink. The structure of images. *Biological Cybernetics*, 50(5):363—370, 1984.

[70] O. H.J Koning, O. R Oudegeest, E. R Valstar, E. H Garling, E. van der Linden, J. W Hinnen, J. F Hamming, A. M Vossepoel, and J. H. van Bockel. Roentgen stereophotogrammetric analysis: An accurate tool to assess Stent-Graft migration. *Journal of Endovascular Therapy*, 13(4):468—475, 2006.

[71] Olivier H J Koning, Bart L Kaptein, Eric H Garling, Jan Willem Hinnen, Jaap F Hamming, Edward R Valstar, and J H van Bockel. Assessment of three-dimensional stent-graft dynamics by using fluoroscopic roentgenographic stereophotogrammetric analysis. *Journal of Vascular Surgery: Official Publication, the Society for Vascular Surgery [and] International Society for Cardiovascular Surgery, North American Chapter*, 46(4):773–779, October 2007.

[72] Dirk-Jan Kroon and Cornelis H. Slump. MRI modalitiy transformation in demon registration. In *Proceedings of the Sixth IEEE international conference on Symposium on Biomedical Imaging: From Nano to Macro*, pages 963–966, Boston, Massachusetts, USA, 2009.

[73] R K Kwan, A C Evans, and G B Pike. MRI simulation-based evaluation of image-processing and classification methods. *IEEE Transactions on Medical Imaging*, 18(11):1085–1097, November 1999.

[74] G. Langs, N. Paragios, R. Donner, P. Desgranges, A. Rahmouni, and H. Kobeiter. Motion analysis of endovascular Stent-Grafts by MDL based registration. In *Proc IEEE Int Conf Comput Vis*, pages 1–8, 2007.

[75] Georg Langs, Nikos Paragios, Pascal Desgranges, Alain Rahmouni, and Hicham Kobeiter. Learning deformation and structure simultaneously: In situ endograft deformation analysis. *Medical Image Analysis*, 15(1):12–21, February 2011.

[76] Jack Lee, Patricia Beighley, Erik Ritman, and Nicolas Smith. Automatic segmentation of 3D micro-CT coronary vascular images. *Med Image Anal*, 11:630—647, 2007.

[77] Seungyong Lee, George Wolberg, Kyung-yong Chwa, and Sung Yong Shin. Image metamorphosis with scattered feature constraints. *IEEE Transactions on Visualization and Computer Graphics*, 2:337—354, 1996.

[78] Seungyong Lee, George Wolberg, and Sung Yong Shin. Scattered data interpolation with multilevel b-splines. *IEEE Transactions on Visualization and Computer Graphics*, 3(3):228—244, 1997.

[79] David Lesage, Elsa D. Angelini, Isabelle Bloch, and Gareth Funka-Lea. A review of 3D vessel lumen segmentation techniques: Models, features and extraction schemes. *Med Image Anal*, 13(6):819–845, December 2009.

[80] Zhonghua Li and C. Kleinstreuer. Analysis of biomechanical factors affecting stent-graft migration in an abdominal aortic aneurysm model. *J Biomech*, 39:2264—2273, 2006.

[81] Zhonghua Li and C. Kleinstreuer. Computational analysis of type II endoleaks in a stented abdominal aortic aneurysm model. *J Biomech*, 39:2573–2582, 2006.

[82] J. Lu, J. Egger, A. Wimmer, S. Grosskopf, B. Freisleben, Michael I. Miga, and Kevin R. Cleary. Detection and visualization of endoleaks in CT data for monitoring of thoracic and abdominal aortic aneurysm stents. In *Proc. of SPIE*, volume 6918, page 69181F, San Diego, CA, USA, March 2008. SPIE.

[83] Antoine J. B Maintz and Max A Viergever. A survey of medical image registration. *Medical Image Analysis*, 2(1):1—37, 1998.

[84] Armando Manduca, Joel G. Fletcher, Robert J. Wentz, Raymond C. Shields, Terri J. Vrtiska, Hassan Siddiki, Theresa Nielson, Xiaoping P. Hu, and Anne V. Clough. Reproducibility of aortic pulsatility measurements from ECG-gated abdominal CTA in patients with abdominal aortic aneurysms. In *Proc. of SPIE*, volume 7262, page 72620L, Lake Buena Vista, FL, USA, February 2009. SPIE.

