

This electronic thesis or dissertation has been downloaded from the King's Research Portal at <https://kclpure.kcl.ac.uk/portal/>



Non-Rigid 2D-3D Registration for Use in Computer-Assisted Abdominal Aortic Aneurysm Repair Procedures

Guyot, Alexis

Awarding institution:
King's College London

The copyright of this thesis rests with the author and no quotation from it or information derived from it may be published without proper acknowledgement.

END USER LICENCE AGREEMENT



Unless another licence is stated on the immediately following page this work is licensed

under a Creative Commons Attribution-NonCommercial-NoDerivatives 4.0 International

licence. <https://creativecommons.org/licenses/by-nc-nd/4.0/>

You are free to copy, distribute and transmit the work

Under the following conditions:

- Attribution: You must attribute the work in the manner specified by the author (but not in any way that suggests that they endorse you or your use of the work).
- Non Commercial: You may not use this work for commercial purposes.
- No Derivative Works - You may not alter, transform, or build upon this work.

Any of these conditions can be waived if you receive permission from the author. Your fair dealings and other rights are in no way affected by the above.

Take down policy

If you believe that this document breaches copyright please contact librarypure@kcl.ac.uk providing details, and we will remove access to the work immediately and investigate your claim.

Division of Imaging Sciences and
Biomedical Engineering
School of Medicine



Non-Rigid 2D-3D Registration for Use in Computer-Assisted Abdominal Aortic Aneurysm Repair Procedures

by Alexis GUYOT

A thesis submitted in partial fulfilment of the requirements for the
degree of Doctor of Philosophy of King's College London

September 2015

Supervisors:

Dr. Graeme PENNEY

Mr. Tom CARRELL

Examiners:

Prof. Daniel RUECKERT

Dr. Matt CLARKSON

Abstract

This thesis contributes to knowledge by describing three methods to non-rigidly register 2D X-ray images acquired during a Complex Endovascular Aneurysm Repair (CEVAR) procedure to a 3D pre-operative CT scan.

The first part of the thesis presents an interpolation framework (thin-plate spline) that is tailored to accurately register 3D CT scan data to 2D X-ray projection data. Registering the 3D to the 2D images proves challenging, due to the lack of information perpendicular to the imaging plane. A method to interpolate manually selected displacements of 3D points located on the aorta surface has been tailored to model the known error distribution along the X-ray projection directions.

The second part of the thesis describes the intra-operative use of finite element based algorithms to deform the aorta surface based upon the positions of a guide-wire during a CEVAR procedure. The aorta is sequentially deformed so as to encompass a simulated wire which was initialised at the centrelines of the aorta and which is dragged towards the guide-wire. Experiments were conducted on the mechanical parameters of the finite-element model and showed the influence of the Young Modulus and the Poisson's ratio on registration accuracy.

The third part of the thesis focuses on the novel use of interventional digital tomosynthesis images to extract intra-operative information on the calcifications of the aorta and drive non-rigid registration of the aorta during CEVAR. Calcium-based correspondences were established between the pre-operative and the intra-operative scene. A similarity measure has been defined as a weighted sum of a bending energy term and a second term that estimates how well pre-operative and intra-operative patches of calcium match. Erroneous correspondences are corrected using a simulated annealing optimisation on this similarity measure.

Using the three methods, large rigid registration errors of 9 *mm* were brought down to 4 *mm* or below the clinical target of 3 *mm* (half the diameter of the renal arteries). The proposed methods fit well with current clinical workflows. The first method presented above requires little manual input during the operation and the two other methods are/can be fully automated. The work presented in this thesis has the potential to increase the availability of image guidance systems for CEVAR procedures and for minimally invasive surgery where soft tissues are involved.

Acknowledgements

Completing this PhD thesis would not have been possible without the inspirational guidance of my supervisors Dr. Graeme Penney and Mr. Tom Carrell. It is hard to overstate the role Graeme played during the last four years: he constantly supported me and his enthusiasm was a great source of motivation in the toughest times of my research. Tom's clinical input was very much valued and this dissertation would not be the same without it.

I gratefully acknowledge Philips Healthcare for funding this research. Special mention goes to my clinical science contacts at Philips, Dr. Maya Barley, Ms. Kirsten Zuurmond and Dr. Thijs Elenbaas for their valuable insight and for welcoming me to their headquarters in the Netherlands.

Last, but certainly not least, I wish to express my heartfelt gratitude to my family. My mother, father and brothers have always been there for me when I needed them. For this, I cannot be thankful enough.

Contents

1	Introduction	12
1.1	Motivation	12
1.2	Summary of the Thesis	14
1.3	Original Contributions	15
2	Clinical Background	17
2.1	Introduction	18
2.2	Minimally Invasive Surgery	18
2.2.1	History	19
2.2.2	Minimally Invasive Surgery vs. Open Surgery	21
2.2.3	Minimally Invasive Surgery Environment	21
2.2.4	Endovascular Aneurysm Repair	22
2.2.5	Deformations of the Aorta during Surgery	26
2.3	Image-Guidance for Endovascular Aneurysm Repair	28
2.3.1	Information Required for Endovascular Aneurysm Repair	28
2.3.2	Image Requirements	29
2.3.3	Interventional Images	30
2.3.4	Use of Images during Endovascular Aneurysm Repair	35
2.4	Computer Assisted Aortic Surgery	35
2.4.1	Computer-aided Endovascular Aneurysm Repair	36
2.4.2	Limitations of Current Computer Assistance	38
2.4.3	Future of Computer-Aided Endovascular Repair	39
2.5	Conclusion	41
3	Literature Review	42
3.1	Introduction	43

3.2	Background on Registration	43
3.2.1	Registration	43
3.2.2	Medical Image Registration	44
3.2.3	Registration for Image Guided Interventions	46
3.3	Non-Rigid Registration for Medical Images	47
3.3.1	Main Applications of Non-Rigid Medical Image Registration	48
3.3.2	Challenges of Non-Rigid Medical Image Registration	48
3.3.3	Transformation Models	51
3.4	2D-3D Non-Rigid Registration	61
3.4.1	Statistical models	61
3.4.2	2D-3D non-rigid registration using algebraic reconstruction	62
3.4.3	Monocular model to image registration	64
3.4.4	Graph-based deformable 2D-3D registration	66
3.4.5	Thin-plate splines based 2D-3D registration	69
3.4.6	Summary	70
3.5	Conclusion	73
4	Spline-based 2D/3D Registration	75
4.1	Introduction	76
4.2	Theory	76
4.2.1	Interpolating Thin-Plate Splines	77
4.2.2	Approximating Thin-Plate Splines	79
4.2.3	Approximating Thin-Plate Splines for Medical Image Registration	81
4.2.4	Approximating Thin-Plate Splines for Computer Assisted Aneurysm Repair	83
4.3	Experiments	88
4.3.1	Data	88
4.3.2	Methods	90
4.3.3	Validation	92
4.4	Results	93
4.5	Discussion	96
4.6	Conclusion	100

5	Non-Rigid 2D/3D Registration using Finite Element Methods	101
5.1	Overview	102
5.2	Method	104
5.2.1	Registration Pipeline	104
5.2.2	Finite-Element Methods for Aorta Registration	110
5.3	Experiment 1: Using Physically Based Mechanical Parameters	114
5.3.1	Data	115
5.3.2	Method	116
5.3.3	Validation	117
5.3.4	Results	121
5.3.5	Discussion	122
5.4	Experiment 2: synthetic data	123
5.4.1	Data	123
5.4.2	Method	124
5.4.3	Results	124
5.4.4	Discussion	126
5.5	Experiment 3: Optimal Parameter Values	127
5.5.1	Data	127
5.5.2	Method	127
5.5.3	Validation	128
5.5.4	Results	128
5.5.5	Discussion	131
5.6	Overall Discussion	133
5.7	Conclusion	137
6	Digital Tomosynthesis	139
6.1	Introduction	140
6.2	Theory	140
6.2.1	Digital Tomosynthesis	141
6.2.2	Enhancement of 2D Projection Images during EVAR with Digital Tomosynthesis	142
6.2.3	Retrieving and Using 3D Information during EVAR with Dig- ital Tomosynthesis	148

6.3	Interventional Tomosynthesis Registration Method	152
6.3.1	Automatic Registration of Manually Selected Patches	153
6.3.2	Automatic Registration of Automatically Selected Patches	156
6.3.3	Automatic Correction of Outliers in Registered Automatically Selected Patches	158
6.4	Experiments	163
6.4.1	Data	163
6.4.2	Method	164
6.4.3	Validation	166
6.5	Results	167
6.5.1	3D ground truth	167
6.5.2	2D ground truth	167
6.6	Discussion	169
6.7	Conclusion	174
7	Conclusion	176
7.1	Summary	176
7.2	Computer Assistance and Clinical Constraints	180
7.3	Future Work	181
7.4	Overall Summary	182

List of Figures

2.1	Abdominal Aortic Aneurysm	22
2.2	Endovascular aneurysm repair	24
2.3	Abdominal aortic aneurysm cases	25
2.4	Stent-graft types	26
2.5	Causes of renal ostia motion	27
2.6	Intra-operative X-ray images	33
2.7	Computer assistance workflow	37
2.8	Computer assistance example	38
2.9	Limitations of computer assistance systems	39
2.10	Future of computer assistance	40
3.1	Back-projection uncertainties	47
3.2	rigid/non-rigid registration	49
3.3	Motion models	53
3.4	Statistical shape models	54
3.5	Aorta deformations and anatomical constraints	66
4.1	Thin-plate spline interpolation	77
4.2	Thin-plate spline registration	77
4.3	Thin-plate spline: analytic solution	79
4.4	Interpolating/Approximating thin-plate spline point picking	79
4.5	Landmark error ellipses	80
4.6	Same dimensionality thin-plate spline registration	82
4.7	Back-projection line landmark error	85
4.8	Thin-plate splines + Projection Uncertainties	86
4.9	Sequential thin-plate spline warp	86

4.10	1D-2D simulated data	88
4.11	2D-3D simulated data	89
4.12	Interventional data	91
4.13	Validation	92
4.14	1D-2D simulated data: registration results	93
4.15	2D-3D simulated data: registration results	94
4.16	Interventional data: registration results	96
4.17	Information from a second view	99
5.1	Abdominal aortic aneurysm repair simulation	103
5.2	Registration pipeline	104
5.3	Wire initialisation	106
5.4	Intra-operative wire computation	107
5.5	Deformation field computation	109
5.6	Collision detection algorithm	110
5.7	Aorta surface/mesh deformation	115
5.8	Aorta registration: two stages	118
5.9	Point-to-surface registration error	119
5.10	Validation: manual selection of edge points	120
5.11	Mechanical parameters from the literature	121
5.12	Synthetic data experiment	124
5.13	Synthetic data experiment results	125
5.14	Mechanical parameters optimisation on one patient	129
5.15	Registration results for optimised mechanical parameters	131
5.16	Anatomical constraints	136
6.1	Shift-and-add tomosynthesis	142
6.2	Interventional digital tomosynthesis	143
6.3	Digital tomosynthesis with de-boning	144
6.4	Enhancement of 2D projections using digital tomosynthesis	146
6.5	De-boned digital tomosynthesis volume	148
6.6	Comparison of digital tomosynthesis and X-ray imaging	149
6.7	Digital tomosynthesis planes	150
6.8	Digital tomosynthesis registration pipeline	151

6.9	Correspondences between real and synthetic digital tomosynthesis volumes	153
6.10	Alignment of patches of calcium	154
6.11	Displacement field generation	156
6.12	Automated selection of calcium patches: method	157
6.13	Automated selection of calcium patches: results	158
6.14	Correction of displacement field	159
6.15	Clinical data	165
6.16	3D validation	168
6.17	2D validation	170

List of Tables

3.1	method summary: type, data, accuracy and validation method . . .	71
3.2	method summary: number of views, application and computation time	72
4.1	Simulated 1D/2D data target registration errors for 8 sets of deformations	94
4.2	Simulated 2D/3D data target registration errors for 8 sets of deformations	95
4.3	Reprojection errors for interventional data	95
5.1	Left renal ostium reprojection errors (mechanical parameters from the literature)	121
5.2	Right renal ostium reprojection errors (mechanical parameters from the literature)	122
5.3	Mean and std point-to-surface distance errors (mechanical parameters from the literature)	122
5.4	mean points-to-surface distance errors for different (Young’s modulus, Poisson’s ratio)	130
5.5	mean point-to-surface distance errors using optimised mechanical parameters	132
6.1	3D distances between ground truth and different registration methods (expert 1)	167
6.2	3D distances between ground truth and different registration methods (expert 2)	167
6.3	Reprojection errors for different registration methods (expert 1) . .	169
6.4	Reprojection errors for different registration methods (expert 2) . .	169

List of Abbreviations

2D	Two dimensional
3D	Three dimensional
AAA	Abdominal Aortic Aneurysm
AP	Anterior-Posterior
CT	(X-ray) Computed Tomography
CBCT	(X-ray) Cone-Beam Computed Tomography
CPU	Central Processing Unit
DRR	Digitally Reconstructed Radiograph
DSA	Digital Subtraction Angiography
DTS	Digital Tomosynthesis
FLE	Fiducial Localisation Error
FRE	Fiducial Registration Error
CEVAR	Complex Endovascular Aneurysm Repair
EVAR	Endovascular Aneurysm Repair
GPU	Graphics Processing Units
MR	Magnetic Resonance
PCA	Principal Component Analysis
PET	Positron Emission Tomography
PU	Projection Uncertainty
RMS	Root Mean Square
ROI	Region Of Interest
TPS	Thin-Plate Spline
TRE	Target Registration Error

Chapter 1

Introduction

Contents

1.1 Motivation	12
1.2 Summary of the Thesis	14
1.3 Original Contributions	15

The work presented in this thesis is concerned with calculating from 2D X-ray images the deformations that a 3D aorta undergoes during an endovascular aneurysm repair procedure and is presented in partial fulfilment of the requirements for the degree of Doctor of Philosophy of the University of London. Firstly, this chapter gives an insight into the reasons why this research was conducted, followed by a brief summary of the overall thesis. Lastly, an account is provided of how this document contributes to knowledge.

1.1 Motivation

Surgery is used to treat injuries or disorders of the body by cutting it open and manipulating the diseased anatomy of a patient. Where large incisions are concerned, the patient's discomfort and their recovery time are increased. In addition, there is a greater risk for the wounds to get infected, and particular care has to be given to maintain full sterility of the operating room. Minimally-invasive surgery uses smaller cuts to get access to the diseased anatomy and addresses all the aforementioned issues. Carrying-out minimally-invasive surgery, however, comes at a

cost for the surgeons, who lose direct access (vision and tactile feedback) to the structures they are operating on and have to rely on other information to guide the instruments they are using. The collection of data inside the operating theatre is limited in several ways: manoeuvring the imaging apparatus should not disrupt the clinical setting (in particular, the devices in use must not obstruct the surgical site from which the surgeons introduce their instruments) and the imaging devices should allow for quick acquisition if required. For a wide range of minimally invasive surgery, especially the main clinical application for the work described in this thesis, Endovascular Aneurysm Repair (EVAR), images are acquired prior to the intervention for diagnostic purposes and to plan the operation. The acquisition of these images is not subject to the same constraints as those inherent to interventional imaging and, as a result, valuable information can be extracted from pre-operative images which cannot easily be obtained during the intervention. During surgery, the clinical practitioners can make use of that information via a process of mental mapping or by using more sophisticated solutions, such as tools that automatically register the pre-operative and intra-operative data.

EVAR procedures involve the deployment of a prosthesis at the site of an aortic aneurysm to prevent its expansion and rupture. During those interventions, instruments can cause large deformations of the vessels they are inserted into, particularly so in the case of initially tortuous aortas. The current generation of computer-assistance systems which rigidly match the 3D pre-operative data from a CT-scan to the 2D X-ray images acquired during EVAR interventions do not account for instrument-induced motion. The research detailed in this thesis intends to devise strategies to recover the deformation that occurs during surgery. What motivates this work is the need to find registration methods that:

- account for the difference of dimensionality between the (3D) pre-operative and (2D) intra-operative data
- deform the soft-tissues from pre-operative data in accordance with the information available during surgery
- align images where different types of information are visible: from the pre-operative CT scan, an operator can see the soft-tissues, calcifications and

bones of a patient, while poor soft-tissue differentiation in X-ray images means that intra-operative images often do not contain useful soft-tissue information, unless nephrotoxic contrast agents are used.

1.2 Summary of the Thesis

This thesis is divided into seven chapters, which are all summarised below.

Chapter 2 details the clinical context of the research carried out in this thesis. A general review of the field of minimally invasive surgery is first given, followed by a description of endovascular aneurysm repair procedures. A discussion ensues on the nature of the images that are used before and during these procedures. Last, current and future computer-assistance systems for use in EVAR surgery are presented to the reader.

Chapter 3 gives a literature review of registration algorithms, with a focus on 2D-3D non-rigid registration for computer-assistance. A basic overview of registration methods is given in the first part of the chapter. Those are narrowed down successively to approaches used in image-guided interventions and to strategies devised to non-rigidly align anatomical structures. The last section describes previous works that specifically deal with 2D-3D non-rigid registration.

Chapter 4 proposes the novel use of thin-plate splines for 2D-3D non-rigid registration purposes. To date, thin-plate splines have been proposed to register data that have the same dimensionality by establishing (absolute or approximate) point correspondences between them. The method proposed in this chapter uses thin-plate spline to loosely match 3D points picked on CT data acquired before an EVAR procedure to 2D points manually selected on X-ray data collected during the intervention. The first part of the chapter explains how this approximating thin-plate spline framework was tailored to work with a combination of 2D+3D data, while the second part describes experiments carried out on synthetic and real data to assess the validity of this approach. Results, discussion and conclusions are presented in subsequent paragraphs.

Chapter 5 describes finite-element methods to deform an aorta surface based upon the position of instruments during an EVAR procedure. The approach that has been proposed consists of initialising an idealised wire at the centreline of the aorta surface extracted from the pre-operative scan, computing the positions of a stiff guide-wire inserted during the operation, moving the idealised wire towards the intra-operative wire and deforming the aorta alongside so that it eventually encompasses the intra-operative wire. This method was described in a first section, with additional background on the use of finite-element methods to deform the aorta. Experiments were then conducted on real and synthetic data which are presented in the three following sections. The last part of the chapter discusses the results that were obtained.

Chapter 6 describes a new approach to use digital tomosynthesis reconstruction methods to drive a non-rigid registration of the aorta. The first part discusses the ways to obtain 3D information on the calcifications of the aorta, as well as how to use this information to register pre- and intra- operative information for an EVAR procedure. The next section describes the methods that were devised to extract information on calcium deposits from digital tomosynthesis volumes acquired before and during the intervention and how this information is input into a novel 3D-3D non-rigid registration algorithm. Experiments on clinical data are subsequently presented, along with the results of these experiments. A discussion on digital tomosynthesis for EVAR concludes this chapter.

Chapter 7 summarises the work that was carried out for this thesis and discusses future directions for non-rigid registration of the aorta in EVAR.

1.3 Original Contributions

The work presented in this thesis contributes to knowledge in three ways:

1. A mathematical framework that incorporates knowledge of the known highly anisotropic error distribution when carrying out 2D/3D registration was developed and tested.

Impact: The framework enables 2D/3D non-rigid registration using manually picked landmarks on a 3D pre-operative surface and 2D intra-operative X-ray images. Only a few points (e.g. less than 10 pairs of source/target points) are needed, which can be selected within a reasonable time-frame. Interactive update and display of the (deformed) aorta surface every time a new pair of points are selected (on the 3D surface and 2D image respectively) allows the surgeons to visually assess the validity of each step of the manual registration process.

2. A method was developed to use the positions of instruments during image guided surgery to drive finite element methods for use in EVAR.

Impact: If coupled with automatic extraction of instruments this approach could provide fully automated non-rigid 2D-3D registration. The only additional disruption to clinical workflow would be the acquisition of a second rotated fluoroscopy view which is only a minor disruption of the clinical setting. Using such a method that only relies on the 3D positions of instruments bypasses the need to inject nephrotoxic contrast-agents into a patient to visualise their blood vessels.

3. A new approach was proposed to automatically enhance specific anatomy in interventional digital tomosynthesis by making use of pre-operative data and a novel non-rigid 3D-3D registration algorithm was devised to use this data.

Impact: Little information is visible on the calcifications of a patient in X-ray images due to limited contrast. An interventional digital tomosynthesis method has been devised to specifically highlight calcium deposits on the aorta wall. The visible 3D intra-operative information can be used to drive a 3D-3D non-rigid registration algorithm.

Chapter 2

Clinical Background

Contents

2.1	Introduction	18
2.2	Minimally Invasive Surgery	18
2.2.1	History	19
2.2.2	Minimally Invasive Surgery vs. Open Surgery	21
2.2.3	Minimally Invasive Surgery Environment	21
2.2.4	Endovascular Aneurysm Repair	22
2.2.5	Deformations of the Aorta during Surgery	26
2.3	Image-Guidance for Endovascular Aneurysm Repair	28
2.3.1	Information Required for Endovascular Aneurysm Repair	28
2.3.2	Image Requirements	29
2.3.3	Interventional Images	30
2.3.4	Use of Images during Endovascular Aneurysm Repair	35
2.4	Computer Assisted Aortic Surgery	35
2.4.1	Computer-aided Endovascular Aneurysm Repair	36
2.4.2	Limitations of Current Computer Assistance	38
2.4.3	Future of Computer-Aided Endovascular Repair	39
2.5	Conclusion	41

2.1 Introduction

Surgical procedures are operations where injuries or disorders of the body are treated by incision or manipulation. Large incisions may be required to visualise and access the anatomy to be treated and are responsible for significant pain and lengthened recovery times. With medical imaging, surgeons can see structures of a patient that are hidden below the surface anatomy. Minimally invasive surgery relies on medical imaging and only requires a small incision to insert instruments to operate on a patient. The aorta is the largest blood vessel in the human body and is vulnerable to aneurysms (swelling that can potentially rupture), which can be treated either with conventional surgery or with minimally invasive surgery (endovascular aneurysm repair). Endovascular aneurysm repair (EVAR), which is the main application area for the work described in this thesis, allows surgeons to insert stent-grafts from an incision in the femoral arteries to the site of the aneurysm where large abdominal incisions were previously needed.

This chapter starts with an overview of the field of minimally invasive surgery, where information is given about its history and the pros and cons with respect to open surgery. Broad information is then given on the subject of endovascular repair as an introduction to the two sections to follow.

The second part gives an insight of image-guidance for EVAR. It describes how surgeons use images during aortic surgery, what it means in terms of image requirements and what types of image modalities fulfill these requirements.

A description of computer-aided EVAR concludes this chapter. The computer assistance workflow is illustrated before anything else, followed by a discussion on the topic of registration. Lastly, elements are given as to how visualisation of registered data is performed.

2.2 Minimally Invasive Surgery

Minimally Invasive surgery is the use of ‘Minimal access techniques [...] to perform surgical procedures while avoiding the morbidity of conventional surgical wounds’ [44].

This section first gives a history of both the broad field of minimally invasive

surgery and the more specific topic of endovascular repair. Next, minimally invasive surgery is compared to open surgery in terms of advantages and disadvantages. This is followed by a description of the clinical requirements for the operating room used in minimally-invasive surgery. Information is then delivered to the reader about the clinical background of EVAR, which is the main application area for all the work described in this thesis. A few paragraphs about the deformations the aorta is subjected to during EVAR conclude this section.

2.2.1 History

The history of surgery can be roughly divided into three phases [99]. The first phase spans from ancient times to the mid-nineteenth century, when surgery was ‘rough, rapid, brutal, ablative, and had only limited applications.’ The second phase, which lasted until about 1960, saw complicated and invasive procedures carried out thanks to the advent of anaesthesia and improved resuscitation techniques. The last phase, from 1960 onwards, deals with minimally invasive surgery, i.e. the realisation that ‘operations could be performed more elegantly and less traumatically with advanced instruments’. Although [99] traces back the advent of minimally invasive surgery to the second half of the 20th century, endoscopic procedures, where instruments are introduced into the body to allow internal vision of the organs, had already been developed since ancient times, as explained in the paragraph to follow. A general introduction is given here to both the wide theme of minimally invasive surgery and to the more specific topic of EVAR.

Minimally Invasive Surgery The earliest mention of minimally-invasive surgery is the rectal speculum found in Hippocrates’ treaty on fistula [8]. At the end of the eighteenth century, Bozzini designed a light guide (*lichtleiter*) made of a housing, a candle, tubes and an eyepiece, to reflect the light towards organs while avoiding reflections into the eyes. Progress was subsequently made in the field of endoscopy, as Nitze, a general practitioner, developed the first cystoscope in 1887 and many flexible esophagoscopes were designed during the first half of the nineteenth century. The light transmission, viewing angles and diameters of endoscopic systems were improved in the latter half of the 20th century following innovations in design from Hopkins and van Hell [8]. The first laparoscopic cholecystectomy was performed

by Mouret in France in 1987 [44]. Fluoroscopic control for endovascular repair was reported at about the same time and is covered more in depth in the following paragraph. The late twentieth century brought digitally enhanced images, transmission of high-resolution video images to the surgeon, further advances in the visualisation of the surgical field and improved hand instrumentation [52].

Endovascular Aneurysm Repair Endovascular aneurysm repair is a minimally invasive type of surgery used for treatment of abdominal aortic aneurysm, a disease which is covered more in depth in section 2.2.4. ‘Initial exploration of a solution to the problems involved in intraluminal graft placement [for the treatment of abdominal aortic aneurysm repair] begun in 1976’ [65]. Stents-grafts were designed as prosthetic material to be anchored to the aortic wall and experimental animal studies showed how they could replace surgical suture. Volodos performed a ‘remote transfemoral stent-graft procedure of the thoracic aorta using a self-fixating synthetic prosthesis’ in 1987. The operation was performed under fluoroscopic control and completion aortography showed that the endoprosthesis was positioned as planned. Parodi et al. reported clinical outcomes for 57 patients treated between September 1990 and April 1994 with endovascular stented grafts [66]. Out of these patients, 50 presented with abdominal aortic aneurysms. For 40 of these patients the procedures were considered a success: ‘the stented graft was in contact with normal aortic wall proximally and distally or iliac wall distally, and all flow was observed to traverse the graft without leakage.’ Conventional endovascular repair procedures, where a ‘simple’ stent-graft is used, are unsuited for patients with unfavorable anatomy, i.e. patients with ‘short or angulated necks, aneurysmal extension into either internal iliac artery, or complex aneurysmal involvement of the juxtarenal, paravisceral, and thoracoabdominal aorta’ [73]. Complex endovascular repair procedures deal with these problematic patients, using more evolved stents grafts (the design of which is described in section 2.2.4). These devices were designed to ‘extend the proximal sealing zone from the infrarenal segment to the juxta and suprarenal aorta’, and their first use was reported in 1996.

2.2.2 Minimally Invasive Surgery vs. Open Surgery

As mentioned earlier, minimally invasive surgery is used to avoid the morbidity of conventional surgical procedures. What follows is a comparison of minimally invasive surgery and open surgery with reference to their pros and cons.

Advantages

The use of minimally invasive surgery can be beneficial for patients in that large wounds are avoided and postoperative pain is therefore limited. Decreased postoperative pain means that patients can resume physical activity more quickly, which reduces morbidity due to immobility. In addition to this, hospital discharges are quicker and convalescence shorter. Access of instruments through small portholes also means improved cosmesis [44]. Medical complications and infection risks are also reduced [63].

Disadvantages

The lack of 3-dimensional vision, limitation in tactile sensation and distance separating the surgeons' hands from the target organs are noted in [36] and constitute the main drawbacks of minimally invasive surgery when compared to open surgery. Other downsides include the 'learning curve' of minimally invasive surgery or inadvertent perforation of viscera during the initial puncture [44].

As surgeons cannot directly see the anatomy of the patient they are operating on, they rely on intra-operative imaging for the purpose of visualisation. This introduces modality-specific disadvantages, such as exposing the patient to the risks of ionisation radiation when performing X-ray images guided endovascular aneurysm repair [20].

2.2.3 Minimally Invasive Surgery Environment

Specialised facilities are required for minimal access interventions, especially in the case of endovascular repair procedures, where infections can have severe consequences. Guidelines for carrying out minimally invasive surgery include the need for a sterile environment, as the surgical sites can be exposed to airborne bacteria by microbes settling on the surgeons' clothes or the operative instruments. In addition




Figure 2.1: (a) The aorta is healthy as there is no swelling (b) The aorta is diseased, a bulging can be seen in the abdominal region

to environmental sterility, movement within the operation theatre and the number of personnel should be minimised in order to reduce the count of airborne microbes. Prolonged procedures carry increased risk for surgical site infection, therefore the duration of interventions ought to be limited. Last, the operation theatres should be designed in such a way that minimal access surgery can quickly and safely be converted to open-surgery, should any complications or technical difficulties arise during interventions [43].

2.2.4 Endovascular Aneurysm Repair

The aforementioned benefits of minimally invasive surgery are directly applicable to treating aneurysms of the abdominal aorta with EVAR. The following paragraphs describe abdominal aortic aneurysm and give an account on the different ways this pathology is treated by the medical community. Further attention is then given to the different types of endovascular procedures that are clinically used.

Abdominal Aortic Aneurysm

Abdominal aortic aneurysm (AAA) is a bulging of the aorta in the region of the abdomen as illustrated in figure 2.1, that can potentially rupture and cause major internal bleeding. With an overall mortality rate of 90 percent after rupture, AAA is a major issue, and is dependent on age-related factors, such as atherosclerosis, genetically predisposed alterations of the arterial wall, history of cigarette smoking

and haemodynamic mechanical factors. Risk factors for rupture include the diameter of the aneurysm, elevated blood pressure, the presence of obstructive pulmonary disease and family history [45].

Symptoms that are indicative of AAA are described in [25] and primarily consist of abdominal pain, of the presence of an abdominal mass and of a thrill or bruit over the mass. Suspected aneurysms can be confirmed using medical imaging, such as contrast angiography, echocardiography, computed tomography and magnetic resonance imaging (MRI) [38]. Which technique is more suited to diagnosis is decided by factors such as operator skills (in particular for ultrasound imaging) and clinical urgency (with MRI less likely to be available for very sick patients).

Ultrasonography is a noninvasive (or minimally-invasive) technique, but cannot capture the aorta along many view angles, because of the surrounding bony anatomy, which reduces acoustic access, and due to the deep lying nature of the aorta. On the other hand, CT scanners give more precise details on the shape and size of the aneurysm and also on the relationship between the aorta and its surrounding anatomy. In addition, CT is a particularly appropriate modality to show calcification on the aortic wall when this information is relevant to surgical planning. Last, MRI can be used to evaluate the aorta and to size aortic aneurysms [38], [24].

Abdominal Aortic Aneurysm: Medical Options

Patients presenting with AAA can be offered three treatment options.

Continuous medical management can be followed when the aneurysm has been detected and its diameter does not exceed a certain value (usually 5cm). Several approaches are generally taken to prevent the growth of the aneurysm, such as change in lifestyle and prescription of drugs to control haemodynamics of the patient and to reduce elastin degradation of the vessels [7].

If the risk of aneurysm rupture is deemed too high, open surgery or endovascular repair can be suggested to the patient. In the case of open surgery on both unruptured and ruptured aneurysms, the surgeon typically gains access to the aortic aneurysm by means of a midline incision of the abdomen. The neck of the aneurysm is then exposed and a vertical incision of its anterior wall is done. After this incision, a prosthesis-graft is inserted inside the aorta and sewn to the walls of aorta.




Figure 2.3: (a) The aorta is healthy as there is no swelling (b) The aorta is diseased, the aneurysm does not span the visceral arteries (c) The junctions of the renal arteries are located over the aneurysm.

diameter of the proximal and distal neck using pre- (from the CT scan) and intra- (from an angiogram or ultrasound measurements) operative information, the undeployed stent-graft is advanced into the proximal neck, positioned inferior to the lowest renal artery and subsequently deployed. A completion angiogram (or a bi-plane angiogram) is used to check for endoleaks and if none is detected the delivery system is removed from the patient's body and the femoral arteries are repaired [49].

EVAR operations are performed using image guidance, which is discussed in detail in section 2.3.

As it is a type of minimally invasive surgery, EVAR has the same benefits over open surgery as those described in 2.2.2. More specifically, EVAR has lower incidence of spinal cord injury (as open surgery may interrupt blood circulation to the spine), renal insufficiency and respiratory failures.

Complex Endovascular Aneurysm Repair

Two main anatomical configurations for AAA can occur on a patient: 'Standard' cases are the cases where the aneurysm lies on a part of the aorta where there is no junction with visceral vessels. 'Complex' cases on the other hand are cases where the aneurysm extends above or below sections of the aorta where there is one or several junctions of visceral vessels, or within 15 mm of such vessels (see figure 2.3 for an illustration of a 'standard' and a 'complex' case).




Figure 2.4: ‘Standard’ stent-grafts are sufficient to prevent the blood from exerting pressure on the walls of the aneurysm. For a ‘complex’ aneurysm procedure, a ‘standard’ stent graft would occlude the vessels that are connected to the aorta at the location of the aneurysm, hence the need for branched or fenestrated stent-grafts.

‘Standard’ cases are treated by inserting ‘standard’ stent grafts at the site of the bulging to prevent the blood from exerting pressure on the walls of the aneurysm. For more complex cases, branched or fenestrated stent-grafts are required to also allow the blood to continue to flow down the visceral arteries (see figure 2.4 for an illustration of stent-grafts requirements for the different aneurysm configurations).

Branched or fenestrated stent-grafts require more accurate positioning so that access to the blood vessels is not obstructed. Intra-operative images are used to help the positioning of the graft. Image-guided technologies to assist with the task of deploying prostheses at the right location are discussed in both section 2.3 and section 7.2.

2.2.5 Deformations of the Aorta during Surgery

In image guided surgery, some organs such as the brain (prior to a craniotomy) are often considered as rigid. In addition, individual bones can almost always be considered as single rigid structures [27]. Other organs, in particular tissues in the abdominal area (the area of interest for this thesis), exhibit elastic behaviour, the causes of which generally being breathing, organ-specific motility and external




Figure 2.5: Left (respiratory motion): maximum reported displacement of the left renal ostium in the superior-inferior direction during the respiratory cycle. Centre (cardiac motion): maximum reported displacement of the proximal renal artery during the cardiac cycle. Right (instrument-induced motion): maximum reported displacement of the left renal artery ostium caused by instrument motion during an EVAR procedure.

forces applied by medical instruments [19].

Precise deployment of complex prostheses (i.e. fenestrated and branched stent grafts) is required in the area of visceral vessels, to avoid obstruction of the blood flow to the organs connected to the aorta, as emphasised in section 2.2.4. Renal ostia motion makes stent placement more difficult and happens during respiratory cycles, cardiac cycles and EVAR. Maximum displacements of the left renal ostia (for seven patients) of 2.9mm in the superior-inferior direction and 1.2mm in the anterior-posterior direction (alongside average displacements of 1.6mm and 0.3mm respectively) have been reported in [22] for breathing-induced motion and were higher than corresponding average displacements of the right renal ostia. A study on eight patients found maximum displacements of 2.4mm and mean displacements of 1.2mm for the proximal renal artery (i.e. the section of the renal artery near the aorta) during the cardiac cycle [61]. Instruments were found to be the main cause of deformation in the region of the visceral vessels. A median ostium motion of 6.7mm for the superior mesenteric artery, 6.2mm for the right renal artery and 6.4mm for the left renal artery was reported in [55] for 20 patients undergoing endovascular aortic surgery. Maximum displacements for the above ostia were 13.5mm , 13.5mm and 14.5mm respectively. A summary of maximum displacements for respiratory, cardiac and instrument-induced motion is provided in figure 2.5.

As seen in the previous paragraph, the insertion of rigid instruments leads to significant motion in the area of the visceral vessels. This is particularly the case for highly angulated aortas, where the tortuous geometry is straightened by the rigid structures inserted inside it during endovascular procedures [12].

2.3 Image-Guidance for Endovascular Aneurysm Repair

As introduced in 2.2.4, the structures surgeons are operating on are constantly visualised and monitored using intra-operative images that are updated regularly during the intervention.

This section is concerned with the use of images during surgery. It firstly compares pre-operative and intra-operative imaging for EVAR and then discusses the requirements for the images acquired during the surgery. Next different types of images used during interventions are discussed in particular in relation to the aforementioned requirements. Last, the clinical workflow for image-guided EVAR is given as summary of the previous sections.

2.3.1 Information Required for Endovascular Aneurysm Repair

Endovascular aneurysm repair procedures involve different types of images, which provide crucial information to the surgeon. Information on the aorta can be obtained pre-surgery or inside the operation theatre. Below paragraphs discuss pre-operative and intra-operative imaging.

Pre-operative Imaging Abdominal aortic aneurysm diagnosis is facilitated by the use of images to obtain clinically relevant data such as the maximum diameter of the aneurysm or its peak wall stress [85] [26]. After the surgeons have deemed endovascular repair appropriate for a patient, interventions are planned using pre-operative images to identify the proximal and distal sealing zones and to describe them in terms of size, length and wall characteristics. For complex AAA cases (as described in section 2.2.4), the distance and rotation of the branching vessels is

analysed precisely to design prostheses that accommodate the complex anatomy of the patient [88].

Intra-operative Imaging Endovascular aneurysm surgery is performed under image guidance, where images are required to display the location of instruments inside the patient and are acquired regularly throughout the operation. Instrument positions can be determined with respect to either the aorta or the surrounding anatomy, such as bones (i.e. vertebrae and pelvis) or nearby soft-tissues (nearby blood vessels and organs).

The deployment of a stent-graft inside the diseased aorta requires particular care on the surgeons' part, as the prosthesis can potentially block blood flow to surrounding organs. Stiff instruments can deform the aorta (as mentioned in section 2.2.5), therefore surgeons cannot check for obstruction of the vessels using information on soft-tissue that was acquired prior to the insertion of certain wires. Careful monitoring is needed for stent placement (before, during and after the deployment) and at the end of the intervention to ensure that the blood circulation to visceral vessels is not impaired. This monitoring can be done through live imaging and updated display of soft-tissues [92].

2.3.2 Image Requirements

The below paragraphs discuss imaging requirements with respects to the ways EVAR procedures are carried out (as described in section 2.3.1) and to the need for a highly sterilised environment that was highlighted in section 2.2.3.

To maintain the sterility of the operation theatre, only limited interaction with standard computers is possible on the surgeons' side. This implies that no complex manipulation of images (in particular 3D volume images) can easily be done, such as scrolling through anatomical planes of an MRI or CT scan. Three-dimensional information has to be provided in a form that does not disrupt the surgeon's workflow. Surface rendering of the clinically-relevant anatomy and display of 2D projection images are ways to tackle this issue.

Motion of the surgeons and instruments during endovascular surgery should be hampered as little as possible. Some imaging apparatuses are better suited for the operation theatre than others and are discussed more in detail in section 2.3.3.

Typically, 3D imaging modalities are less appropriate for intra-operative imaging.

The required clinical information, i.e. the instruments, the aorta and surrounding anatomy, must be easy to observe. This involves the ability to discriminate between different soft-tissues and to visualise bone structures inside intra-operative images. Which modalities are a good match for each of these requirements is debated in section 2.3.3.

The need for continuous, and in certain cases real-time, imaging of the patient undergoing surgery was described in 2.3.1. Acquisition times should be low and potentially small enough to allow real-time imaging if necessary. Additionally, ionisation and injection of nephrotoxic contrast agents should be monitored and kept within limits, as images may be obtained regularly throughout the intervention.

Lastly, another factor to take into account is the expense of the modalities in use during surgery, as interventions can take several hours, which reduces access to the units used within the operating theatre.

2.3.3 Interventional Images

The below paragraphs describe the different types of images that are commonly used for minimally-invasive surgeries in terms of dimensions and modalities and describe which types of images are suited for EVAR.

Acquisition time

The images used in minimally-invasive surgery consist of both images acquired during the surgery and images acquired prior to the surgery.

X-ray images and ultrasound images are routinely used during image-guided interventions for reasons discussed in section 2.3.3. Acquisition of tomographic images (like the ones produced using CT and MRI units) during surgery is limited to key points of the operations due to a variety of factors which are described in section 2.3.3).