[85] R. Manzke, T. Kohler, T. Nielsen, D. Hawkes, and M. Grass. Automatic phase determination for retrospectively gated cardiac CT. *Medical Physics*, 31:3345—3362, 2004.

[86] R. Manzke, P. Koken, D. Hawkes, and M. Grass. Helical cardiac cone beam CT reconstruction with large area detectors: a simulation study. *Physics in Medicine and Biology*, 50:1547—1568, 2005.

[87] Stephen Marsland, Carole J. Twining, and Chris J. Taylor. A minimum description length objective function for groupwise non-rigid image registration. *Image and Vision Computing*, 26(3):333–346, March 2008.

[88] Julian Mattes, Iris Steingruber, Michael Netzer, Karl Fritscher, Helmut Kopf, Werner Jaschke, and Rainer Schubert. Spatio-temporal changes and migration of stent grafts after endovascular aortic aneurysm repair. *Int. Congr. Ser.*, 1281:393–397, May 2005.

[89] 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*, 15(2):238–249, April 2011.

[90] Bart E Muhs, Arno Teutelink, Matthias Prokop, Koen L Vincken, Frans L Moll, and Hence J. M Verhagen. Endovascular aneurysm repair alters renal artery movement: a preliminary evaluation using dynamic CTA. *Journal of Endovacular Therapy*, 13(4):476—480, 2006.

[91] Bart E. Muhs, K. L. Vincken, J. van Prehn, M. K. Stone, L. W. Bartels, M. Prokop, F. L. Moll, and H. J. Verhagen. Dynamic Cine-CT angiography for the evaluation of the thoracic aorta; insight in dynamic changes with implications for thoracic endograft treatment. *Eur J Vasc Endovasc Surg*, 32(5):532—536, 2006.

[92] K. Murphy, B. van Ginneken, S. Klein, M. Staring, B.J. de Hoop, M.A. Viergever, and J.P.W. Pluim. Semi-automatic construction of reference standards for evaluation of image registration. *Medical Image Analysis*, 15(1):71–84, February 2011.

[93] V. Noblet, C. Heinrich, F. Heitz, and J.-P. Armspach. Retrospective evaluation of a topology preserving non-rigid registration method. *Medical Image Analysis*, 10(3):366–384, June 2006.

[94] Bernd Ohnesorge, Thomas Flohr, Christoph Becker, Andreas F Kopp, Uwe J Schoepf, Ullrich Baum, Andreas Knez, Klaus Klingenbeck-Regn, and Maximilian F Reiser. Cardiac imaging by means of electrocardiographically gated multisection spiral CT: initial experience. *Radiology*, 217:564—571, 2000.

[95] Josien P. W. Pluim, J. B. Antoine Maintz, and Max A. Viergever. Image registration by maximization of combined mutual information and gradient information. In Scott L. Delp, Anthony M. DiGoia, and Branislav Jaramaz, editors, *Medical Image Computing and Computer-Assisted Intervention – MICCAI 2000*, volume 1935, pages 452–461. Springer Berlin Heidelberg, Berlin, Heidelberg, 2000.

[96] Andrew N Primak, Cynthia H McCollough, Michael R Bruesewitz, Fie Zhang, and Joel G Fletcher. Relationship between noise, dose and pitch in cardiac Multi-Detector row CT. *RadioGraphics*, 26(6):1785—1794, 2006.

[97] M. Prinssen, E. Buskens, and J. D. Blankensteijn. The dutch randomised endovascular aneurysm management (DREAM) trial. background, design and methods. *J. Cardiovasc Surg. (Torino)*, 43(3):379—84, 2002.

[98] Mathias Prokop. New challenges in MDCT. *European Journal of Radiology*, 15:35—45, 2005.

[99] F.K.H. Quek and C. Kirbas. Vessel extraction in medical images by wave-propagation and traceback. *IEEE Trans Med Imaging*, 20(2):117–131, 2001.

[100] Justus E Roos, Jeffrey C Hellinger, Richard Hallet, Dominik Fleischmann, Christopher K Zarins, and Geoffrey D Rubin. Detection of endograft fractures with multidetector row computed tomography. *J. Vasc. Surg.*, 42(5):1002–1006, November 2005.

[101] D. Rueckert, L. I Sonoda, C. Hayes, D. L.G Hill, M. O Leach, and D. J Hawkes. Nonrigid registration using Free-Form deformations: Application to breast MR images. *IEEE Transactions on Medical Imaging*, 18(8):712—721, 1999.