The imaging apparatus used before an operation is not subject to the same constraints as imaging devices in use during surgery. In particular, 3D tomographic modalities can be used prior to an intervention for diagnosis purposes and to plan the

operation. These images can be used inside the operation theatre, either via a mental mapping process, where the surgeons use their knowledge of the pre-operative anatomy of the patient to guide their instruments during the intervention, or with more sophisticated computer assistance systems which are the topic of section 2.4.3.

Image Dimensionality

The images used during any surgery can have a different dimensionality, depending on their modalities.

3D volumes The intra-operative use of 3D tomographic images can be achieved through different strategies:

- Scrolling through the different anatomical planes of the volume is usually excluded, as it requires interacting with a computer (which generally is hard to achieve in a standard operating room owing to the sterility requirements mentioned in 2.3.2) and disrupts the clinical workflow.
- Three-dimensional images can be rendered as surfaces on a screen. This approach requires pre-segmenting the anatomy of interest. Automatic segmentation is an active topic of research, therefore most interventions require manual segmentation. Manual segmentation is a lengthy process and is likely to be too cumbersome to fit with an intraoperative workflow. For that reason, most approaches involving rendering surfaces extracted from volumetric data require that data to have been acquired pre-operatively.
- Volume images can also be projected onto 2D planes. The resulting images suffer from the same limitations as 2D projections acquired during the interventions, which are discussed in the below paragraph.

2D images Two-dimensional projection images project the whole anatomy of a patient from a particular view direction. Projection beams can go through different parts of the anatomy, which results in an overlap of structures perpendicular to the imaging plane. The interpretation of 2D projection images is therefore problematic if the anatomy that is imaged is complex along the projection direction.

Slices from a 3D volume do not show overlapping structures, but are restricted in that they only present a minor section of a patient's anatomy, and may fail to give enough information to the surgeons when structures of interest are not aligned on the same plane.

Image Modalities

What follows is a succinct review of the main image modalities in use during image-guided surgery. Basic imaging principles are given for each case. The benefits and limitations of interventional use of each modality are then discussed.

X-ray imaging X-ray images are produced by placing a part of the patient in front of an X-ray detector and then illuminating it with a short X-ray pulse. X-rays are absorbed by the structures that fall within the radiation beam, and images are produced as a summation of the X-ray absorption [57].

X-ray images can be visualised on-line and produced real-time, which makes them a modality of choice for EVAR interventions, where images may need to be acquired regularly (as described in 2.3.2). They give clear visualisation of bones and instruments (as those have a high attenuation coefficient) and can display soft-tissues when contrast agents are injected inside the patients. Digital Subtraction Angiography (DSA) highlights the soft-tissue anatomy of a patient by subtracting a mask (i.e. an image where no contrast medium has been introduced) to images where the contrast agent is visible. Examples of intra-operative X-ray images are given in figure 2.6.

X-ray images are projection images and suffer from the problems mentioned in section 2.3.3. To a certain degree, spatial orientation limitations can be overcome through the use of multi-planar fluoroscopy [57]. This approach is restricted by a number of factors: to minimise the harm inflicted on the patient, the number of X-ray images acquired during the intervention must be kept low (because of ionisation), especially when those images are contrast-enhanced (to limit nephrotoxicity). On top of that, the arm of an X-ray machine can only be displaced within a certain range of imaging angles without significantly disrupting the clinical workflow.

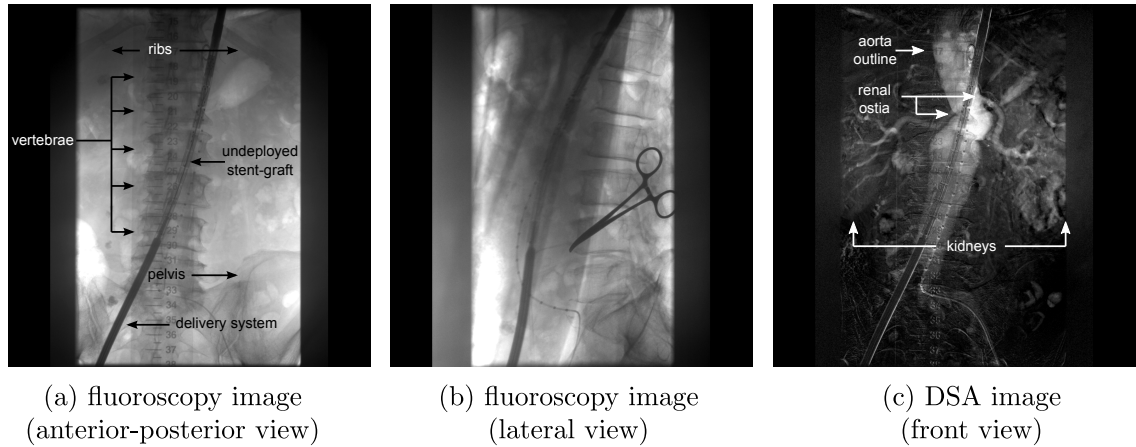


Figure 2.6: (a) X-ray image acquired from an anterior-posterior (AP) view. One can see the bony anatomy of the patient (vertebrae, ribs and pelvis) and instruments inside the aorta (delivery system with undeployed stent-graft) (b) X-ray image acquired at a similar point in time with a different position (lateral view) of the imaging system. The same structures can be seen from a different angle. (c) DSA image from an AP view. Information is shown on the aorta and the blood vessels connected to it and surrounding organs.

Computed tomography Computed tomography (CT) is a procedure whereby X-ray images of a patient are taken from different angles, and a cross-sectional 3D image of the patient is computed using digital image processing. CT scans are three-dimensional and slices can be extracted in any plane, but are mostly extracted in the standard axial, coronal and/or sagittal planes to be viewed by the surgeons [57].

CT scanner acquisition is based on the same physical principles as X-ray projection imaging and is therefore subject to the same harmful ionisation. On average, radiation doses are much higher than those observed for radiographs (e.g. a patient surface dose from a typical abdominal CT scan is 10-20 times more than that of a typical abdominal radiograph [64], and the injection of contrast agents in CT angiography for better visualisation of blood-vessels can induce nephropathy. Attempts have been made to generate 3D volumes of the anatomy of the patient with reduced radiation doses. For instance, cone-beam computed tomography (CBCT) uses a cone-shaped X-ray beam centred on a 2D detector and rotates the beam around the patient's anatomy to generate a series of 2D images, which are used to reconstruct a 3D volume. While the exposure to ionising radiation is substantially lower than for an average CT scan, it remains higher than the dose required to

acquire an X-ray image [32]. Besides causing harmful radiation exposure, the use of CT scanners can disturb the clinical scene as patients may have to be moved in and out and the opening of the CT gantry limits surgical intervention [57]. The two above points are the main limitations to the use of CT scanners during surgery (in particular EVAR). However, CT imaging remains a valuable tool when used prior to interventions, for purposes of diagnosis and operation planning, as scans need not be acquired several times and motion of surgeons about the unit is not an issue.

Magnetic Resonance Imaging Magnetic Resonance Imaging (MRI) is a medical imaging technique that relies on strong magnetic fields to form 3D volume images of the body. The oscillation of the magnetic fields at a particular resonance frequency leads to the excitation of hydrogen nuclei in tissues containing water, which as a consequence emit a radio frequency signal. A receiving coil measures that signal and determines the rates at which excited atoms return to the equilibrium state. The contrast between different tissues is deduced from the different rates obtained at different locations [9].

MRI units produce 3D images and do not emit ionising radiation. Thus, they can be used at little cost to the patient. Spatial and temporal resolution were reported (along with practicability, costs and procedure time) in [11] to be factors that favour the use of conventional X-ray angiography and fluoroscopy over MR imaging to guide endovascular procedures. Most MR units have a closed configuration, which restricts motion of the surgeons around the unit, and makes it harder to manipulate instruments introduced into the patients. Lastly the operating room has to be devoid of ferromagnetic elements, and probably the biggest drawback is the requirement for all the instruments used to be both MR safe and not to effect image quality [29].

Ultrasound Imaging Ultrasound imaging is a modality that uses interaction of high-frequency sound waves with a patient's anatomy to create an image of that anatomy. A transducer sends a pulse into a tissue and receives the backscattered echo after a short while. Backscattered echoes of the sound waves are recorded and used to differentiate structures with different densities.

Ultrasound machines do not produce ionising radiation, can be used in real-time,

are cost-effective and very portable. However, the spatial resolution of the images they produce is limited and can only be improved at the cost of reducing penetration of the waves, meaning that only superficial structures can be imaged at a high spatial resolution. This, in addition to the difficulties in imaging structures behind bone and air, refraction artefacts and scattering artefacts makes interpretation of ultrasound images during surgery (in particular EVAR) harder [80] [14].

2.3.4 Use of Images during Endovascular Aneurysm Repair

For the reasons enumerated in 2.3.3, Endovascular aneurysm repair procedures are planned using pre-operative imaging (e.g. CT scans) and conducted using X-ray imaging.

During the interventions X-ray images allow visualisation of the bones, of the instruments manipulated inside the aorta and can also display the soft-tissues of a patient when they are enhanced with contrast. Such contrast-enhanced images are acquired at crucial steps of the EVAR procedure (i.e. before, during and after stent deployment) to check that major blood vessels are not obstructed by the prosthesis inside the aorta

Image-guidance that only relies on the use of X-ray images suffers from a few limitations. X-ray images can only display soft-tissue information if used in conjunction with (generally nephrotoxic) contrast agents. X-ray images are 2D projections and can only be acquired within a certain range of imaging angles, thus surgeons lack 3D vision of the aorta. Frequent X-ray imaging puts the patient at risk of ionisation and the use of contrast agents can cause nephropathy.

Computer assistance that makes use of pre-operative imaging has been suggested (see [69]) to remedy on these limitations and is described in the next section.

2.4 Computer Assisted Aortic Surgery

As described in section 2.3.4, current image-guided systems for EVAR suffer from issues inherent to X-ray imaging: the images in use do not show information on the soft-tissues of the patient (unless when potentially nephrotoxic contrast agents are used), such as the aorta, and their 2-dimensional projective nature makes interpretation of 3D positions of visceral vessels harder. Solutions to get around these

problems consist of frequent multiplanar checks and/or increased use of contrast agents and are a source of concern for the patient's safety and the duration of the procedure. Computer assistance has been proposed as an alternative to these two solutions and is the main topic of this section.

This section begins by describing a picture of the typical clinical workflow for computer-assisted EVAR as of 2015. The limitations of the systems currently in use are then discussed in the following paragraphs. A description of the potential improvements that can be made to such systems is given as a conclusion.

2.4.1 Computer-aided Endovascular Aneurysm Repair

EVAR procedures involve two types of images, as explained in section 2.3.1: images acquired prior to the interventions, which are used to diagnose the aneurysm and plan surgery, and images acquired during the surgery to guide the instruments. Pre-operative images, when acquired with computed tomography angiography, possess valuable information about the soft-tissue anatomy of the patient. Surgeons can use these images to visualise the complex 3D geometry of the aorta to be repaired, and use this knowledge during the EVAR procedure (by mentally mapping the pre-operative data onto the intra-operative X-ray images), while they are guiding instruments inside the blood vessels of the patient.

Instead of requiring the surgeons to mentally align the aorta that is seen inside the pre-operative CT scan to the X-ray images acquired during endovascular surgery, [69] proposes to automatically rigidly register the 3D pre-operative CT-scan to the 2D intra-operative X-ray images and to display an overlay of the 3D surface of the pre-operative aorta onto the 2D X-ray images during the surgery. The whole process is illustrated in figure 2.7. Approaches to carry out 2D-3D rigid registration of pre-operative 3D CT data to intra-operative 2D X-ray data that are machine-based and only depend on 1) physically tracking the X-ray apparatus, 2) the source-to-detector distance and 3) calibration data can be applied in real-time [83], but do not account for the movements of the patients, which are a major source of rigid 2D-3D registration errors [79]. Updating the 2D-3D registration based on patient motion is the method followed in [69], using information on the bony anatomy that is available in both pre- and intra- operative images.

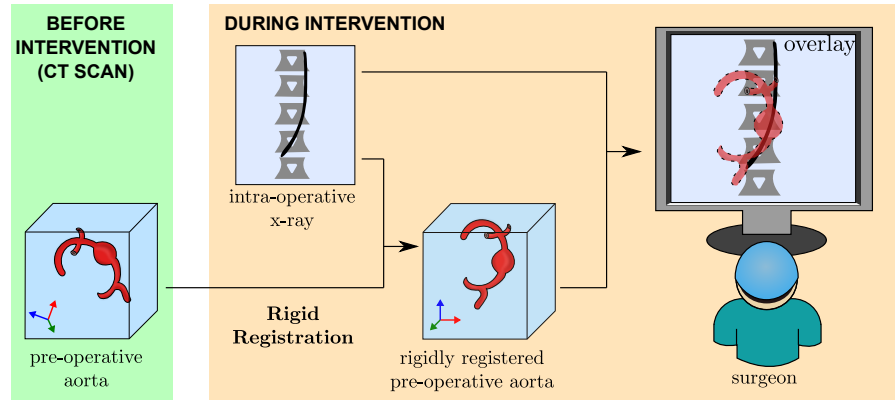


Figure 2.7: The CT scan acquired prior to the intervention is rigidly registered to the intra-operative X-ray images. The updated 3D geometry is overlaid onto the 2D X-ray image. The result is displayed to the surgeons so that they can see the soft-tissue of the patients on top of the projected bony anatomy and instruments.

The acquisition of CT scanner volume images and X-ray images is based on the same physical principles, which makes the registration of CT data to X-ray data essentially a quasi-single modality problem. Different strategies can be applied to bring the 3D information from the CT scan and the 2D information from the X-ray images into dimensional correspondence, i.e. to produce 2D (3D) information from the 3D pre-operative (2D intra-operative) data to align it with the 2D intra-operative (3D pre-operative) information. See [54] for a list of approaches to achieve dimensional correspondence of 3D pre- and 2D intra- operative data, along with a general description of 2D-3D registration methods for image-guided interventions. In [69], digitally reconstructed radiographs (DRRs), i.e. 2D synthetic images projected from the pre-operative CT, are produced and compared to the interventional fluoroscopy images using an intensity-based similarity measure. A global search followed by a hill climbing search strategy is used to retrieve the optimal rigid registration parameters (three translations and three rotations), where for each new set of parameters, a DRR is generated and compared to the X-ray image.

Such a computer assistance system is robust, can be fully automated [97] and has the potential to help the surgeons in guiding the endovascular stent-graft position and orientation to optimise graft and target vessel alignment in order to aid cannulation, selecting landing zones for bridging stents in target vessels and confirming cannulation of correct renal artery in case with multiple renal arteries [69].

Figure 2.8 gives an example of the type of computer assistance that can be




Figure 2.8: (a) Computer assistance allows visualisation of the aorta on top of images where soft-tissues are not normally visible. (b) The 3D geometry can be more easily observed using computer assistance. The position of the visceral ostia can be seen to match well to the DSA image. The 3D rendering provides additional information on the 3D position and orientation of the ostia which can be useful while positioning the stent-graft.

provided using such a system.

2.4.2 Limitations of Current Computer Assistance

The delivery systems and prostheses used during EVAR procedures have limited flexibility, therefore surgeons need to assert large forces to negotiate angulated anatomies. Tortuous aortas tend to be straightened (see figure 2.9) as a consequence of the above, resulting in large (up to $10.5mm$) rigid registration errors when carrying out the 2D-3D rigid registration proposed in section 2.4.1 [12]. Based upon the average $6mm$ diameter of a renal artery ostium, [12] considers $3mm$ the acceptable margin or error for a computer-assisted surgery system to be of clinical use. This means that the benefits of a system described in section 2.4.1 will not be as high for patients presenting with severely angulated aorta as for those whose aorta is not as highly angulated. This is particularly important as it is patients with highly angulated anatomy which are the most difficult surgical cases, and are the cases which would benefit most from the improved visualisation offered by 2D-3D registration. Registration methods that account for the deformable nature of the aorta may be more suitable for these patients and are described in the following section.




Figure 2.9: (a) A stiff wire (in white above) has been introduced inside the aorta and is not encompassed by the highly angulated pre-operative aorta. (b) A DSA image acquired at a similar point in time shows that the aorta has been straightened by the instruments inserted inside it. The positions of the renal ostia (shown with spheres) have been altered.

2.4.3 Future of Computer-Aided Endovascular Repair

Instrument-induced motion is expected to be the largest source of deformation during the operation (see section 2.2.5), with the stiff wires and device delivery systems being the main contribution to rigid registration error [69].

Three-dimensional angiography carried out while the catheter is positioned into the abdominal aorta at the level of the coeliac artery, as mentioned in [92], can account for this main source of deformation. This approach requires the use of CT imaging, which has drawbacks when used inside the operation theatre for the reasons given in section 2.3.3. Disruption of the clinical setting and increased ionisation of the patient during the EVAR procedure mean that 3D angiography cannot be repeated multiple times to recover the deformations that may happen after insertion of the delivery system.

Deforming the 3D pre-operative CT volume has been proposed to represent the intraoperative scene and can be achieved using 2D-3D non-rigid registration [68]. This is illustrated in figure 2.10. The information present within intra-operative 2D X-ray images to drive non-rigid registration algorithms is scarce, as X-ray images only show the bony anatomy of the patients and the instruments inserted inside




Figure 2.10: The computer assistance is similar to the workflow presented previously. However, an additional non-rigid registration is carried out after the rigid alignment of pre-operative CT data to intra-operative X-ray images. A more accurate representation of the soft-tissues is displayed to the surgeons.

them (except when contrast agents are used to produce DSA images). Frequent injections of contrast dye can induce nephropathy, limiting thus the number of DSA acquisitions. Multi-planar angiography is unsuitable due to the above and because it requires rotating the X-ray arm, which disrupts the clinical scene. For these reasons, the soft-tissue anatomy of the patients (in particular, their aortas and visceral vessels) can only be seen for one single plane. As X-ray images are projection images, positions in the 3D space of 2D points picked within the 2D X-ray image cannot be determined with absolute certainty (for each 2D point picked on a 2D fluoroscopy image, the corresponding 3D position can lie anywhere on the back-projection line that connects the X-ray source to the selected 2D point). This makes the task of establishing point-to-point correspondences between pre-operative and intra-operative data harder and potentially diminishes the quality of non-rigid registration in the direction perpendicular to the imaging plane.

The experimental chapters of this thesis are concerned with the implementation of a novel generation of computer-assistance systems for use in EVAR procedures. Three main experiments are presented, which investigate methods that can non-rigidly align a 3D CT scan acquired before surgery to scarce 2D intra-operative information.

2.5 Conclusion

The surgeons do not have direct access (vision and tactile feedback) to the structures they are operating on during minimally-invasive surgery and must therefore use images acquired inside the operation theatre to guide the instruments they are using. Computer assistance has been proposed, and is currently being routinely used in some centres, as a method to bring information available before surgery to the intra-operative scene, so as to improve the visualisation of anatomical structures, and in turn reduce the interventional use of harmful modalities [30].

Endovascular Aneurysm Repair illustrates the points mentioned above. Surgeons use X-ray images to guide instruments within a patient's aorta, as bony structures and metallic objects are projected onto 2D planes. Three-dimensional visualisation of soft-tissues can be achieved by means of computer assistance, that brings the information on soft-tissues visible in a CT scan to the intra-operative scene. During EVAR surgery, stiff instruments are likely to deform the blood vessels they are inserted into, especially when patients present with challenging anatomies (patients with tortuous aortas and patients whose aneurysms extend over visceral vessels). The current generation of computer-assistance system rigidly registers the pre-operative CT scan to the intra-operative X-ray images and fails to account for the changes of shape of the aorta during EVAR procedures. A variety of methods that deform the aorta to match the intra-operative scene are presented in this thesis.

Chapter 3

Literature Review

Contents

3.1	Introduction	43
3.2	Background on Registration	43
3.2.1	Registration	43
3.2.2	Medical Image Registration	44
3.2.3	Registration for Image Guided Interventions	46
3.3	Non-Rigid Registration for Medical Images	47
3.3.1	Main Applications of Non-Rigid Medical Image Registration	48
3.3.2	Challenges of Non-Rigid Medical Image Registration	48
3.3.3	Transformation Models	51
3.4	2D-3D Non-Rigid Registration	61
3.4.1	Statistical models	61
3.4.2	2D-3D non-rigid registration using algebraic reconstruction	62
3.4.3	Monocular model to image registration	64
3.4.4	Graph-based deformable 2D-3D registration	66
3.4.5	Thin-plate splines based 2D-3D registration	69
3.4.6	Summary	70
3.5	Conclusion	73

3.1 Introduction

As described in the previous chapter, a ‘complex’ aneurysm repair procedure can be aided by the use of non-rigid 2D-3D registration. To date, few algorithms have been designed to tackle this problem, subsequently 2D-3D non-rigid registration remains a research area [15]. The first part of this chapter is dedicated to general registration algorithms, while the second part focuses on the broad topic of non-rigid 3D-3D registration. A description on current approaches for the narrower and more recently explored field of 2D-3D non-rigid registration concludes this chapter.

3.2 Background on Registration

In this section, the reader is first given a basic overview of the subject of registration. This topic is then narrowed down to medical image registration, then narrowed further to describe image guided interventions.

3.2.1 Registration

Image registration was defined in [56] as ‘the determination of a one-to-one mapping between the coordinates in one space and those in another such that points in the two spaces that correspond to the same anatomical point are mapped to each other’. Different imaging conditions (e.g. different times of acquisition, different viewpoints, and/or different sensors used) result in differences in the images to be registered (for instance the source and the target images do not display exactly the same structures, having been acquired using different modalities, are altered by different types of noise, etc.). Registration is used in a large range of application, such as environmental sciences, computer vision and medical imaging to name a few [102].

Typically, registration of a source image \mathcal{I}_S to a target image \mathcal{I}_T is the optimisation task of finding the transformation T_{reg} that maximises a similarity value S , which is formalised as follows:

$$T_{reg} = \arg \max_T S(T(\mathcal{I}_S), \mathcal{I}_T) \quad (3.1)$$

Solving this problem can be done in four steps [102]. Firstly, features (such as pixel

intensities or a geometrical object) can be manually or automatically extracted from the source and target images. Secondly, a similarity metric has to be selected in order to measure how well corresponding features from the source and the target images are matched. Thirdly, a transformation model has to be chosen, which defines the class of transformations that can align the source and target images by matching their corresponding features. Transformation models vary in complexity, some involving a relatively small number of degrees of freedom (for example 3D rigid transformations are made of three translations and three rotations giving six degrees of freedom), or a much higher number of degrees of freedom (e.g. non-rigid transformations). The last step is to define a search strategy. At this stage, a transformation that belongs to the transformation model is optimised to maximise the similarity measure in the second step. Different optimisation methods can be applied. Exhaustive search of the solution is computationally prohibitive and is restricted to transformation models with a small number of degrees of freedom. For transformations where an initial estimate is available (using mechanical tracking for example), convex optimisation methods can be used provided the initial estimate is within a convex range near the actual optimum. If the initial estimate is outside the convex range, non-convex optimisations should be carried out.

3.2.2 Medical Image Registration

Medical images are used for a variety of purposes ranging from diagnosis to treatment planning and guidance and also to monitoring of disease progression. Complementary information can be obtained through imaging a patient at different points in time and/or with different imaging apparatuses or also imaging cohorts of patients rather than a single subject. A non-exhaustive list is given below of the ways medical image registration can be employed to benefit from that complementary information. Firstly, multi-modal registration can help overlaying information available from one modality on top of images acquired with a different modality. For example, Positron Emission Tomography (PET) images can be registered with MR images so that PET information can be added to MR images. Secondly, carrying out intersubject registration allows studying the variability of the anatomy (such as the brain anatomy) between subjects or comparing normal subjects with volun-

teers. Lastly, images of the same deformable structure acquired at different times can be registered together, to analyse changes over time to that organ (e.g. serially acquired MR images of children can be registered together to better understand the rate of growth of different structures in the brain) [40].

Medical images are available from different modalities and can have different dimensionalities (2D, 3D, 4D). This, in addition to different fields of view makes the task of finding optimal transformations between source and target images harder. Registration of rigid structures in the 3D space requires searching for a rigid transformation, which is made of six parameters: three translations along the X, Y and Z axis and three rotations around these three axes. Different formulations can be used to represent the class of rigid transformations, which can be described for example: 1) as a rotation (e.g. expressed as a composition of three rotations around the principle orthogonal axes or in the form of quaternions) followed by a translation, 2) in a matricial form, using homogeneous coordinates, or 3) using dual-quaternions, an extension of the quaternions that represents not only the rotational part of the rigid transformation but also its translational part [48]. Alignment of deformable structures, on the other hand, can involve an infinite number of degrees of freedom, vastly increasing the complexity of the optimisation procedure mentioned in 3.1.

Registration methods can roughly be divided into three different types: intrinsic, extrinsic and calibration-based registration methods. Intrinsic registration methods rely on images of anatomical structures and maximise a similarity value between a source and a target image. This similarity measure can be defined on features present inside both source and target images (e.g. points, lines, curves, etc.). Pixels and their intensities form a subset of those features, and registration methods that use this information are called intensity-based. Extrinsic registration methods are those where pre- and intra- interventional markers are inserted into the patient's body. For these methods, correspondences are established between the markers (typically prior to and during an intervention) and a feature-based registration is subsequently carried out. Last, calibration-based registration methods rely on a careful pre-calibration of the imaging devices and the establishment of the position of the operation table with respect to the imaging apparatus during image acquisition, with the components of the imaging devices being tracked in real-time. This family of registration methods was proposed for registration of 3D MR images to

X-ray images and also to integrate CT and DSA images as part of a multifunctional image-guided therapy suite [54].

3.2.3 Registration for Image Guided Interventions

As explained in a previous chapter, computer assistance during image-guided interventions has been proposed to bring information on the 3D pre-operative geometry of a patient inside the operation theatre. The first generation of image-guided system stems from stereotaxy in the field of neurosurgery, where stereotactic frames were attached to the patient's skull to define a space within their brain and use that space to guide instruments. Coupled with stereotaxy, computed tomography allowed the surgeons to match the position of voxels in the CT volume to positions in physical space. Frameless stereotactic systems use trackers to relate the pre-operative scan to 3D points visualised during the intervention (such as fiducial markers attached to the skin or anchored to the bones) or images [15]. The application described in this thesis, i.e. computer assistance for use in EVAR procedures, is concerned with the registration of a pre-operative CT scan to images acquired during the intervention. To date, most image-guidance approaches involve registering pre-operative data (from a CT scan or an MR scan) with same-dimensionality data (e.g. cone-beam CT, 3D ultrasound images, 3D digitised points/surfaces). A detailed review of 3D-3D registration is beyond the scope of the work presented in this thesis, which is primarily concerned with the more difficult problem of registering a 3D pre-operative CT scan with 2D intra-operative data.

What follows is a succinct description of the issues encountered when carrying-out 2D/3D registration. To compare 3D pre-interventional data with 2D intra-interventional data, they must be brought into dimensional correspondence (i.e. 3D data must be converted to 2D or the other way around). Several methods can be applied to that effect. Three-dimensional information can be brought into the 2D space by means of projections. Conversely, 2D information can be brought into the 3D space by means of back-projections. Finally, 3D information can be reconstructed (when using several 2D images). In the case of 2D projective intra-operative images (such as the X-ray projections that are used in endovascular aneurysm repair, the object of this thesis), the dimension perpendicular to the imaging plane is




Figure 3.1: Left: a 3D sphere object is projected as a 2D circle onto a 2D plane. Right: the 3D position of the projected sphere cannot be retrieved from the sole knowledge of the 2D projected data. The only 3D information available from the 2D projected data is that the 3D projected object lies on the back-projection line that connects the source image to the 2D circle.

lost, making the task of finding 3D correspondences between the 3D pre-operative data and the 2D intra-operative images much more difficult in that direction (see figure 3.1 for an illustration of this point). 2D-3D registration methods can be split into all the general types described in part 3.2, that is intrinsic, extrinsic and calibration-based methods. A good review article on current state of the art 2D-3D registration is provided in [54].

An article, [70], divides image guided surgery procedures into two categories, traditional stereotactic approaches where targets deep inside the brain are approached via a probe and open craniotomies, and argues that 'All conventional image-guided surgery systems operate on the assumption that the tissue being operated on can be treated as a rigid body'. Surgery carried out on the deformable anatomy of a patient leads to tissue movement, which rigid registrations cannot account for accurately. This type of minimally invasive surgery consequently requires the use of non-rigid registration, a field which is described in general terms in the following section.

3.3 Non-Rigid Registration for Medical Images

Non-rigid registration has many uses in the field of medical imaging. The main applications of this type of registration are given here. A description of the challenges posed by the non-rigid nature of the registration follows. The last part investigates the main type of transformations that have been proposed as an answer to the

highlighted challenges.

3.3.1 Main Applications of Non-Rigid Medical Image Registration

Non-rigid image registration can be used for a variety of medical applications, which are a subset of the registrations mentioned in section 3.2.2. Deformations, as far as soft tissues are concerned, can be caused by a variety of factors. Outside any clinical setting, motion of organs and their surrounding anatomy occurs naturally as a response to a patient's breathing or their heart beating. This has an impact on the monitoring and treatment of diseases located in those areas subjected to constant physiological motion, as it makes it harder to track and/or specifically target diseased soft-tissues (see [74] for an example of how to model liver motion during the respiratory cycle). On top of respiratory- or cardiac- induced motion, simple movements of the patients (such as coughing, moving head or moving arm) and changes in posture have a potential impact on the shape of certain organs (for instance [82] studies the non-rigid motion of breast between MR images acquired at different times).

Closer to the object of this thesis, deformations of the patient anatomy can occur during surgery. As an example, [10] investigates the use of non-rigid registration to compensate for the deformations (brain-shift) that occur after opening the dura mater during neurosurgery. Contact between instruments and tissues is a common source of changes in shape and is one of the main reasons why deformable alignment between pre-operative and intra-operative structure is required. A method proposed in [90] illustrates this by presenting a concept to combine an intra-operatively extracted surface model of the liver with a biomechanical model extracted from pre-operative imaging to account for the respiratory- and instrument-induced deformations occurring during minimally invasive liver surgery.

3.3.2 Challenges of Non-Rigid Medical Image Registration

Aside from a feature space, most algorithms for non-rigid registration make use of three components, which are first, a similarity measure to quantify how well source and target images are aligned, second, a transformation model that defines the




Figure 3.2: Left: a source image (red) is rigidly registered to a target image (blue). The registration process involves looking for only three degrees of freedom in 2D (one translation on the X axis, one translation on the Y axis and one rotation). Right: a source image (red) is non-rigidly registered to a target image (blue). Several control points on a grid are displaced to deform the source image (an approximate representation of the displaced control nodes is shown with the green arrows in the picture). The number of control points to be displaced is arbitrary, making the search of the optimal transformation harder.

types of transformation applicable to the source image and third, a process to optimise over the aforementioned similarity measure. Similarity measures as described above are defined according to the type of registration that is adopted (intrinsic or extrinsic, as defined in 3.2.2). Due to the constraints inherent to imaging during surgery, only limited information can be used to align the images acquired prior to the intervention with those obtained inside the operation theatre. Many possible deformation fields can be applied to deform the pre-operative data to represent the intra-operative scene. Transformation models are used to constrain these deformation fields so that the data are registered in a physically accurate manner and are significantly more complicated than the models accounting for rigid registrations, as is shown in figure 3.2. The increase in complexity of the family of transformation that are searched for means it is harder to find a model of deformations that provides reasonable results while being computationally inexpensive to fit clinical requirements. As shown in figure 3.2, non-rigid transformation models can be significantly more complicated than those used in rigid-registration, which increases computation time for both the similarity criterion and the optimisation procedure. The main varieties of transformation models are discussed in section 3.3.3. The surgical setting requires fast registration or even real-time updating of the registered anatomy, where ‘real-time’ means ‘sufficiently fast for the application’. Requirements in terms of computation times vary depending on the application: certain tasks need to be computable in a few minutes (e.g. during patient set-up and dur-

ing preparation of the operating table) while delays have to be kept lower than the video frame rate (30-40 *ms*) at critical stages of the surgery. This, according to [15], pushes the limit of what is feasible as of 2008, though efforts have been made since then to accelerate non-rigid registrations. A method proposed in 2008, [94], makes use of graphics-processing units (GPU) for biomechanical modelling of soft-tissues during surgical simulation. Significant speed gains (up to 16.8 \times) were reported using a midrange GPU compared with equivalent central processing unit (CPU) implementation. Solution times lower than 5 *ms* were achieved per time step for the simulation of the indentation of a brain model with 46655 tetrahedral elements.

Validation methods of non-rigid 3D-3D registration are another issue which is an active research area. To date, most non rigid 3D-3D registration have not been validated as thoroughly as rigid registration algorithms. Validation can be carried out using interventional ground truth, by identifying anatomical point landmarks or fiducial markers inside intra-operative images, and checking the differences between predicted and actual positions of those points. With the exception of salient features on anatomical structures, point landmarks can be hard to locate, while implanting fiducial markers is invasive to the patient. In addition, as explained in the clinical background of this thesis, 3D imaging is hard to perform during image-guided surgery (especially so in the case of EVAR), which results in the intra-operative images being 2D and in the 3D intra-operative positions of anatomical landmarks being hard to retrieve. Less invasive but potentially less realistic techniques can also be developed. 1) a physical phantom can be deformed in a way similar to the anatomy of a patient undergoing surgery would be deformed. The deformed phantom is not subject to the same intra-operative constraints as those affecting an actual patient and can therefore be imaged using a CT scanner or an MR unit to produce an ‘intra-operative’ 3D volume. Corresponding positions of landmarks can be selected within the 3D pre-operative and the 3D intra-operative volumes to assess whether the transformation obtained through non-rigid registration returns the right positions for landmarks in the intra-operative space. 2) a simulated deformation can be applied to the pre-operative anatomy, yielding simulated intra-operative data. Registration of the pre-operative images to the simulated intra-operative images produces a transformation that can subsequently be compared to the actual

simulated deformation [15].

3.3.3 Transformation Models

The transformation models used for non-rigid registration of medical images have to satisfy several conditions. Great care should be taken to choose an appropriate number of parameters, to ensure that transformations are flexible enough to model the clinical deformations while limiting parameter determination-induced optimisation time. Additional strategies can be devised to alleviate computation times. B-splines, for example, are only defined in the vicinity of points controlling the overall deformation and, as a result, perturbing the position of one control point only affects the transformation in the neighbourhood of the point [18]. However, registration schemes that use B-splines can suffer from folded deformation fields, an undesirable effect as no one-to-one mapping is achieved between the source and target image and therefore the transformation from source space to target space cannot be inverted [93]. Unfolding can be achieved but requires special measures to be taken (such as those in the method proposed in [60]), that can be harder to enforce at finer resolutions [18]. It should be noted that the bijective nature of a transformation is not always desirable. Objects visible in the source image can be absent from the target image (or vice-versa): this may happen when registering images acquired with different fields of view or when aligning data acquired pre- and post- surgery [40]. Last, folding may be necessary to describe some physical transformations, such as those occurring in the intersubject brain registration where one subject has a large extrinsic tumour and in abdominal registration where fluid and gas filled spaces can appear and disappear between scans [18].

The following section describes a variety of transformation models that have been used in medical image non-rigid registration: motion models, statistical shape models, spline models and biomechanical models. Motion and statistical shape models restrict likely deformations by producing statistical likelihoods from training data. Spline models use functions that produce smooth deformation fields between the images to be registered, while biomechanical models attempt to model the underlying physics behind the deformation of an anatomical structure.

Production and Use of Motion Models

An example of the use of motion models, close to the field of image guided surgery, is their application in radiotherapy treatment. Here their use to align the shapes of deformable organs at different stages of the respiratory cycle has been suggested (see [86]) to track the positions in time of tumours. A patient's breathing results in their organs (e.g. lungs and liver) moving, which can have an effect on the image acquisition, treatment planning and radiation delivery. Wide margins can be used to compensate for this type of movement, but this results in unnecessary irradiation of the anatomy surrounding the diseased tissues, and in turn limits the total doses that can be delivered (thereby limiting the effectiveness of the treatment).

Transformation models are used in the context of matching the shape of an organ at different stages of the breathing cycle, which has been described in [86] and broadly consists of three steps. First, (MR or CT) images are acquired at specific times while the patient is breathing (end-inhale, end-exhale, deep-inspiration). Second, correspondences between the images acquired at different times are established using features such as boundary points (in relation to the surrounding anatomic structures), selected points and/or extracted surfaces. Third, computation of deformation fields is carried out using finite-element models, spline interpolation (thin-plate splines and B-splines) or intensity-based methods.

Motion models have been proposed as an extension of these approaches to the whole respiratory cycle (as opposed to the sole end-inhale, end-exhale and deep-inspiration states) and are illustrated in figure 3.3. This approach involves acquiring a sequence of low spatial resolution 3D images of the patient at different points of the breathing cycle, alongside a high spatial resolution breath-hold 3D image. Once the low-resolution volumes are acquired, they are all non-rigidly registered to the breath-hold scan, yielding a continuous sequence of high resolution volumes of the anatomy of interest (the lesion to be treated and surrounding tissues) [15].

As explained previously, the largest deformations occurring during minimally invasive aortic surgery are ascribable to the instruments inserted inside the vessels of a patient and are not linked to cyclic motion provoked by breathing or the heart beating. Motion models are therefore not fully relevant to the problem of non-rigidly registering the aorta during surgery.

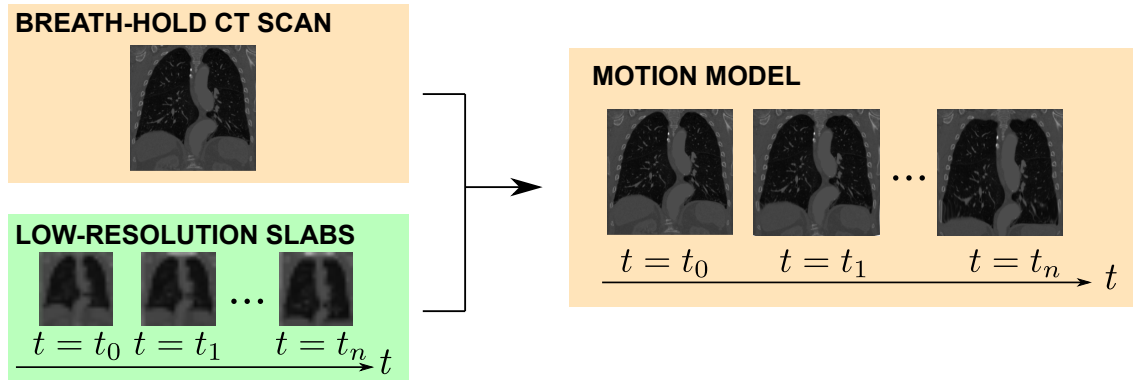


Figure 3.3: A high-resolution CT scan is non-rigidly registered to each slab of a temporal sequence of continuous low-resolution CT scans. This process yields a temporal sequence of continuous high-resolution CT scans that gives a clear picture of the changes in the anatomy of deformable structures.