[102] Daniel Rueckert, Paul Aljabar, Rolf A Heckemann, Joseph V Hajnal, and Alexander Hammers. Diffeomorphic registration using b-splines. *Proceedings of MICCAI*, 9(Pt 2):702–709, 2006.

[103] Dieter Seghers, Emiliano D'Agostino, Frederik Maes, Dirk Vandermeulen, and Paul Suetens. Construction of a brain template from MR images using State-of-the-Art registration and segmentation techniques. In *Proceedings of MICCAI*, volume 3216, pages 696–703. Springer Berlin Heidelberg, Berlin, Heidelberg, 2004.

[104] Dag Sverre Seljebotn. Fast numerical computations with cython. Proceedings of the 8th Python in Science Conference, pages 15—22, Pasadena, CA USA, 2009. Proc. of the 8th Python in Science Conference.

[105] J. A Sethian. Fast marching methods. *SIAM Review*, 41:199—235, 1999.

[106] S. William Stavropoulos and Sridhar R Charagundla. Imaging techniques for detection and management of endoleaks after endovascular aortic aneurysm repair. *Radiology*, 243(3):641—655, 2007.

[107] Abigail Swillens, Lieve Lanoye, Julie De Backer, Nikos Stergiopulos, Pascal R Verdonck, Frank Vermassen, and Patrick Segers. Effect of an abdominal aortic aneurysm on wave reflection in the aorta. *IEEE Trans Biomed Eng*, 55(5):1602—1612, May 2008.

[108] Arno Teutelink, Bart Muhs, Koen L Vincken, Lambertus W Wartels, Sandra A Cornelissen, Joost A. van Herwaarden, Mathias Prokop, Frans L Moll, and Hence J. M Verhagen. Use of dynamic computed tomography to evaluate pre- and postoperative aortic changes in AAA patients undergoing eneurysm repair. *J. Endovasc. Ther.*, 14(1):44–49, 2007.

[109] J P Thirion. Image matching as a diffusion process: an analogy with maxwell's demons. *Medical Image Analysis*, 2(3):243–60, September 1998.

[110] Hans Henrik Thodberg. Minimum description length shape and appearance models. *Image Processing Medical Imaging*, pages 51—62, 2003.

[111] M Truijers, J A Pol, L J Schultzekool, S M van Sterkenburg, M F Fillinger, and J D Blankensteijn. Wall stress analysis in small asymptomatic, symptomatic and ruptured abdominal aortic aneurysms. *European Journal of Vascular and Endovascular Surgery: The Official Journal of the European Society for Vascular Surgery*, 33(4):401–407, April 2007.

[112] Dale Umbach and Kerry N Jones. A few methods for fitting circles to data. *IEEE Transactions on instrumentation and measurement.*, 52(6):881—1885, 2000.

[113] Michael Unser. Splines: A perfect fit for signal and image processing. *IEEE Signal Processing Magazine*, 16(6):22—38, 1999.

[114] Gilbert R Upchurch and Timothy A Schaub. Abdominal aortic aneurysm. *American Family Physician*, 73(7):1198–1204, April 2006.

[115] Maarten J van der Laan, Arno Teutelink, Rudy Meijer, Christopher L Wixon, and Jan D Blankensteijn. Noninvasive evaluation of the effectiveness of endovascular AAA exclusion. *Journal of Endovascular Therapy: An Official Journal of the International Society of Endovascular Specialists*, 10(3):458–462, June 2003.

[116] Tom Vercauteren, Xavier Pennec, Aymeric Perchant, and Nicholas Ayache. Non-parametric diffeomorphic image registration with the demons algorithm. In *Proceedings of MICCAI*, pages 319–326, Brisbane, Australia, 2007. Springer-Verlag.

[117] He Wang, Lei Dong, Jennifer O'Daniel, Radhe Mohan, Adam S Garden, K. Kian Ang, Deborah A Kuban, Mark Bonnen, Joe Y Chang, and Rex Cheung. Validation of an accelerated 'demons' algorithm for deformable image registration in radiation therapy. *Physics in Medicine and Biology*, 50:2887—2905, 2005.

[118] William M Wells, III, Paul Viola, and Ron Kikinis. Multi-Modal volume registration by maximization of mutual information. *Medical Image Analysis*, 1(1):35–51, 1996.