Use of Statistical Shape Models

Statistical shape models are designed as flexible models of a structure, the shape of which varies across a population. The principal modes of variation are captured from a training set of the structure shapes. Representation of the objects to be modelled is most commonly done under the form of a set of points which are placed in the same way on each item of the training set, as described in [16]. The sets of points of all items are aligned and a ‘Point Distribution Model’ is derived from the statistics of the positions of the labeled points. This model gives the average positions of the labeled points, and Principal Component Analysis (PCA) returns the main modes of variation found in the training set. This procedure, in particular for 3D structures, requires segmenting all the shapes of the training set and manually identifying corresponding landmarks on these surface. Work has been conducted to reduce the tediousness, and manually introduced errors, of both segmenting and picking corresponding points by non-rigidly registering all the volume images that constitute the training set to a reference volume picked from the set (using free-form deformations) and constructing a statistical deformation model by applying PCA on deformation fields rather than on landmarks. Briefly, [81] takes n deformation fields \mathbf{d}_i produced to non-rigidly align the anatomy of a reference subject \mathcal{S}_r to the anatomy of other individuals \mathcal{S}_i and approximates the distribution of the




Figure 3.4: Prior to an intervention, shape models can be generated using a training set of shapes of an object (for example the same bony structures across different patients or the variations in shape of an organ during the respiratory cycle). A distribution of points is used to represent the object for all its different shapes. A shape model is derived from the known correspondences between these points, which yields an average shape and the main modes of variations of said shape. During the intervention, a registration is performed using intra-operative imaging modalities to retrieve the shape of an object. A great reduction of the number of degrees of freedom that are searched for is achieved using the main modes of variations mentioned above.

displacement field \mathbf{d} using a parameterised linear model

$$\mathbf{d} = \hat{\mathbf{d}} + \Phi \mathbf{b} \quad (3.2)$$

In the above equation $\hat{\mathbf{d}}$ is the average of all displacement fields, \mathbf{b} is the model parameter vector, and the columns of Φ are formed by the eigenvectors ϕ_i of the covariance matrix \mathbf{S} , with:

$$\mathbf{S} = \frac{1}{n-1} \sum_{i=1}^n (\mathbf{d}_i - \hat{\mathbf{d}}) (\mathbf{d}_i - \hat{\mathbf{d}})^T \quad (3.3)$$

Incomplete information on a structure (in particular, information acquired during surgery, with all the limitations that intra-operative modalities imply) can be interpolated or extrapolated using such statistical shape models [15]. This is illustrated in figure 3.4. Applications have been reported inside the operation theatre to build models of bony structures (e.g. the vertebrae, pelvis, bones of the lower limbs) using training sets of the same structure across different patients. The mo-

tion of non-rigid structures during cycles (respiratory, cardiac) can be modelled in a similar way. Using a point distribution model at the start of a time sequence (e.g. start to end of respiratory cycle), subsequent point distributions can be propagated at each step of the time sequence using non-rigid registration. A statistical shape model can be generated from the set of all point distributions [39].

Instruments manipulated inside soft-tissue can cause deformations in different directions and with variable strengths, which are difficult to predict using information available before interventions. Generating a statistical shape model of the aorta is a complex task and due to time constraints this would probably not be feasible during surgery of the aneurysm. Interaction of soft-tissues with instruments can be simulated with multiple runs of biomechanical models prior to an intervention and this can help capture the main modes of variation of a patient-specific anatomy during an intervention. This was proposed in [59], where the authors run finite-element simulations of the insertion of a transrectal ultrasound probe to observe different changes in the shape of the prostate of a patient, and use this information to carry-out non-rigid registration between two orthogonal cross-sectional ultrasound images and the pre-operative prostate. This approach, however, requires pre-operative knowledge of how the intervention is likely to unfold and that the interaction between different structures (anatomy of the patient and instruments) can be modelled easily using finite-element analysis. Many instruments are inserted during EVAR procedures from either the left or the right iliac artery, which makes it hard to predict the course of the surgery and to produce a model that can account for the wide variety of possible clinical scenarios. To the knowledge of the author of this thesis, no article that follows this approach for the specific application described in this document has yet been published.

Spline Models

Spline models use spline functions to interpolate a smooth displacement field between a source and a target image.

Affine transformations can be computed as a linear function of all points (x, y, z) of the source image. Polynomial transformations generalise affine transformations and can be used to model non-rigid motion, but their behaviour is unpredictable (especially when the degree of the polynomials used is high) and they do not always

adequately capture local deformations of anatomical structures [3].

Piecewise polynomial representations can be used to describe more accurately local deformations and their use has been reported for various registration tasks such as motion compensation, patient positioning, multi-modality imaging and brain warping [96]. Of all the piecewise polynomial representations, B-spline representations have attracted a great deal of interest in the medical imaging community. Briefly, B-splines can be used to build piece-wise functions where the pieces are low-order (e.g. lower than 3) polynomial curves (Bézier curves) which approximate a set of ordered control points and are anchored to the first and the last point. The piece-wise functions can be used to interpolate data over a regular grid of control points or to approximate irregularly-spaced data (the data do not match the control points of the regular grid for the latter case). B-spline interpolation requires solving a linear system to compute a set of weights based on the input data, which can be done efficiently as B-spline basis functions have local support (this results in a linear system with a sparse matrix). In addition to reducing the computation times to solve the B-spline interpolation linear system, local support makes it easier to optimise a displacement field produced while using a regularly spaced B-spline grid: a translation of a control point from the grid results in local deformation around this node only, so deforming a source image to a target image using B-spline interpolation only requires computing the position of voxels that lie in the neighbourhood of translated control points (as opposed to all the voxels) [40].

In general, B-spline interpolation requires using a regularly-spaced grid whereas scattered input data can only be approximated using B-splines. For some clinical scenarios the exact matching of corresponding landmarks inside a source and a target image is necessary. Such landmarks are not always regularly spaced and B-spline interpolation is therefore not suitable to interpolate a displacement field from these landmarks. Thin-plate spline interpolation generates a function that interpolates those data while minimising a bending energy (defined as the squared norm of the second derivative) [40]. Thin-plate spline interpolation requires solving a linear system, where the dimension of the matrix is directly related to the number of selected landmarks. A more complete description of thin-plate spline interpolation is beyond

the scope of this section, as the method is further detailed in the first experimental chapter of this thesis. The selection of landmarks within both source and target image is prone to errors whether it is carried out manually or automatically. Typically, points lying on flat regions of an image/surface are harder to select and match than points that are located on salient features (e.g. a corner) of an image. Approximating thin-plate spline schemes have been proposed as an answer to this issue. Approximating thin-plate spline registration allows corresponding source and target points to be loosely matched to get a smoother displacement field (with respect to the aforementioned bending energy). Strategies that use anisotropic landmark error have been devised to tailor the approximate matching, so that the looseness of the matching be direction-specific [77]. Additional information, such as orientation attributes, can be incorporated into thin-plate splines to preserve the shape of rigid structures embedded in elastic materials [76]. Work has been further conducted to extend thin-plate splines by controlling the locality of the transformation used for registration purposes [78]. This new framework, *Gaussian Thin-plate splines*, derives radial basis functions from Gaussian forces, where a standard deviation parameter is associated to each control point to define its spatial influence.

The first experimental chapter of this thesis investigates a point-based registration of the 3D aorta surface from a pre-operative CT scan to 2D intra-operative X-ray images. The main goal of this chapter is to model the large position uncertainties in the 3D space that comes from registering the 3D pre-operative CT scan to 2D intra-operative data. With this purpose in mind, thin-plate splines were used with anisotropic landmark errors, the directions of which were tailored to run parallel to the X-ray projection direction. Adjusting the orientation or spatial influence of particular landmarks was not viewed as directly relevant to the projection uncertainty mentioned above and was therefore not carried out.

Biomechanical Models

Biomechanical models are based upon the biomechanics of soft biological tissues. Soft tissues are composed of cells and an extracellular matrix, which react to change in the mechanical environment. Continuum biomechanics are concerned with the study of fundamental concepts, postulates and principles, with the formulation of

constitutive relations that describe material behaviour and with the solution of initial-boundary-value problems. Soft-tissues, in particular arteries, share similar mechanical properties to those of rubber, such as the anisotropic response, viscoelastic character and thermoelastic behaviour. With unloaded natural configurations rubber-like materials exhibit an isotropic behaviour while most soft-tissues have an anisotropic behaviour. Constitutive relations have been studied within several theoretical frameworks, such as finite elasticity, membrane theory, mixture theory, models that account for growth and remodelling, viscoelasticity, thermodynamics models, porous media models [42]. Different types of biomechanical models are available, such as mass-spring models, finite-element models and boundary-element models [13].

Biomechanical models can be used for registration purposes, where the knowledge of the biomechanical properties of soft-tissues is used to produce smooth deformation fields. Non-rigid registration using biomechanical models requires discretising the anatomy of interest into small volumetric elements, a non-trivial task, and assigning material properties to all the structures that are being modelled. Defining the material properties presents significant difficulties, since measurements are generally not readily available during interventions, soft-tissues exhibits different behaviours in-vivo and in-vitro and difference in properties can be significant between patients (in particular, diseased tissues often have different mechanical properties than the surrounding healthy anatomy). In addition to those difficulties, the information necessary to drive biomechanical models is hard to retrieve during image-guided interventions. Stress measurements cannot easily be obtained during operations. Sparse information on the displacements of the anatomy can be retrieved using trackers (tracked pointers, laser-range scanners, stereopsis) or intra-operative imaging modalities (ultrasound and X-rays), albeit with potential measurement errors. The best way to apply measured displacements to a model is still a topic of research, as imposing erroneous displacements to a model generates unrealistic forces and potentially unnatural deformations. A last practical consideration is the execution times.

For most operations, model-updated images must be available within, at most, a few minutes, and in some cases have to be displayed in real-time. Biomechanical models usually involve systems of equations composed of large matrices, requiring a

lot of memory. Processing power can also be in high demand in the case of dynamic biomechanical models where integration schemes are used with a lot of iterations, making it harder to quickly find solutions. Simplifications of the models can be used to alleviate these requirements but may come at the cost of fidelity [13]. The use of Graphics-Processor Units (GPUs) coupled with a finite-element formulation suited to soft-tissue simulations is reported in [95] and allows real-time solution for models made with up to 16000 tetrahedral elements.

Biomechanical models have previously been applied to endovascular aneurysm repair procedures to predict aneurysm rupture, analyse migration and endoleaks at the stent and also to deal with catheterisation issues. The novel use of finite-element models to anticipate the position of a guide-wire inserted inside a patient and to predict the deformation of the aorta during endovascular repair has been proposed in [47]. Particular attention is given in this article to the intra-operative positions of the iliac arteries and to the changes in size of vessels, as those affect the sizing of the devices (i.e. stent-grafts) that are inserted inside the patient. Roughly, the authors of [47] adopted the following approach:

- First, the aortailiac structure was extracted and tissue was classified in accordance with the amount of calcification
- second, a mesh of the aorta was generated
- lastly, a guidewire was initially constrained to the pre-operative vessels centerlines of the patient and allowed to relax.

Validation was conducted by projecting the simulated guide wire onto intra-operative X-ray images of the patient and comparing the positions to those of the real guide-wire. The model parameters were tuned on 10 patients and validated retrospectively on 4 patients, with a reported mean error of $2.3mm$. The method described above was used in [35], where position errors were reported between the simulated and the intra-operative guide-wire. An additional correction step was suggested to correct the simulated guide-wire positions based upon the intra-operative 2D positions of the wire. These studies were concerned with the intra-operative positions of the guide-wire mainly as a means of validation, and their authors acknowledge that displacements of the ostia of visceral vessels (which are the main

clinical points of interest for accurate positioning of branched or fenestrated prostheses inside complex aneurysms) were not taken into account in the validation process.

Endovascular aneurysm repair procedures are carried out using 2D images for intra-operative guidance. Most of the 3D-3D registration methods presented above can be performed off-line or necessitate acquiring 3D images during surgery. They cannot therefore be directly translated to the work presented inside this thesis.

Summary: Transformation Models for Computer Assisted Aneurysm Repair

- Motion models are primarily used to model cyclic motion such as respiratory or cardiac motion. Most of the motion that occurs during EVAR is instrument-induced and therefore acyclic, so motion models are not relevant to EVAR procedures.
- Statistical shape models describe the variations in shape of a structure and have been used in image-guided surgery by pre-operatively running finite-element simulations on the likely course of an intervention, deriving the main modes of variation of the anatomical structure of interest from all the outcomes of the simulations and using the main modes of variation to constrain the registration that takes place during the intervention. EVAR procedures vary in the types of instruments they involve and the ways these instruments are manipulated inside the patient. For these reasons, running pre-operative simulations that try to replicate what is happening during an intervention is not practically feasible.
- Spline models have been used to obtain smooth displacement fields by interpolating sparse displacements from a source to a target image. These approaches vary in their interpolating and/or approximating behaviour. A method is presented in this thesis that relies on the use of approximating thin-plate splines, where anisotropic errors are used to model the uncertainty along the X-ray projection direction.
- Biomechanical models use the biomechanics of soft-biological tissues to predict the deformations that occur when constraints are applied to the anatomy of

a patient. To date, many approaches that involve biomechanical modelling of anatomical structures require extracting correspondences from pre-operative images and intra-operative images or tracked features. The second experimental chapter of this thesis presents a method to use intra-operative information on the instruments in use during EVAR to produce constraints that drive a registration which uses a biomechanical model of the aorta.

3.4 2D-3D Non-Rigid Registration

The main applications of non-rigid 2D-3D registration are reconstruction of 3D volumes from generic anatomical 3D models and patient-specific 2D images, or warping of pre-operative 3D volumes to match intra-operative 2D images.

3.4.1 Statistical models

In [28] the shape of an anatomical structure (a femur) is recovered from a 3D shape model and a few 2D intra-operative X-ray images. First, statistical information (i.e. mean shape and main modes of deformation) is inferred from 3D data from a population of patients, using PCA, so as to generate the shape model. Once the shape model is generated, the rigid parameters and non-rigid parameters of the transformation for a new patient are recovered by:

1. extracting contour points from both the X-rays (2D points) and the shape model (3D points)
2. computing 3D X-ray contour points from the 2D X-ray contour points and the shape model
3. establishing correspondences between a first 3D set of points (X-ray extracted points) and a second 3D set of points (shape model points)
4. finding the non-rigid transformation that minimises the sum of square differences (of corresponding points) by
 - (a) using an Iterative Closest Point approach to find the rigid parameters
 - (b) using the Downhill Simplex algorithm to find the non-rigid parameters

Registration experiments are carried out on simulated data that consist of a 3D model (with 4 deformation modes) of the distal part of a femur and two simulated X-ray images taken from different view points around the 3D model. The Root Mean Square (*RMS*) distance between the reference model and the deformed model are computed for several X-ray projections (where each projection ray corresponds to a 2D contour point from an X-ray image). Using 10 ray projections leads to a RMS distance equal to 1.3 *mm*, 50 ray projections to 1.13 *mm* and 200 ray projections to 0.99 *mm*.

3.4.2 2D-3D non-rigid registration using algebraic reconstruction

Prümmer et al. ([71]) propose a method to register a high-contrast CT volume to a series of low-contrast fluoroscopic X-ray images.

An energy functional is first defined in the case of a mono-dimensional registration, that consists of a distance measure \mathcal{D} (between source – template T – and target – reference R – data) and a regularisation term \mathcal{S} (which measures the smoothness of the deformation field):

$$\mathcal{J}[u] := \mathcal{D}[R, T; \mathbf{u}] + \alpha \mathcal{S}[u] \quad (3.4)$$

where \mathcal{D} and \mathcal{S} can be defined as the sum of squared differences and the curvature-regulariser respectively:

$$\mathcal{D}^{SSD}[R, T; \mathbf{u}] := \frac{1}{2} \int_{\Omega} (T_u(x) - R(x))^2 dx \quad \text{and} \quad \mathcal{S}^{curv}[u] := \frac{1}{2} \sum_{l=1}^d \int_{\Omega} (\Delta u_l(x))^2 dx \quad (3.5)$$

where $T_u(x)$ denotes the deformed source volume and $\Omega =]0, 1[^d$. The optimal displacement field is the solution of the Euler-Lagrange equation:

$$f_u(x) + \alpha \Delta^2[u](x) = 0 \quad (3.6)$$

where $f_u(x)$ is an external force term, computed from the intensity of the deformed source image and of the target image. In the mono-dimensional case, $f_u(x)$ is a function of a dissimilarity measure $D_u^I(x)$ (e.g. $D_u^I(x) = R(x) - T_u(x)$ if a sum

of squared differences is used as a similarity measure between source and target volumes). When a 3D source volume is registered to one or several 2D target images, such a dissimilarity measure cannot be used.

For a non-rigid registration of a 3D volume to a series of 2D images, Prümmer et al. propose to compute a 3D dissimilarity measure between the deformed 3D volume and the series of 2D images, based upon the computation of all the dissimilarity measures between each 2D image and corresponding 2D deformed volume projection. An external force is derived from this 3D dissimilarity, which is used to produce an updated displacement field by solving the modified Euler-Lagrange equation:

$$f_u^R(x) + \alpha \Delta^2[u](x) = 0 \quad (3.7)$$

with $f_u^R(x)$ is the external force that is computed from the 3D dissimilarity measure between the 3D source volume and the 2D target images.

An iterative optimisation process is applied to retrieve the optimal displacement:

- an initial deformation field $u_{(0)}$ is set to 0 (no displacement)
- a 3D dissimilarity measure is computed between the 3D deformed volume (i.e. the initial volume deformed using $u_{(k)}$) and the series of 2D images
- an external force is computed from the 3D dissimilarity measure
- a new ‘optimal’ displacement $u_{(k+1)}$ is computed from this external force by solving equation 3.7
- the process is repeated until convergence of the algorithm

Experiments conducted on simulated data (first, a 3D cube deformed to a sphere, and second spheres aligned along a helical trajectory with an artificial deformation) and real CT data (with a simulated sinuodal synthetic deformation that simulates the breathing of a patient) show a quantitative improvement of the distance (sum of squared differences of intensities) between the ground truth and the 3D source volume after it has been registered to the reference volume. The error between the source and target volumes before and after non-rigid registration is reduced by a factor comprised between 2 and 3.

The authors of [71] acknowledge that their approach can only retrieve small deformations (such as those induced by the breathing of a patient). As discussed in the clinical background chapter of this thesis, the instruments inserted inside the aorta during endovascular aneurysm repair can be large, especially when patients present with tortuous vessels. In addition to this, the method in [71] intends to match structures from CT data to X-ray images using voxel/pixel intensities. Soft tissues cannot be well visualised inside X-ray images, unless when contrast agents are introduced into the arteries. Contrast injection can induce nephropathy and thus should be kept to a minimum during the interventions. The approach proposed in [71] is expected to produce best results when a few such images are used, something that can rarely happen in minimally invasive surgery of the aorta for the reasons mentioned above.

3.4.3 Monocular model to image registration

Groher et al. ([33]) propose a new method for 2D-3D non-rigid registration of vascular structures that are observed in a 3D angiographic scan and in a single 2D Digital Subtraction Angiography (DSA) image. First, a vascular model is extracted from the 3D angiographic scan, combining both a region-growing algorithm to segment the scan, and a topological thinning of the resulting segmentation, which yields centreline points ($\{X_i\}$) and associate radii of the vessels. Next, the 2D DSA image is enhanced: its noise is reduced and its tubular structures are enhanced. The latter step is performed following the approach proposed in [50]: linear filters are non-linearly combined and elongated symmetric line structures are searched for at different scales (maximum responses are obtained at a scale proportional to the width of a line-type structure).

Lastly, a 2D distance map is produced that yields high intensity values around the vascular structure. The centrelines of the 3D model are then projected into 2D and matched with the 2D vessels centrelines: the 3D centrelines $\{X_i\}$ are moved to their new 2D positions $\{Y_i\}$, with a deformation field ϕ . For this purpose, an energy function composed of two terms is defined as follows:

- an external energy term that measures the similarity between the 3D projected points and their projected 2D counterparts. The similarity is computed using

the neighbourhood of projected 3D points inside the 2D distance map.

- an internal energy term that aims at preserving the length between each centreline point

The energy function is minimised with two gradient-descent based optimisation techniques: At each step of the optimisation process, approximating thin-plate splines (with $\{X_i\}$ and updated $\{Y_i\}$ as source and target points respectively) are used so as to obtain a smooth deformation field of the volume.

Registration experiments have been conducted on synthetic data (artificial vascular model with a simulated deformation that occurs perpendicular to the image plane) and simulated data from real data (first a 3D pre-operative CT scan and a 3D intra-operative CT scan are acquired, then a reference deformation is computed via an interpolating thin-plate spline. A synthetic DSA is then generated, using the pre-operative CT scan and the reference deformation).

Accuracy of the method is computed for the synthetic and the simulated anatomical data by calculation of the Euclidean and angle error. On synthetic data, the algorithm yields a mean Euclidean error of 0.55 mm (min: 0.24 mm , max: 1.04 mm) and a mean angle error of 7.9 degrees (min: 1.1 degrees , max: 19.5 degrees). On simulated anatomical data, the algorithm yields a mean Euclidean error of 10.4 mm (min: 9.13 mm , max: 11.67 mm) and a mean angle error of 12.9 degrees (min: 12.6 degrees , max: 13.2 degrees). The authors of [33] claim that these results are an improvement of the error by 16.3% on the previous feature-based approach they presented in [34].