[119] Robert Wentz, Armando Manduca, J. G. Fletcher, Hassan Siddiki, Raymond C. Shields, Terri Vrtiska, Garrett Spencer, Andrew N. Primak, Jie Zhang, Theresa Nielson, Cynthia McCollough, and Lifeng Yu. Automatic segmentation and co-registration of gated CT angiography datasets: measuring abdominal aortic pulsatility. In *Proc. of SPIE*, volume 6511, page 65111I, San Diego, CA, USA, 2007.

[120] Eric E Williamson, Jacobo Kirsch, Philip A Araoz, Whitney B Edmister, Daniel D Borgeson, James F Glockner, and Jerome F Breen. ECG-Gated cardiac CT angiography using 64-MDCT for detection of patent foramen ovale. *American Journal of Roentgenology*, 190:929—933, 2008.

[121] O. Wink, W.J. Niessen, and M.A. Viergever. Multiscale vessel tracking. *IEEE Trans Med Imaging*, 23(1):130–133, 2004.

[122] Onno Wink, Alejandro F Frangi, Bert Verdonck, Max A Viergever, and Wiro J Niessen. 3D MRA coronary axis determination using a minimum cost path approach. *Magnetic Resonance in Medicine: Official Journal of the Society of Magnetic Resonance in Medicine / Society of Magnetic Resonance in Medicine*, 47(6):1169–1175, June 2002.

[123] Stefan Worz, William J. Godinez, Karl Rohr, Josien P. W. Pluim, and Benoit M. Dawant. Segmentation of 3D tubular structures based on 3D intensity models and particle filter tracking. In *Proc. of SPIE*, volume 7259, page 72591P, Lake Buena Vista, FL, USA, February 2009.

[124] Chenyang Xu and Jerry L Prince. Snakes, shapes and gradient vector flow. *IEEE Transactions on Image Processing*, 7(3):359—369, 1998.

[125] C K Zarins, R A White, D Schwarten, E Kinney, E B Diethrich, K J Hodgson, and T J Fogarty. AneuRx stent graft versus open surgical repair of abdominal aortic aneurysms: multicenter prospective clinical trial. *J. Vasc. Surg.*, 29(2):292–305, February 1999.

[126] Chuan Zhou, Heang-Ping Chan, Aamer Chughtai, Smita Patel, Prachi Agarwal, Lubomir M. Hadjiiski, Berkman Sahiner, Jun Wei, Jun Ge, Ella A. Kazerooni, Maryellen L. Giger, and Nico Karssemeijer. Automated segmentation and tracking of coronary arteries in ECG-gated cardiac CT scans. In *Proc. of SPIE*, volume 6915, page 69150O, San Diego, CA, USA, March 2008.

List of publications

Peer reviewed and book chapters

Almar Klein, Luuk J. Oostveen, Marcel J. W. Greuter, Yvonne Hoogeveen, Leo J. Schultze Kool, Cornelis H. Slump, and W. Klaas Jan Renema. *Detectability of motions in AAA with ECG-gated CTA: a quantitative study.* Medical Physics, 36(10): 4616–4624, October 2009.

Almar Klein, W. Klaas Jan Renema, J. Adam van der Vliet, Luuk J. Oostveen, Yvonne Hoogeveen, Leo J. Schultze Kool, and Cornelis. H. Slump. *Motion calculations on stent grafts in AAA.* in: Diagnosis, Screening and Treatment of Abdominal, Thoracoabdominal and Thoracic Aortic Aneurysms. ISBN 978-953-307-466-5, 2011.

Almar Klein, J. Adam van der Vliet, Luuk J. Oostveen, Yvonne Hoogeveen, Leo J. Schultze Kool, W. Klaas Jan Renema, and Cornelis. H. Slump. *Automatic segmentation of the wire frame of stent grafts from CT data.* Medical Image Analysis, 2011.

Peer reviewed in preparation

Almar Klein, W. Klaas Jan Renema, Luuk J. Oostveen, J. Adam van der Vliet, Yvonne Hoogeveen, Leo J. Schultze Kool, , and Cornelis. H. Slump. *Quantitative motion analysis of endovascular stent grafts by segmentation and image registration.* Medical Image Analysis, in preparation

Almar Klein, W. Klaas Jan Renema, J. Adam van der Vliet, Luuk J. Oostveen, Yvonne Hoogeveen, Cornelis. H. Slump, and Leo J. Schultze Kool. *Automated motion analysis of stent grafts in AAA.* Journal of Vascular Surgery, in preparation

Conference proceedings

Almar Klein, F. van der Heijden, and C. H Slump. *Alignment of diabetic feet images.* In the 2007 The third annual IEEE Benelux/DSP Valley Signal Processing Symposium, page 185–192, 2007.