During an EVAR procedure, the anatomy that surrounds the aorta (e.g. the vertebrae and the structures connecting these to the aorta) probably have an influence on instrument-induced deformations, which do not spread evenly through the entire arteries, but rather, are much more localised, as illustrated in figure 3.5. The deformations caused by stiff instruments during endovascular aneurysm repair are likely to violate the assumption of minimum bending energy made by [33], which makes this approach un-applicable to our application.




Figure 3.5: Pre-operative aorta: the intra-operative wire is outside the pre-operative aorta shape (drawn at the site of the aortic bifurcation in this picture), meaning that the aorta surface was deformed during the intervention. The aorta is connected to surrounding vertebrae by ligaments and lumbar arteries, represented in blue. Minimisation of bending energy: the aorta shape is deformed in such a way that the length of vessels is preserved if a bending energy is minimised. This hypothetical situation results in deformations of the anatomical structure linking the aorta to the vertebrae that are unlikely to take place during an EVAR procedure. Anatomical constraints: the connections between the aorta and the vertebrae restrict the aorta up-and-down motion. From conversation with a surgeon at St Thomas' Hospital, in such situations where the iliac artery straightens, a shortening occurs by the artery walls rucking up as shown in (right). The method proposed in [33] would reject this shortening.

3.4.4 Graph-based deformable 2D-3D registration

Liao et al. ([51]) present a new method to non-rigidly register a 3D CT volume with either one or two 2D fluoroscopy images for abdominal aortic aneurysm interventions. First, both 3D and 2D data are used to extract information about the location of the aorta in space:

- The abdominal aorta CT volume is segmented with a graph-cut approach. The centrelines of the aorta are then obtained with a sequential topological thinning process. Last, a 3D graph is generated from the centrelines.
- The fluoroscopy image is segmented, and a 2D distance map of the aorta is derived from the segmentation.

Registration of the 3D graph to the 2D distance map is performed in two steps.

1. A rigid registration is performed on three segments of the 3D graph, these are 1) the renal arteries, 2) the iliac arteries and 3) the abdominal aorta, with the

connectivity of these segments being maintained. This provides the non-rigid registration step with an initialisation.

2. Non-rigid registration of the 3D data to the 2D data is done as follows:

- An energy functional is defined as the weighted sum of a difference measure term, a length preserving term and a smoothness term. The difference measure term accounts for the distances of the projected 3D centrelines to the 2D segmented aorta. The length preserving term accounts for the difference in lengths of original edges and corresponding new edges. The smoothness term accounts for the differences in lengths of neighbouring node displacements.
- A registration of the 3D volume and the 2D fluoroscopy image is done, by minimising the energy functional with a gradient descent approach (Broyden-Fletcher-Goldfarb-Shanno method). Such an optimisation scheme is guaranteed to converge globally in case the function being optimised is convex. Global convergence cannot necessarily be achieved in the non-convex case[101].

The method that the authors propose is computationally efficient, since the use of a 2D distance map does not require the correspondences between 2D and 3D graphs to be explicitly established. The weights given to each term composing the energy functional have to be chosen carefully to ensure that privileged deformations both match the intra-operative scene and abide with smoothness requirements. It is unclear from [51] how appropriate weights were chosen.

Registration experiments have been conducted on phantom data (a 3D phantom of the abdominal aorta and DSA acquisitions of the phantom after introducing an artificial deformation) and simulated data from real data. To simulate these data, first a 3D pre-operative CT scan is registered to a 3D intra-operative cone beam CT volume. This yields a ‘natural’ deformation field. The 3D pre-operative data is then deformed either using a synthesised deformation field, or the ‘natural’ deformation field described above. A synthetic DSA is then generated, using the pre-operative CT scan and the reference deformation.

The known deformation fields are used as ground truth for quantifying the registration performance. The registration for phantom data is evaluated using Target Registration Errors (*TRE*): the distance between all transformed nodes of the aorta graph and corresponding ground truth is computed as the average of all individual TREs, where individual TREs are computed as the 3D Euclidean distance between a transformed node and the corresponding ground truth node [51]. Fiducial markers are not used in [51], which prevents the computation of Fiducial Localisation Errors (errors in identifying the correct position of a fiducial marker in either source or target images) or Fiducial Registration Errors (TREs computed at the positions of fiducial markers). Fiducial points are only a subset of all the points that can be used to compute TREs and therefore FREs are less clinically relevant than TREs [40]. TRE values are all below 2 *mm* when registering the 3D CT volume with two views, and depend on the initial offset for a single-view registration scenario. Accuracy of the non-rigid registration of anatomical data is evaluated and compared to the accuracy of rigid registration. Mean errors go from 8.55 *mm* (phantom data), 15.19 *mm* (clinical data with synthetic deformation) and 9.72 *mm* (clinical data with ‘natural’ deformation) to 0.76 *mm*, 0.95 *mm* and 0.82 *mm* respectively.

Low non-rigid registration errors (below 2 *mm*) have been achieved in [51] using a two-view settings, while much larger non-registration errors are reported for a single view (when the initial rigid registration error is high). Repeated injection of contrast should be limited due to nephrotoxicity of the dye agent. This makes the scenario of acquiring two subsequent angiograms unsuited for most endovascular aneurysm repair. Capturing two different views of the same scene with a single contrast injection can potentially be done using a bi-plane angiography system. This, however, involves ‘aligning two planes to position the area of interest within the isocentre of both imaging tubes’, which is a cumbersome procedure likely to disrupt the clinicians’ workflow[84]. This, in addition to the extra radiation dose emitted by the second X-ray tube, is probably the main reason why bi-plane angiography systems are normally not used during EVAR surgery. As a result, the scenario described in [51] (that is, using two contrast-enhanced X-ray images acquired at a different angle) is not suited for most endovascular aneurysm repair procedures.

Additionally, automatic segmentation of 2D DSA images is hard to achieve because of the non-uniform contrast propagation and the difficulty to isolate the aorta from surrounding structures (blood vessels, nearby organs and artefacts resulting from the subtraction of instruments and bones). Manual segmentation is time-consuming and can also be hard to achieve for the same reasons as those above. The approach proposed in [51] was tested either on X-ray images from a phantom undergoing a deformation or on simulated X-ray images projected from a simulated deformed CT scan. These two scenarios allow segmentations of the X-ray images but they do not account for the more complicated case of interventional contrast-enhanced X-ray images segmentation.

3.4.5 Thin-plate splines based 2D-3D registration

Raheem et al. ([72]) propose to use interpolating thin-plate splines in a 2D-3D framework in order to register a 3D segmented pre-operative aorta with a 2D intra-operative X-ray image. A set of 3D source points $\{s_i^{3D}\}$ and 2D target points $\{t_i^{2D}\}$ are manually picked at anatomically relevant positions on the 3D segmented aorta surface and in the 2D images. Additional ‘fixed’ points are also picked, which [72] assumes not to move during the operation. From each 3D source point s_i^{3D} and 2D target point t_i^{2D} a 3D target point t_i^{3D} is computed as the closest point to s_i^{3D} on the back-projection line that connects t_i^{2D} to the centre of projection. A registration on the aorta using the two sets of 3D source points and 3D target points is performed, that yields a warped aorta, which is aligned with the 2D images. Experiments show an improvement in the proposed non-rigid registration over a rigid registration, when focusing on the particular position of the visceral ostia.

Registration experiments have been conducted on interventional data (a 3D pre-operative volume and a 2D DSA image) from 8 patients.

Registration errors are calculated using sets of ‘gold-standard’ 2D positions (u_{su}) and 3D positions (x_{su}) of renal ostia picked by an experienced surgeon in the 3D CT volume and the 2D DSA image. Non-rigid registration errors are computed as the distance from the position of the ostia after an observer-based thin-plate spline warp to the back-projection line that joins u_{su} and the centre of projection.

Mean reported registration error was 3.11 *mm*.

The approach proposed in [72] adapts a mono-dimensional (3D-3D) registration method to a multi-dimensional 2D/3D case. It does so by making assumptions on the positions in the 3D space of points that have been manually picked on a 2D DSA image. These assumptions (that the 3D positions of target 2D points are the closest points to corresponding 3D source points on back-projection lines) are likely to be incorrect and result in warps with a higher bending energy than that of the deformations taking place inside the operation theatre.

3.4.6 Summary

A summary of the non-rigid 2D-3D registration methods is provided with table 3.1 on page 71 and table 3.2 on page 72.

The methods described in this section are all expected to have difficulties handling the case of abdominal aortic aneurysm repair for the following reasons:

- Statistical Shape Models (e.g. [28]) can be used to reconstruct a rigid body from a set of 2D images of this body and a 3D generic model of this body, where the model has been derived from a population of corresponding rigid bodies. The same approach can be followed for non-rigid structures which exhibit predictable motion in particular directions. The 2D/3D non-rigid registration problem presented inside this thesis is concerned with retrieving the deformation that an aorta undergoes during an EVAR procedure. A large set of aortas shape from the anatomy of a same patient is needed to build the shape model of the patient's aorta. These shapes cannot be acquired during the operation, since intra-operative imaging during EVAR is limited to 2D imaging. Prior to the intervention, a substantial collection of shapes can be generated through the use of biomechanical models, using a method similar to the one proposed in [59]. However, defining the biomechanical simulation prior to the intervention can prove particularly difficult, as one cannot easily predict the number of instruments that will be used during the operation at a specific point in time, their type or the femoral artery (left or right) into which they will be introduced. As a consequence, obtaining a warped aorta from a set of 2D X-ray images and a 3D statistical model is not the preferred direction of the work presented in this thesis.

Method	Intensity-based Feature-based	Data	Accuracy	Validation
Fleute and Lavallée 3.4.1 [28]	Feature-based	simulated data: 3D model of a femur and 2D simulated X-rays	RMS 1.3mm to 0.99mm	ground truth (simulated model of the femur)
Prümmer et al. [71] 3.4.2	Intensity-based	synthetic and anatomical data	weighted sum of square differences measure. anatomical data mean error, before registration: 1.1, after registration 0.6	gold-standard transformation (user-defined transformation to simulate breathing)
Groher et al. 3.4.3 [33]	Feature-based/intensity-based	synthetic and real anatomical data	anatomical data: mean Euclidean error 10.4mm, mean angle error 12.9 degrees	gold standard transformation (deformation field extracted from pre-operative and intra-operative 3D data)
Liao et al. 3.4.4 [51]	Feature-based	phantom data and simulated/real anatomical data	real deformation, mean error: 9.72mm (prior to registration), 0.82 (after registration)	gold standard transformation (deformation field extracted from pre-operative and intra-operative 3D data)
Raheem et al. 3.4.5 [72]	Feature-based	interventional data	mean registration error 2.7 mm	2D ground truth (DSA image)

Table 3.1: Summary of the non-rigid 2D-3D methods: name of the method, type of the method (intensity-based/feature-based), type of data used for validation (synthetic data, anatomical data) accuracy of the method on anatomical data, validation method.

Method	Number of views	Application	Computation time
Fleute and Lavallée 3.4.1 [28]	$1 \leftrightarrow \infty$	3D shape reconstruction from X-ray images	$< 1 \text{ min}$
Prümmer et al. 3.4.2 [71]	$1 \leftrightarrow \infty$	CT + X-ray image(s) registration	$2 \text{ min to } 120 \text{ min}$
Groher et al. 3.4.3 [33]	1	Vascular	$0.39 \text{ s to } 679.3 \text{ s}$
Liao et al. 3.4.4 [51]	1	Aortic Aneurysm Repair	$< 0.1 \text{ s}$
Raheem et al. 3.4.5 [72]	1	Aortic Aneurysm Repair	$< 5 \text{ s} + \text{ manual point picking}$

Table 3.2: Summary of the non-rigid 2D-3D methods: name of the method, number of 2D views, application of the method, computation time

- Computationally highly-intensive methods (e.g. [71]) are excluded to avoid disrupting the clinical procedure. According to the clinicians' knowledge the insertion of a stiff guide-wire is the main cause of aorta deformation during minimally-invasive aortic surgery, therefore only one registration has to be conducted during the procedure to account for this deformation. As registrations need not be performed continuously, a limited interruption of the procedure (up to about 10 minutes) can be allowed once.
- Methods that rely solely on the intensities of pixels within X-ray images can only make use of one DSA image with all additional X-ray images not being contrast-enhanced (to reduce nephrotoxicity). Assumptions have to be made when using only one view, due to the large error distributions parallel to the projection plane. In addition to this, the propagation of contrast is not uniform inside 2D X-ray images, which means that some parts of the aorta might not be visible, making it harder to rely on the intensities of pixels in corresponding areas.
- As mentioned above, reconstructing a 3D intra-operative volume from a single DSA image (see section 3.4.3 and section 3.4.4) requires making assumptions on the deformations that are taking place during the intervention: 1) the

length of the vessels does not change during the intervention 2) the transformation minimises a smoothness criterion [33], [51]. The instruments used during EVAR procedures are responsible for deformations that are unlikely to minimise the bending energy of an aorta.

- Although the approach presented in [72] fits well within the clinical workflow and can realistically be used during an EVAR procedure, it relies on assumptions in the positions in the 3D space of points manually selected on 2D planes. These assumptions can lead to unrealistic deformations, especially since the deformation fields applied to the aorta using this method do not minimise a bending energy.

3.5 Conclusion

The registration approaches that have been presented above cannot be translated directly into the work presented inside this thesis as they would require multiplanar angiography, automatic segmentation of data or use of statistical shape models. Simulations of intra-operative deformations using biomechanical models are hard to perform and require making assumptions about the course of the surgical procedure. The methods presented inside this thesis intend to make the most of the information available inside the operation theatre and therefore do not consider biomechanical simulations that are limited to the sole data acquired before EVAR interventions. Registration methods that make use of 3D information acquired during the operation cannot be directly applied to the interventional case where only 2D images are acquired, unless strategies are devised to convert 2D information into 3D information. Selecting points on an X-ray plane is an easy way to retrieve 2D information from intra-operative data, and methods that make use of a few manually selected points can fit well with the clinical workflow. The 3D positions in space of those 2D points obtained manually are known to lie along the back-projection lines that joins the X-ray source to the 2D point. The work described in [72] assumes the 3D position on the back-projection lines and is a first step towards getting 3D information from 2D data to drive a registration algorithm during EVAR surgery. The next chapter addresses the issue of not knowing the 3D position in space of manually-selected 2D points by allowing 3D points selected on an aorta surface to

be loosely matched with their 2D counterparts along corresponding back-projection lines. This is followed by a chapter that uses biomechanical modelling of the aorta to simulate its interaction with instruments during an EVAR procedure. The last experimental method presents a novel method to extract intra-operative information on the calcifications of the aorta and show how to use that information to drive a non-rigid 3D-3D registration.

Chapter 4

Spline-based 2D/3D Registration using Error Ellipsoids to Account for Projection Uncertainty

Contents

4.1	Introduction	76
4.2	Theory	76
4.2.1	Interpolating Thin-Plate Splines	77
4.2.2	Approximating Thin-Plate Splines	79
4.2.3	Approximating Thin-Plate Splines for Medical Image Reg- istration	81
4.2.4	Approximating Thin-Plate Splines for Computer Assisted Aneurysm Repair	83
4.3	Experiments	88
4.3.1	Data	88
4.3.2	Methods	90
4.3.3	Validation	92
4.4	Results	93
4.5	Discussion	96
4.6	Conclusion	100

4.1 Introduction

This chapter is concerned with the novel use of an interpolation framework, thin-plate splines, for 2D-3D registration purposes. Thin-plate splines minimise a functional which represents the bending energy of a thin-plate passing through a set of points and can therefore be used as a flexible deformation model with intuitive results when registering same-dimensionality data (e.g. 2D-2D, 3D-3D data) [75]. Registration of 2D X-ray images with 3D CT volumes can be achieved by back-projection of the 2D data into the 3D space, which results in position uncertainties parallel to the X-ray projection directions. Previous thin-plate spline registration approaches make assumptions on the three-dimensional positions of back-projected 2D points. The registration method presented in this chapter allows 3D points to be loosely matched to 3D lines back-projected from 2D points and is hypothesised to be more accurate than 1) the previous approaches mentioned above and 2) rigid registration methods that do not account for deformation occurring to soft-tissues.

The first part of this chapter deals with the theory behind registration using interpolating/approximating thin-plate splines. An account of the experiments performed to test the initial hypotheses is then given to the reader. Results of these experiments are last presented and discussed.

4.2 Theory

What follows is a general background of the theory behind the method proposed in this chapter. First, a registration framework to match point landmarks, interpolating thin-plate splines, is presented to the reader. This is continued by a description of approximating thin-plate splines, an extension of the above framework that allows loose alignment of the landmark points in specific physical directions. A few examples of conventional ways to use approximating thin-plate splines for medical image registration are then discussed. Lastly, an explanation of how the approximating thin-plate spline framework was tailored to account for the specificities of 2D-3D non-rigid registration for computer assisted aneurysm repair concludes this section.




Figure 4.1: Left: TPS interpolation yields the function that passes through the points $(s_1, t_1), (s_2, t_2), \dots, (s_4, t_4)$. Minimisation of a bending energy functional ensures the smoothness of the curve. Right: For each point s on the S axis, a corresponding point t can be found on the T axis.




Figure 4.2: Left: a mapping function u is defined that moves all points s_i in s_1, \dots, s_n to $t_i = u(s_i)$. Right: the displacement field is interpolated for all other points s in the source image.

4.2.1 Interpolating Thin-Plate Splines

Thin-Plate Spline Registration

Thin-Plate Spline (TPS) is an interpolation method that originates from Duchon et al. ([23]) that emulates the behaviour of a thin metal plate on point constraints. Thin-plate spline interpolation yields the function that passes through a set of points while minimising a bending energy functional, as illustrated in figure 4.1.

TPS interpolation can be used for registration purposes: displacements from a source image to a target image can be interpolated at each position of the source image. Provided two sets of source points $\{s_1, s_2, \dots, s_n\}$ and corresponding target points $\{t_1, t_2, \dots, t_n\}$, a mapping function u can be defined, with $t_i = u(s_i)$ for all i . For all points s in the source image, the displacement $u(s)$ can be computed by means of TPS interpolation, as described in figure 4.2.

Mathematical Foundations

The notations below correspond to those from Rohr et al. in [77]. For the sake of convenience, changes have been made to the notation of source points, no longer denoted by $\{p_i\}$ but rather $\{s_i\}$. Likewise, target points are no longer denoted by $\{q_i\}$ but rather $\{t_i\}$. As in Rohr et al. ([77]), the mapping function is denoted by u and its components by u_k , with k in $1, \dots, d$, where d is the image dimension (that is, $d = 3$ for complex endovascular procedures). The thin-plate spline bending energy is expressed as

$$J_m^d(u) = \sum_{k=1}^d J_m^d(u_k) \quad (4.1)$$

with m order of derivatives in the functional and

$$J_m^d(u_k) = \sum_{\alpha_1 + \dots + \alpha_d = m} \frac{m!}{\alpha_1! \dots \alpha_d!} \int_{R^d} \left(\frac{\partial^m u_k}{\partial x_1^{\alpha_1} \dots \partial x_d^{\alpha_d}} \right)^2 dX \quad (4.2)$$

The analytic solution of this TPS bending energy takes the form (for all $k = 1, \dots, d$):

$$u_k(s) = \sum_{v=1}^M a_{k,v} \phi_v(s) + \sum_{i=1}^n w_{k,i} U(s, s_i) \quad (4.3)$$

with ϕ_v monomials and U suitable radial basis functions (as defined in [77]). For example

$$U(s, s_i) = \begin{cases} |s - s_i|^{2m-d} \log |s - s_i|, & \text{if } d \text{ is even} \\ |s - s_i|^{2m-d}, & \text{otherwise} \end{cases} \quad (4.4)$$

The coefficients $a_{k,v}$ describe the rigid part of the mapping function (that is the translation and rotation), while the coefficients $w_{k,i}$ represent the non-rigid part of the mapping function. The reader is referred to figure 4.3 for a 2D illustration of the analytical solution.

The coefficients $a_{k,v}$ and $w_{k,i}$ are computed by solving the following system of equations:

$$\begin{aligned} Kw + Pa &= v \\ P^T w &= 0 \end{aligned}$$

In this system, $K = (K_{ij} Id)$, where $K_{ij} = U(s_i, s_j)$ and Id is the $d \times d$ identity

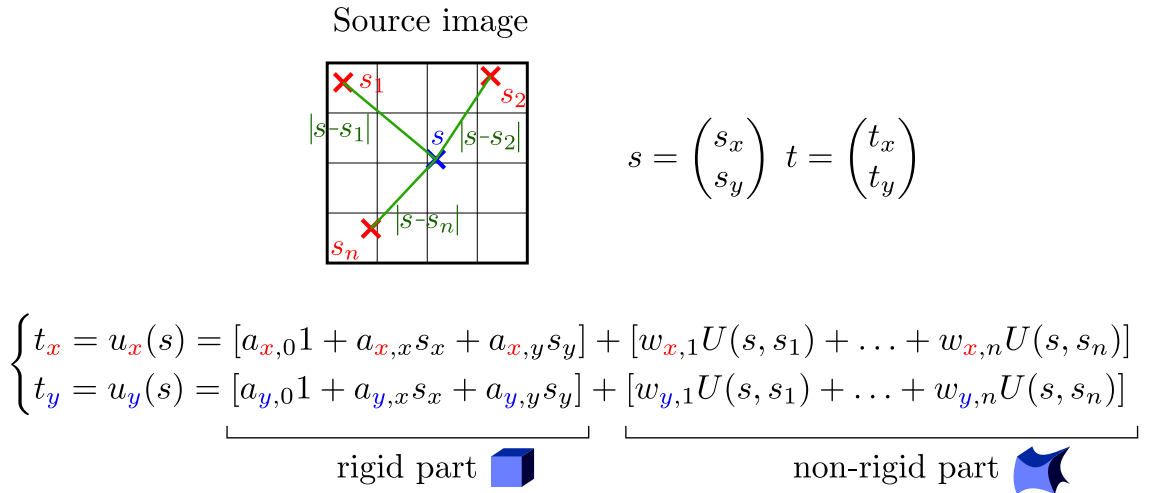


Figure 4.3: 2D representation of the analytical solution of the bending energy functional. The equations are separated into a rigid part, function of the x and y component of the source point s , and a non-rigid part, function of the distances of s to each TPS source point s_i . t_x , the x component of t is a function of rigid coefficients a_x , and non-rigid coefficients w_x . t_y is a function of rigid coefficients a_y , and non-rigid coefficients w_y .

matrix. $P = (P_{ij}I_d)$, where $P_{ij} = \phi_j(s_i)$. $v = (v_1^T, \dots, v_n^T)$, with $v_i^T = (t_{i,1}, \dots, t_{i,d})$.

4.2.2 Approximating Thin-Plate Splines

The TPS framework can be enhanced to allow the source and corresponding target points to be loosely matched, as illustrated in figure 4.4. In such a case, the energy




Figure 4.4: Interpolating TPS: s_1 and s_2 are matched to t_1 and t_2 respectively. Approximating TPS: s_1 is matched to t_1 while s_2 and t_2 are loosely matched, taking into account the geometry of the source and target images. Loose matching of points is represented by an error ellipsis. The longer the axis of an ellipsis, the looser the matching in the corresponding direction.




Figure 4.5: Left: source points (s_1 and s_2) are picked inside the source image by some user. The user picks corresponding target points (t_1 and t_2) inside the target images. Several target points t_2 can be picked along the edge of the surface as illustrated with the red curves. This uncertainty can be approximately modelled in the TPS framework using anisotropic error ellipses, as seen inside the right rectangle. Right: anisotropic error ellipses describe the multi-dimensional uncertainty in picking points. The corresponding covariance matrix is computed using the length of major axis (\vec{v}_2 - blue), secondary axes (\vec{v}_1 - red), and the rotation matrix R from (\vec{X}, \vec{Y}) to (\vec{v}_1, \vec{v}_2) .

functional to be minimised becomes:

$$J_\lambda(u) = \frac{1}{n} \sum_{i=1}^n (t_i - u(s_i))^T \Sigma_i^{-1} (t_i - u(s_i)) + \lambda J_m^d(u) \quad (4.5)$$

The left part of the sum measures the weighted distance between $u(s_i)$ and t_i for all i . For all i this distance is weighted by the covariance matrix Σ_i (the reader is referred to figure 4.5 for an illustration of covariance matrices). The parameter λ in the equation above can be used either to encourage smooth deformations (when lambda is set to a high value), or to encourage deformations that accurately match the source points to their corresponding target points, taking account of their respective covariance matrices.

Similarly to the interpolating case, the analytic solution of the approximating TPS energy functional can be stated (for all k) as:

$$u_k(X) = \sum_{v=1}^M a_{k,v} \phi_v(X) + \sum_{i=1}^n w_{k,i} U(X, s_i) \quad (4.6)$$

Where the coefficients $a_{k,v}$ and $w_{k,i}$ are computed solving a similar linear system:

$$\begin{aligned}(K + n\lambda W^{-1})w + Pa &= v \\ P^T w &= 0\end{aligned}$$

In the equations above, W represents the error matrix and is defined as

$$W^{-1} = \begin{bmatrix} \Sigma_1 & & 0 \\ & \ddots & \\ 0 & & \Sigma_n \end{bmatrix} \quad (4.7)$$

4.2.3 Approximating Thin-Plate Splines for Medical Image Registration

Approximating thin-plate spline registrations have been used in many fields of medicine. The main medical applications are multifarious, as illustrated with the following (non-exhaustive) list of examples: retrieving inter-patient and multi-modal registrations (MR to CT) of the brain ([77]), modelling metallic-implants-induced distortion fields of MR images ([67]), finding deformations undergone by soft-tissues in mammograms at different stages of breast cancer treatment ([53]), aligning the MR volume images of the prostate to images acquired during transrectal ultrasound-guided prostate biopsy ([58]), registering images (CT) of the brain acquired before surgery to images (CBCT) acquired during image-guided interventions ([91]).

The landmarks used in approximating thin-plate spline registration can be selected manually (i.e. points are picked by an expert inside a 3D volume, on a 3D surface or inside a 2D image), semi-automatically or automatically. Automatic methods rely on information present within the medical images (i.e. areas with high variations in the intensities) and as such are sensitive to noise and to the imaging conditions. Manual picking of points is prone to user-errors and can be more time-consuming, depending on the nature of the images in use and the number of landmarks requested to produce accurate registrations. Errors in the selection of points can result in erroneous correspondences between a few source points and corresponding target points, which in turn alters the overall displacement field. For




Figure 4.6: Thin-plate spline registration for same-dimensionality data. Top: a source point (s) is manually picked on the 3D source aorta. Three different attempts (t_1 , t_2 and t_3) at picking a corresponding target point are made on the 3D target aorta. Bottom: the source point s is smoothly matched to a possible target point t with an error ellipsoid e representing the uncertainty in location that is observed in corresponding above frames. Left: s is picked at a junction of vessels, resulting in a small uncertainty in all directions. All axes of e are small. Centre: s is picked on the surface of the aneurysm, making it harder to accurately localise t . Axes of e are large in all directions. Right: s is picked on the segment of the aorta above the renal arteries. Uncertainty in the localisation of t is high in the direction parallel to the centreline of the segment and so the corresponding axis of e is large while the two other axes are small.

each pair of points that have not been picked correctly it is therefore desirable to allow the source point to be approximately matched to the target point: the source point is ‘exactly’ matched to a point located in the neighbourhood of the target point in such a way that the bending energy of the overall displacement field is minimised. Approximating thin-plate splines rely on this method of loosely matching source and target points and can be used for purposes of regularisation.

Anisotropic local variations inside the images or on surfaces can make it more challenging to pick landmarks in certain directions than in others. The top frames of figure 4.6 illustrate this with an ad hoc example of a 3D-3D thin-plate spline registration of a 3D aorta surface with a corresponding 3D aorta surface (this example differs from the problem of 2D-3D rigid registration this thesis is concerned with). Modelling the uncertainties in picking corresponding landmarks can be achieved using the anisotropic error ellipses mentioned in section 4.2.2: the less easy to pick a point in one direction, the larger the corresponding error ellipse axis. This is depicted in the bottom frames of figure 4.6. In medical image registration, points

located at a junction of anatomical structures can be picked with limited error in any directions, which translates to a small error ellipse. The location of surface points is well defined in the direction perpendicular to the surface but is unknown in all other directions, resulting in a large error ellipse except in the direction perpendicular to the surface.

The approach described in this thesis is not as concerned with uncertainties in manual point picking as with modelling the large unknown error distribution in the direction parallel to the X-ray projection direction. In particular, the definition of the error ellipse axes is subjected to a different process. How the approximating thin-plate spline framework was tailored to account for this error distribution is explained in the following section.

4.2.4 Approximating Thin-Plate Splines for Computer Assisted Aneurysm Repair

A registration approach based on interpolating thin-plate splines has been presented in [72] for computer assisted endovascular aneurysm repair (see the literature review of this thesis). This approach derives 3D positions of target points from the manually selected 3D source points (picked on the aorta surface) and 2D target points (picked on a single contrast-enhanced fluoroscopy image). Each 3D target point is computed as the closest point to its source point counterpart on the back-projection line that joins the corresponding 2D target point to the X-ray source. These assumed 3D target point positions can result in deformation fields which are not smooth and therefore lead to incorrect registrations.

This chapter presents a novel approach for 2D-3D non-rigid registration designed to tackle the above issue: Thin-Plate Splines (TPS) plus Projection Uncertainties (which is referred in this document as TPS+PU). Information is manually selected from a 3D aorta surface and two fluoroscopy images acquired from different views during the intervention, prior to stent-graft deployment. More accurate 3D positional information can be retrieved from two 2D images if they are separated by a large angle. A study in [5] was concerned with assessing the influence of the angle between two views on the quality of the 3D reconstruction of a wire visible within two views. To this purpose, a series of images (100 X-ray views going from $+90^\circ$

to -90° around the caudal-cranial axis) were acquired around a phantom of the cerebral vasculature into which a guide wire was introduced. A 3D representation of the vasculature and wire was generated from the series of X-ray images and the 3D ground truth position of the wire was extracted from the 3D volume. Different pairs of X-ray images (where different angles separated the first from the second views) were used to reconstruct a 3D wire, the position of which was subsequently compared to the 3D ground truth. Results showed smaller reconstruction errors for angles comprised between 30 and 150 degrees.

Acquiring fluoroscopy images at a different position requires moving the X-ray arm, which can disrupt the clinical setting when large displacements are involved (see the clinical background of this thesis). Angles circa 30 degrees to an AP view are considered in this study, as such views are routinely used during operations and so have very little effect on clinical workflow, while still enabling accurate reconstruction of 3D information. Contrast agents improve the visualisation of soft-tissues inside X-ray images, but their injection can induce nephropathy and as such should be kept to a minimum. For this reason, the work presented here assumes that one of the above X-ray images will be contrast-enhanced, and accurately show vessel ostia positions. However the other view may or may not be contrast-enhanced and is unlikely to accurately show ostia positions, but will contain some information on the boundary of the aorta, either from contrast, or inferred from instrument positions, or calcium visible on the aortic wall, or using interventional digital tomosynthesis which is discussed later. After an initial rigid registration, aortic deformation is based on preoperatively determined fixed landmarks and a small number (less than 10) of manually identified moving landmarks in the two fluoroscopy images. As shown in figure 4.7, 3D source points $\{s_i\}$ and 2D target points $\{t_i^{2D}\}$ are picked on the aorta surface and in the fluoroscopy images. Each 3D target point $\{t_i^{3D}\}$ is calculated from $\{t_i^{2D}\}$ as the closest point to $\{s_i\}$ on the back-projection line $\{l_i\}$ that joins $\{t_i^{2D}\}$ to the X-ray source. The way covariance matrices $\{\Sigma_i\}$ that represent the uncertainty in matching points s_i to t_i^{3D} along lines l_i are computed is best described in two points. Firstly, the major axis of the error ellipsoid is set to the direction of l_i . In order to allow s_i to be matched to t_i^{3D} anywhere along l_i an infinite length should be given to the major axis. Infinity values cannot be input into the approximating thin-plate spline scheme, and therefore a large value of 1000




Figure 4.7: Approximating TPS can be used to account for the uncertainty along the back-projection lines: 3D-3D TPS registration is performed using landmarks $\{s_i\}$ and $\{t_i^{3D}\}$ with associated error ellipsoids with their major axes along the back-projection lines.

was given, which was physically based on the approximate fluoroscopy focal length. Secondly, the two other axes are set to zero length. The method devised for this chapter does not aim to model the uncertainty in manually picking point. Rather, it puts forward a solution to the 2D-3D registration problem, where 2D intra-operative images are aligned with 3D CT data acquired before the surgery takes place. In this context, using non-zero lengths for the two minor axes is impractical, as the influence of the first axis cannot be as easily assessed if errors due to the manual selection of points are also accounted for. Using these parameters for the ellipsoids and setting a weighting function λ to 1 allows loose matching of landmarks along the back-projection lines. Unlike standard backprojection methods, the approach described here does not require corresponding landmarks to be identified in both fluoroscopy images. Such corresponding landmarks are difficult to find in clinical images. Instead TPS+PU models the lack of information perpendicular to the 2D fluoroscopy images using non-isotropic error ellipsoids, where the major axis lies along the back-projection lines.

The solution of this system is computed and applied to the aorta surface. This novel method is referred here as the ‘Thin-Plate Splines plus Projection Uncertainties’ strategy (TPS+PU) (see figure 4.8), and is compared to an ‘Interpolating Thin-Plate Splines’ strategy (subsequently referred as interpolating TPS) which registers the aorta surface to each 2D image in turn and does not use non-isotropic error ellipsoids (see figure 4.9).

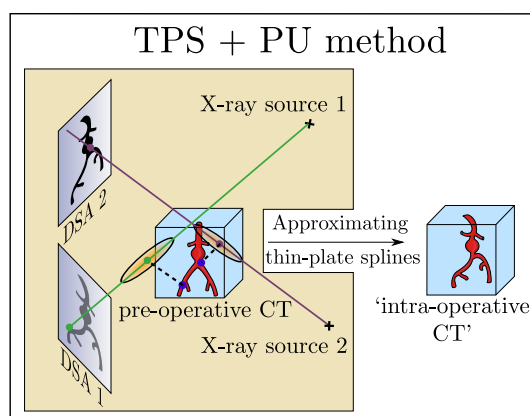


Figure 4.8: Thin-Plate Splines + Projection Uncertainties: 3D source points are picked on the aorta surface, corresponding 2D target points are picked within the two DSA images. 3D target points are computed as the closest points to 3D source points on the back-projection lines that join the X-ray source and the 2D target points. Projection uncertainties take the form of ellipsoids, with a main axis which lies on the back-projection lines and two secondary axes with a 0 length. Registration is carried out using the 3D source points, 3D target points and error ellipsoids.

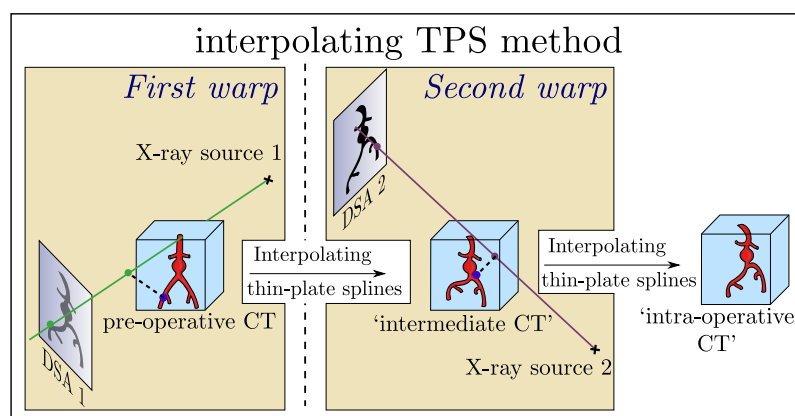


Figure 4.9: Interpolating Thin-Plate Splines: (left) 3D points are picked on the aorta surface and corresponding 2D target points are picked on the first view (DSA 1). 3D target points are computed as the closest points to the 3D source points on the back-projection lines joining the X-ray source and the 2D target points. A first warp of the aorta surface is produced by registering the aorta surface using interpolating thin-plate splines with the 3D source points and the 3D target points. The updated aorta surface is registered to the second DSA image using the same method as described above. This second warp of the aorta yields the registered aorta surface.

Software (built upon the image-guidance system for endovascular repair described in [69]) has been designed to perform 2D-3D non-rigid registration during an endovascular repair procedure, with the imperative of keeping clinical disruption to a minimum. Several steps are performed prior to the intervention to minimise disruption inside the operation theatre. These are based on the knowledge available before the operation: a high-quality volume image (from a CT scan) has been acquired, in which points can be picked and structures of interest segmented. The surface of the aorta is segmented from a pre-operative CT scan and covers the junction of the internal and external iliac arteries and the ostia of the left renal artery, the right renal artery, the superior mesenteric artery and coeliac artery. Fixed points (i.e. the 3D points where the target equals the source), are picked at the centre of lumbar vertebrae, based on anatomical knowledge of the physicians. A rigid registration, which aligns the vertebrae of the patients from the coordinate system of the pre-operative CT scan to the coordinate system of the endovascular surgery, is done before carrying out the non-rigid registration proposed in this chapter. This makes the selection of landmark points on the vertebrae a valid approach, as these vertebrae positions are found automatically during the intervention.

As described in the clinical background section of this thesis, displaying an aorta surface on top of fluoroscopic images enables visualisation of soft-tissues and provides the surgeons with valuable information on the 3D positions of the ostia of the vessels of interest. The software that has been designed, allows manual selection of a limited number of source and target points (on the 3D aorta surface and on the 2D fluoroscopy images) after the insertion of a stiff guide wire inside the aorta, which occurs immediately before the deployment of a stent-graft at the site of the aneurysm. Since this event only takes place once during the intervention, a small interruption (up to a few minutes) of the clinical workflow can be allowed to provide the surgeons with better means of visualising the anatomy they are treating. A limited number of points (up to about 10) needs to be selected by a skilled expert, which can be done well within the aforementioned time constraints. To assist the expert, an automatic update of the aorta surface is provided every time a 3D source point is manually selected on the aorta surface along with its 2D counterpart on an intra-operative plane. While the aorta is being registered to the fluoroscopy images, its 3D surface is overlaid on the 2D images, and different levels of transparency

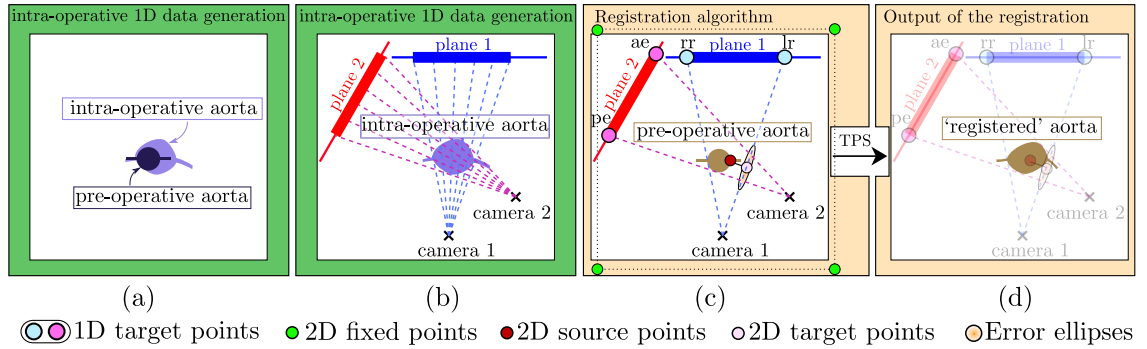


Figure 4.10: (a) interpolating TPS is applied to the pre-operative aorta to obtain the intra-operative aorta. The ostium of the left renal artery is fixed while the ostium of the right renal artery is moved. (b) The simulated intra-operative aorta is projected onto two planes. (c) 1D intra-operative points: left, right renals (lr, rr) and anterior, posterior edges (ae, pe) are visible on the projected intra-operative aorta on each 1D intra-operative plane. 1D target points are selected on the 1D intra-operative planes and the points on the relevant aortic area closest to the back-projection lines are selected as corresponding 2D source points. The points on the back-projection lines closest to the 2D source points are selected as corresponding 2D target points. For the TPS+PU registration, error ellipses are produced along the back-projection lines. (d) Result of applying non-rigid TPS + PU registration on the aorta using 2D source and 2D target points.

can be applied to see either the soft-tissues or the bony anatomy and instruments.