Almar Klein, W. Klaas Jan Renema, Luuk J. Oostveen, Leo J. Schultze Kool, and Cornelis H. Slump. *A segmentation method for stentgrafts in the abdominal aorta from ECG-gated CTA data.* In Proc. of SPIE, page 69160R, San Diego, CA, 2008.

Almar Klein, Luuk J. Oostveen, Marcel J. W. Greuter, Yvonne Hoogeveen, Leo J.

Schultze Kool, Cornelis H. Slump, and W. Klaas Jan Renema. *Diagnostic quality of time-averaged ECG-gated CT data.* In Proceedings of SPIE, pages 725836-6, Lake Buena Vista, FL, USA, 2009.

A. Klein, W. K.J. Renema, L. J. Schultze Kool, and C. H. Slump. *Initial steps towards automatic segmentation of the wire frame of stent grafts in CT data.* In Proc. of IEEE-EMBS Benelux Chapter, pages 116–119, Enschede, The Netherlands, November 2009.

C. Kloeze, A. Klein, S. Hazenberg, F. van der Heijden, J. G van Baal, and S. A Bus. *Pixel classification for automated diabetic foot diagnosis.* In Proc. of IEEE-EMBS Benelux Chapter, Enschede, The Netherlands, November 2009.

Almar Klein, Dirk-Jan Kroon, Yvonne Hoogeveen, Leo J. Schultze Kool, W. KlaasJan Renema, and Cornelis˜H. Slump. *Multimodal image registration by edge attraction and regularization using a B-spline grid.* In Proc. of SPIE pages 7962–71, Orlando, USA, 2011.

Summary

Patients with an Abdominal Aortic Aneurysm (AAA) have a high risk of dying due to the rupture of a dilated aorta. Endovascular aneurysm repair (EVAR) is a technique to threat AAA, by which a stent graft prosthesis is implanted in the aorta of the patient, which takes the pressure off the aneurysm wall. Due to its minimal invasive character, this intervention has smaller risks for the patient compared to the conventional approach (in which the unhealthy aorta is replaced with an artificial vessel using open surgery).

Yielding good results on the short term, this technique has increased in popularity. Unfortunately, due to effects such as metal fatigue, leakage and stent migration, this technique is less successful on the long term. These problems are caused by the forces applied by the pressure waves of the blood flow. It is therefore important to understand the motion behavior of the stent graft inside the human body. This idea is the main motivation for the research described in this thesis.

The technique of *ECG-gated CT* can be used to obtain multiple 3D images of the patient. Each of these images corresponds to a different phase of the heart cycle. Therefore this technique enables measuring motions inside the patient. We have performed experiments to study the possibilities and limitations for using this technique to measure the motions of interest. We concluded that this technique is indeed suitable to study the motions of stent grafts in AAA.

Having obtained the set of 3D images, the goal is to measure the motions of the stent graft. This is done in two steps. The first is *segmentation*: detecting where the stent is located and subsequently creating a geometric model of the stent. This geometric model is represented as a graph consisting of nodes that are connected by edges. The nodes represent the corners and crossings in the stent's frame, and the edges represent the wires in between. The second step is *registration*, by which the deformation between the different images is calculated.

Our first attempted segmentation approach was to sample 2D slices (orthogonal to the centerline of the stent) from the 3D images. The points where the metal frame penetrates the slice are then detected. By repeating this process while tracking along the centerline of the stent, a model of the stent can be obtained. Unfortunately, we encountered many problems with the tracking part of this method related to bifurcations and overlapping parts of the stent. To overcome these problems, we decided to work on an approach to segment the stent directly in 3D. First, a set of seed points is detected in the 3D image, based on a few simple criteria. These points are connected to each other using a modified version of the minimum cost path (MCP) method, which is an algorithm that tries to find the optimal route between two or more points. This results in a graph with many edges, which is processed to obtain the final geometric model of the stent.

Next is the registration step, in which the deformation fields between the different

3D images are calculated. For this purpose we developed a groupwise registration method which ensures that the transformation fields do not fold. The motion of the stent is determined from the deformation fields and incorporated in the geometric stent model. The obtained dynamic model can be used for visualization of the motion of the stent, as well as for performing further calculations such as estimating the forces in the stent.

Clearly, the proposed method is just a small step on the road to better care for patients with AAA. The method should be seen as a tool that enables further quantitative research to motions of stent grafts. These studies will provide new insights in the behavior of the stent graft in vivo. We expect that this will enable the detection and prediction of stent failure in individual patients, and can help in designing better stent grafts in the future.

Samenvatting

Patiënten met een abdominaal aneurysma in de aorta (AAA) hebben een grote kans te overlijden aan het scheuren van een uitgerekt bloedvat. Met een techniek die EVAR genoemd wordt, kan AAA behandeld worden. Hierbij wordt een stent graft prothese geïmplanteerd in de aorta van de patiënt. Deze stent neemt de druk van de aortawand weg. Dankzij zijn minimaal invasieve karakter brengt deze operatietechniek een kleiner risico met zich mee dan de conventionele techniek (waarbij de beschadigde aorta wordt vervangen door een kunstvat door middel van open chirurgie).

Door zijn goede resultaten op de korte termijn heeft EVAR de laatste tijd flink aan populariteit gewonnen. Echter, door verschijnselen als metaalmoeheid, lekken en migratie van de stent, is de techniek minder succesvol op de lange termijn. Deze problemen vinden hun oorzaak in de krachten die ontstaan door de drukgolven van de bloedstroom. Om deze reden is het van belang om het bewegingsgedrag van de stent graft in de patiënt beter te begrijpen. Dit is de achterliggende gedachte voor het onderzoek beschreven in dit proefschrift.

De techniek *ECG-gated CT* kan worden gebruikt om meerdere 3D beelden van de patiënt te verkrijgen. Elk van deze beelden correspondeert met een verschillende fase van de hartcyclus. Zo maakt deze techniek het mogelijk bewegingen in de patiënt te meten. Uitgaande van de bewegingen waarin wij geïnteresseerd zijn, hebben wij experimenteel onderzoek verricht om de mogelijkheden en beperkingen van deze techniek in kaart te brengen. We konden concluderen dat deze techniek inderdaad geschikt is om de bewegingen van stent grafts in AAA te onderzoeken.

Het doel is om uit deze set van 3D beelden de bewegingen van de stent graft te meten. Dit doen we in twee stappen. De eerste stap is *segmentatie*. Dit is het detecteren van de stent om vervolgens een geometrisch model van de stent te genereren. Dit geometrische model wordt gerepresenteerd met een graaf bestaande uit nodes en verbindingen. De nodes representeren de bochten en kruisingen in het frame van de stent, en de verbindingen representeren de draden daartussen. De tweede stap is *registratie*, waarbij de vervorming tussen de verschillende beelden wordt berekend.

Onze aanvankelijke segmentatie methode was erop gericht 2D plakken (loodrecht op de middellijn van de stent) uit de 3D beelden te halen. De punten waar het metalen frame de plak doorkruist worden vervolgens gedetecteerd. Door dit te herhalen terwijl de middellijn van de stent gevolgd wordt, kan een model van de stent verkregen worden. Echter, wij kwamen bij deze methode veel problemen tegen met betrekking tot splitsingen en overlappende delen in de stent. Om deze problemen het hoofd te bieden is besloten om te werken aan een methode die direct in 3D werkt. Eerst wordt er in het 3D beeld een set van "beginpunten" gevonden op basis van enkele simpele criteria. Deze punten worden vervolgens verbonden door middel van een aangepaste versie van de MCP (minimale kosten pad) methode. Deze methode vindt de optimale route tussen twee of meerdere punten. Het resultaat is een graaf met veel

verbindingen, welke wordt opgeschoond om uiteindelijk het geometrische model van de stent te verkrijgen.

In de registratie stap worden de vervormingsvelden tussen de verschillende 3D beelden berekend. Hiervoor hebben wij een registratie algoritme ontwikkeld om meerdere beelden tegelijk te kunnen registreren, op een manier die "vouwen" in de vervormingsvelden vermijdt. Uit de vervormingsvelden wordt de beweging van de stent bepaald, die we vervolgens in het stent model incorporeren. Het verkregen dynamische model kan worden gebruikt om de beweging van de stent te visualiseren, alsmede om verdere berekenen te doen zoals het schatten van de krachten in de stent.

Het is duidelijk dat de voorgestelde methode nog maar een kleine stap is op weg naar een betere behandelmethode voor de patiënten. De methode moet worden gezien als een hulpmiddel dat het mogelijk maakt om kwantitatief onderzoek te doen naar de bewegingen van stent grafts. Dergelijk onderzoek zal leiden tot nieuwe inzichten in het gedrag van stent grafts. Wij verwachten dat dit vervolgens de detectie en voorspelling van stent falen in individuele patiënten mogelijk zal maken, en in de toekomst kan helpen bij het ontwerpen van betere stent grafts.

Dankwoord

Tijdens mijn onderzoek en het schrijven van dit proefschift heb ik ontzettend veel steun en hulp gehad van vele mensen om mij heen.

Als eerste wil ik mijn twee promotoren, Kees Slump en Leo Schultze Kool, alsmede mijn copromotor KlaasJan Renema hartelijk bedanken voor het mogelijk maken van deze leuke opdracht. Ik ben erg dankbaar voor de bijzondere combinatie waarin ik van jullie drie begeleiding heb ontvangen. Kees gaf mij de meeste technische ondersteuning, maar ook veel vrijheid en vertrouwen om mijn werk zelf in te vullen. Hierdoor heb ik onder andere de kans gehad te kijken naar alternatieve ontwikkelomgevingen, wat een leerzame ervaring is geweest die mij ook veel gebracht heeft. Leo gaf mij de "high-level" begeleiding die ervoor zorgde dat ik het uiteindelijke klinische doel niet uit het oog verloor. KlaasJan was degene die mij leerde in hoofdstukken van mijn proefschrift te denken en die mij heeft aangezet tot het publiseren van papers zodra een "hoofdstuk" af was. Deze combinatie van begeleiding heb ik als erg plezierig en nuttig ervaren. Alle drie veel dank hiervoor.

Zonder anderen te kort te doen wil graag een extra woord van dank richten aan KlaasJan: ik heb de samenwerking met jou altijd als zeer prettig ervaren. Je weet altijd heel goed wat belangrijk is en weet dat duidelijk te maken zonder het op te dringen. Ik heb ook ontzettend veel gehad aan onze discussies en het grondige commentaar dat ik ontving op mijn geschreven teksten. Ik heb erg veel van jou geleerd. Dank.

Verder wil ik een aantal mensen van het Radboud bedanken voor hun hulp met het reviseren van teksten, het annoteren van stents, het doen van experimenten, en voor alle gezelligheid: Yvonne Hoogeveen, Daan van der Vliet, Luuk Oostveen, en alle andere mensen van de afdeling klinische fysica.

Graag wil ik ook al mijn collega's van de vakgroep Signalen en Systemen bedanken voor de fijne werksfeer. Geert-Jan en Henny voor de hulp met het bouwen van de fantomen en alle andere technische ondersteuning. Dirk-Jan voor de vele nuttige discussies over registratie algoritmes; deze hebben mij ontzettend veel geholpen om mijn vingers achter de materie te krijgen. Bas voor de interessante discussies tijdens onze tijd in de Hoogekamp. Luuk voor het suggereren van de MCP methode; dat was nou net wat ik nodig had. Anne voor het opstoken van de discussies over Python versus Matlab©. Gerbert, René, Dirk-Jan, Rita en Bianca voor de gezellige momenten tijdens de conferenties in de VS.

Ook dank ik de studenten die ik heb mogen begeleiden tijdens mijn promotie: Carla, Maarten, Sèvrin, Paul en Michel.

Ook een woord van dank aan Rob Reilink, Ludo Visser, Robert Schroll en alle anderen uit de Python open source gemeenschap. De projecten waar we samen aan werken brengen mij veel plezier en maken mij effectiever in mijn dagelijkse werk.

Natuurlijk dank ik ook mijn paranimfen voor hun steun: Jorik en Dirk-Jan.

Pa en ma, bedankt voor jullie jarenlange steun en interesse, en voor het feit dat jullie altijd voor mij en mijn gezin klaarstaan.

Tot slot mijn dank aan Claudia, voor je liefde, steun en zorgzaamheid. En omdat je zo'n goede moeder bent voor Noah.

Almar
Enschede, november 2011