4.3 Experiments

4.3.1 Data

Three sets of data were used: (1) simulated 1D-2D synthetic data, (2) simulated 2D-3D data generated from interventional data, (3) 2D-3D interventional data.

Simulated 1D-2D synthetic data were produced to resemble an axial 2D slice of the aorta at the site of the renal arteries. This 2D slice is the simulated pre-operative image (see figure 4.10 (a)), which was then warped to generate intra-operative datasets as follows: TPS transformations were carried out using five fixed points (the corners of a square around the aorta to anchor the deformation and the ostium of the left renal artery) and a single moving point at the right renal ostium which was displaced by 10 mm at an angle of $i \times 45^\circ, i = 1, \dots, 8$ for each i^{th} deformation. From each of these warped images two 1D simulated intra-operative datasets were generated by a) projecting the ostia of the renal arteries onto a first

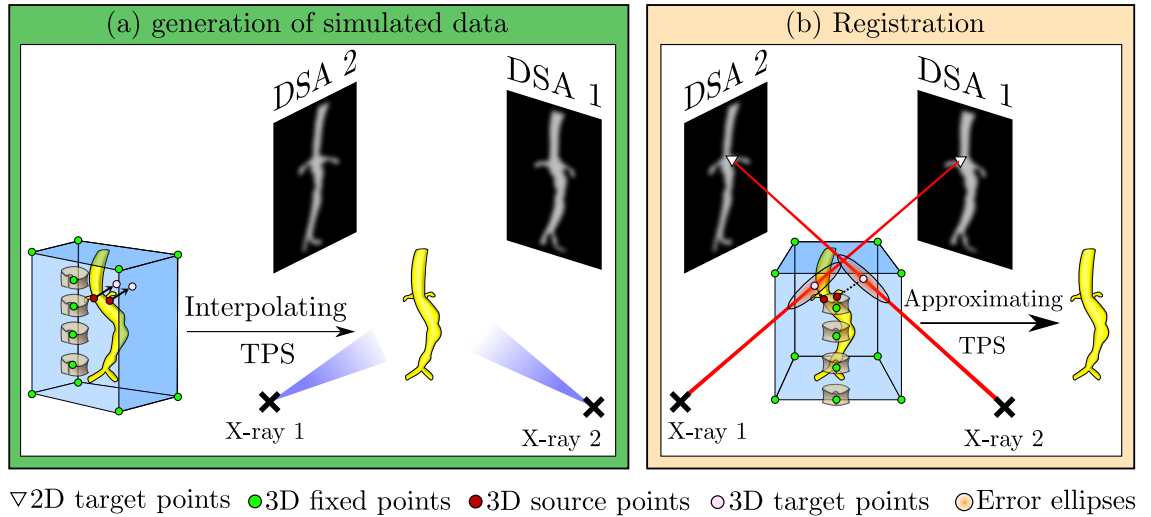


Figure 4.11: (a) generation of simulated data: an interpolating thin-plate spline is used to deform a pre-operative aorta to a simulated intra-operative aorta. Source, target and fixed points are picked as shown in the picture. The simulated 3D intra-operative aorta is projected onto two simulated DSA images with a different view direction. (b) the pre-operative aorta is registered to the simulated intra-operative DSA images using the thin-plate splines with projection uncertainties.

1D plane (rr and lr in figure 4.10 to represent a contrast-enhanced fluoroscopy image (b)) and then b) projecting the edges of the aorta onto a second, 30° rotated, 1D plane (ae and pe in figure 4.10 (b)) to represent an non-contrast-enhanced X-ray image with limited anatomical information on the outline of the vessels.

Simulated 2D-3D data generated from interventional data were produced using a patient CT as the pre-operative data (see figure 4.11). Intraoperative data sets were generated by warping the CT image, eight times, using a TPS transformation. Fixed points were the corners of a large bounding box surrounding the abdominal aorta, the centres of vertebrae L1 to L5 and the bifurcations of the iliac arteries. Moving points were placed on the renal artery ostia and anterior and posterior edges of the aorta level with the renal arteries. These points were displaced 10 mm , a distance chosen to be greater than the medium renal ostium motion (6.4 mm) reported in [55] and smaller than the maximum corresponding displacement (14.5 mm) reported in the same article. From each simulated ‘intra-operative’ CT scan, two intra-operative DSA images, with 30° difference in view direction, were synthesised (the reader is referred to 4.2.4 for an explanation of why this angle was chosen).

2D-3D interventional data were used from a complex EVAR procedure. The data consisted of a pre-operative CT scan, two intra-operative DSA images acquired after insertion of stent delivery device (views differed by 20° rotation).

4.3.2 Methods

For each experiment, the aorta was deformed with either a TPS+PU registration (figure 4.8) or an interpolating TPS sequential registration (interpolating TPS, figure 4.9). Each experiment used a set of fixed points (picked preoperatively) and moving points (picked from intra-operative images). The fixed points were as follows: for the 1D-2D synthetic experiment, corners of a square surrounding the aorta were fixed, while for the 2D-3D experiments the corners of a large cube surrounding the aorta were fixed. Those fixations were input to anchor the corners of a large area (volume) surrounding the aorta in the 1D-2D (2D-3D) experiments, as information outside the convex hull of control points can potentially be extrapolated poorly by the thin-plate splines [68]. In addition, for the 2D-3D experiments, the centre of lumbar vertebrae and bifurcation of common iliac were also fixed, as these points are expected to remain in a reasonable rigid body relationship with surrounding bones.

Moving points for 1D-2D synthetic data were picked as illustrated in figure 4.10 (c). As in the case with real DSA images different information is visible in different views. The renal ostia are assumed to be visible, and so were picked, in plane 1 (equivalent to an anterior-posterior DSA), whereas only the outline of the aorta is visible from the rotated view (which is treated as a non-contrast-enhanced X-ray image, see section 4.2.4), plane 2. 1D target points are selected on the 1D intra-operative planes and the corresponding points on the aorta are selected as 2D source points. For the renal positions corresponding points were the 2D renal artery ostia (see figure 4.10 (c)). For aortic edge positions corresponding points were the closest 2D edge points to back projection lines from the 1D point positions.

Moving points for the simulated 2D-3D data were picked at the ostia of the renal arteries and the Superior Mesenteric Artery (SMA) on the first DSA image and at the outline of the aorta on the second DSA image (the position of the ostia




Figure 4.12: Left: the pre-operative aorta surface is registered to two intra-operative DSA images. Right: 2D points are picked on the ostia of the visceral vessels in the first DSA and on angles of the edges of the aorta on the second DSA image. Corresponding 3D points are picked on the aorta surface and 2D-3D non-rigid registration is carried out.

cannot be retrieved from the second DSA image, since in general no contrast should be used to produce the second view during an actual intervention, see section 4.2.4). Figure 4.12 illustrates how the points were picked at the ostia and on the edges of the aorta, while figure 4.11 (b) gives an overall view of where all fixed and moving points are selected. The location of corresponding 3D source and 3D target points was found in a process equivalent to the one presented in figure 4.10 (c).

Moving points for the interventional data were picked at the ostia of the vessels when visible, generating 3D source and 3D target points as in the case of simulated 2D-3D data (see figure 4.12). The rigidly-registered aortic CT surface was then overlaid onto the first DSA image and the location of largest displacement with respect to the aortic outline was found using visual inspection. An additional moving point was picked at this location, non-rigid registration was performed using all previously picked points and the displayed aortic overlay was updated. This process was repeated, generating one additional moving point at each iteration until the aortic overlay closely matched the DSA image (i.e. within ≈ 2 mm). Moving points were selected in the second DSA image using the same procedure. Overall, 7 moving points were picked within approximately 2 minutes.




Figure 4.13: (a) 1D-2D simulated data: a small region of interest is selected around the right renal ostium and a larger region of interest is selected around the whole aorta. Validation is conducted computing Target Registration Errors on the regions of interest. (b) 2D-3D simulated data: two small regions of interest are selected around the right and the left renal ostium respectively and a larger region of interest is selected in a region that covers both ostia. Validation is conducted computing Target Registration Errors on the regions of interest. (c) 2D-3D interventional data: reprojection errors, i.e. the distances between a 3D point picked on the registered aorta surface and corresponding back-projection lines from corresponding 2D points picked on the intra-operative contrast-enhanced fluoroscopy images, are computed as means of validation.

4.3.3 Validation

For simulated data experiments Target Registration Errors (TRE) were calculated to the known ground truth using the following regions of interest: *ROI 1* was $26 \times 26 \text{ mm}^2$ or $30 \times 30 \times 30 \text{ mm}^3$ region centred on the aorta at the level of the renal arteries for the 1D-2D and 2D-3D data respectively. *ROI 2* was a $2 \times 2 \text{ mm}^2$ centred on the left renal artery ostium (for the 1D-2D data) or $2 \times 2 \times 2 \text{ mm}^3$ region centred on the right renal artery ostium (for the 2D-3D data). For the 2D-3D data *ROI 3* was a $2 \times 2 \times 2 \text{ mm}^3$ region centred on the left renal artery ostium.

No ground truth was available for the interventional data. Instead the position of the CT overlay was compared to a DSA image acquired later in the intervention. Reprojection error distances (i.e. the minimum distance a 3D point need to be moved to overlay its 2D counterpart, see [69]) were calculated at the left (*ROI 1*) and right (*ROI 2*) renal artery ostia.

All validation methods are summarised in figure 4.13.

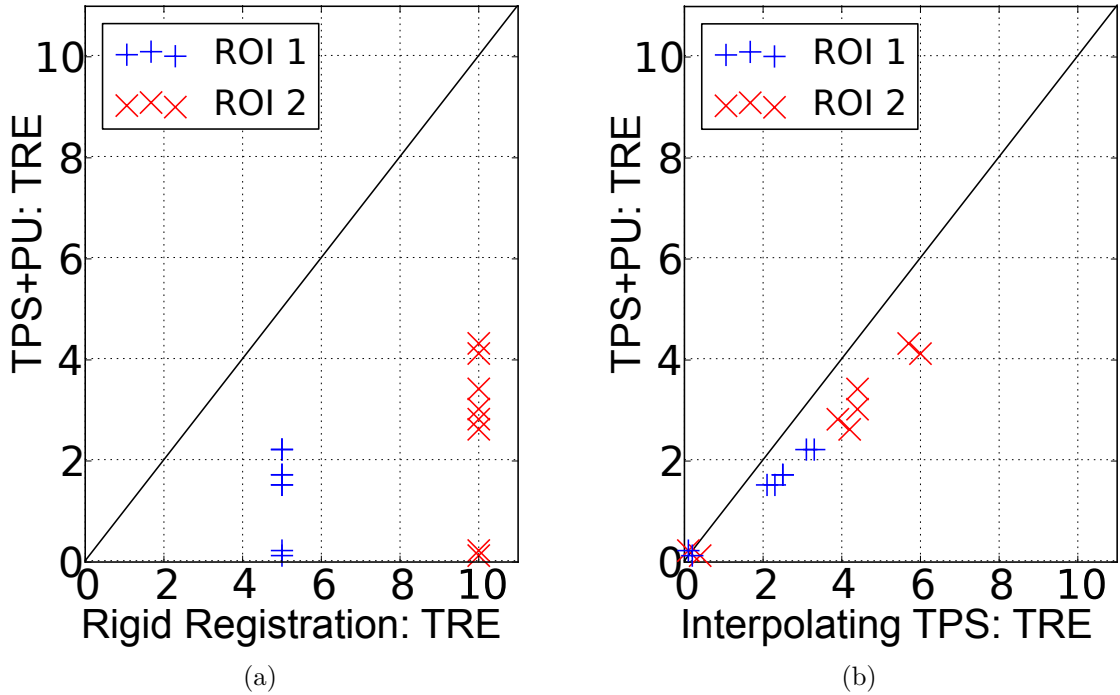


Figure 4.14: 1D-2D simulated data (all TRE values are expressed in mm). (a) TPS+PU vs. Rigid Registration TRE (b) TPS+PU vs. interpolating TPS TRE

4.4 Results

Results comparing TRE values on synthesised data for rigid registration, TPS+PU registration and interpolating TPS sequential warp are shown in table 4.1 and figure 4.14 for the 1D-2D case, and in table 4.2 and figure 4.14 for the 2D-3D case.

The TPS+PU registration results show a major improvement compared to rigid registration: registration accuracy is on average 3.8 times higher for the 1D-2D data and 3.5 times higher for the 2D-3D data. The TPS+PU registration results also show an improvement compared to the interpolating TPS sequential warp: registration accuracy is on average 1.4 times higher for the 1D-2D data and 1.8 times higher for the 2D-3D data.

Using results from the simulated 1D/2D data experiment, a paired t-test was carried out between the rigid registration target registration errors and the TPS+PU target registration errors which returned a p-value of 9.72×10^{-8} . For the same experiment, another paired t-test was carried out between the interpolating TPS registration target registration errors and the TPS+PU target registration errors

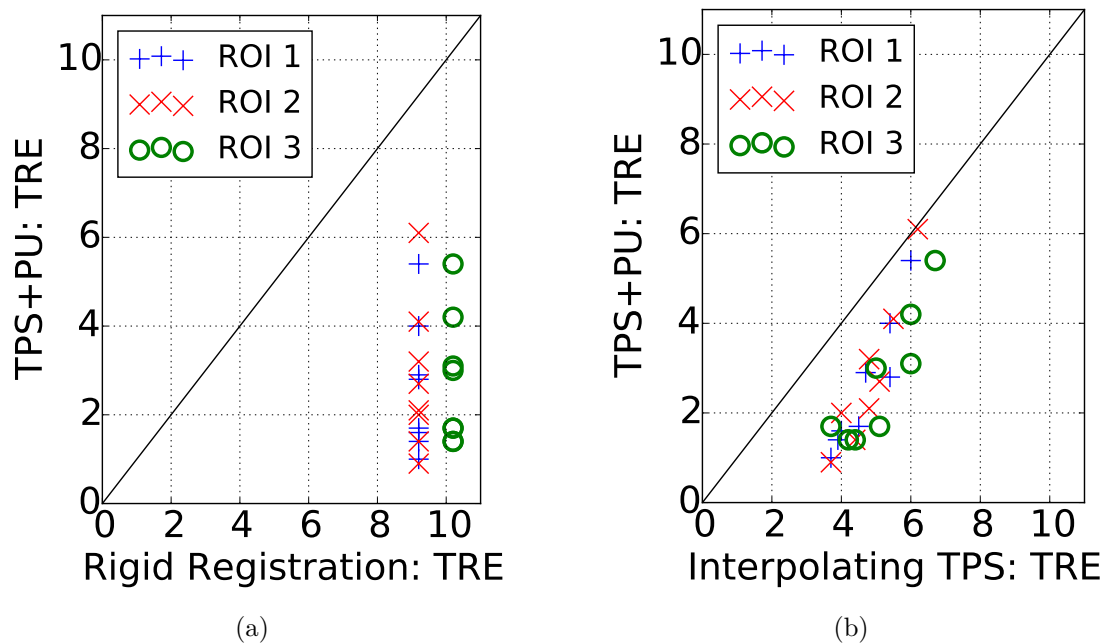


Figure 4.15: 2D-3D simulated data (all TRE values are expressed in mm). (a) TPS+PU vs. Rigid Registration TRE (b) TPS+PU vs. interpolating TPS TRE

ROI 1	set 1	set 2	set 3	set 4	set 5	set 6	set 7	set 8	mean(std)
Rigid registration	5.0	5.0	5.0	5.0	5.0	5.0	5.0	5.0	5.0(0.0)
Interpolating TPS	0.2	2.3	3.3	2.5	0.1	2.5	3.1	2.1	2.0(1.1)
TPS + PU	0.1	1.5	2.2	1.7	0.2	1.7	2.2	1.5	1.4(0.8)
ROI 2	set 1	set 2	set 3	set 4	set 5	set 6	set 7	set 8	mean(std)
Rigid registration	10.0	10.0	10.0	10.0	10.0	10.0	10.0	10.0	10.0(0.0)
Interpolating TPS	0.4	4.2	6.0	4.4	0.1	4.4	5.7	3.9	3.6(2.1)
TPS + PU	0.1	2.6	4.1	3.0	0.2	3.4	4.3	2.8	2.6(1.5)

Table 4.1: Simulated 1D/2D data target registration errors for 8 sets of deformations using rigid registration, interpolating TPS registration and TPS+PU registration

ROI 1	set 1	set 2	set 3	set 4	set 5	set 6	set 7	set 8	mean(std)
Rigid registration	9.2	9.2	9.2	9.2	9.2	9.2	9.2	9.2	9.2(0.0)
Interpolating TPS	5.4	6.0	3.9	3.7	4.0	4.5	5.4	4.7	4.7(0.8)
TPS + PU	4.0	5.4	1.4	1.0	1.6	1.7	2.8	2.9	2.6(1.4)
ROI 2	set 1	set 2	set 3	set 4	set 5	set 6	set 7	set 8	mean(std)
Rigid registration	9.2	9.2	9.2	9.2	9.2	9.2	9.2	9.2	9.2(0.0)
Interpolating TPS	5.5	6.2	3.7	4.4	4.0	4.8	5.1	4.8	4.8(0.8)
TPS + PU	4.1	6.1	0.9	1.4	2.0	2.1	2.7	3.2	2.8(1.6)
ROI 3	set 1	set 2	set 3	set 4	set 5	set 6	set 7	set 8	mean(std)
Rigid registration	10.2	10.2	10.2	10.2	10.2	10.2	10.2	10.2	10.2(0.0)
Interpolating TPS	6.0	6.7	4.2	3.7	4.4	5.1	6.0	5.0	5.1(1.0)
TPS + PU	4.2	5.4	1.4	1.7	1.4	1.7	3.1	3.0	2.7(1.3)

Table 4.2: Simulated 2D/3D data target registration errors for 8 sets of deformations using rigid registration, interpolating TPS registration and TPS+PU registration

Left Renal	expert 1	expert 2	expert 3	expert 4	mean(std)
Rigid registration	4.0	4.3	4.0	4.1	4.1(0.1)
Interpolating TPS	2.2	1.9	2.5	1.9	2.1(0.2)
TPS + PU	2.1	2.0	1.5	2.3	2.0(0.3)
Right Renal	expert 1	expert 2	expert 3	expert 4	mean(std)
Rigid registration	9.2	8.3	8.8	8.3	8.7(0.4)
Interpolating TPS	3.8	3.8	4.6	5.0	4.3(0.5)
TPS + PU	4.5	4.1	3.4	3.0	3.8(0.6)

Table 4.3: Interventional 2D/3D reprojection errors on the left and right renal ostia for points manually picked by 4 experts




Figure 4.16: (a) Comparison of CT overlay with DSA image before and after warp. Aorta outline from overlay (red) matches more accurately to DSA (cyan) after warp. (b) Registration accuracy at renal ostia using rigid registration, interpolating TPS sequential warp and TPS+PU registration.

which returned a p-value of 3.52×10^{-05} .

Using results from the simulated 2D/3D data experiment, a paired t-test was carried out between the rigid registration target registration errors and the TPS+PU target registration errors which returned a p-value of 1.00×10^{-16} . For the same experiment, another paired t-test was carried out between the interpolating TPS registration target registration errors and the TPS+PU target registration errors which returned a p-value of 3.41×10^{-12} .

Figure 4.16(a) visually shows the improvement in using the TPS+PU registration compared to rigid registration on clinical data. Numerical results are presented in table 4.3 and figure 4.16(b) where applying the TPS+PU registration halved the rigid registration error.

4.5 Discussion

Using a spline-based algorithm, several strategies can be adopted when registering a 3D volume with two images. First, the same landmark can be picked inside the two 2D views and its 3D position can be computed at the intersection of the back-projection lines. This approach is practically unfeasible, because of the differences in contrast and of projection effects, which in real clinical images makes it hard to find the same points inside two different views. A second method is to perform sequential one-view registrations: The 3D volume is registered with the first 2D

view, yielding an intermediate 3D volume, and this intermediate 3D volume is then registered with the second 2D view. Last, a TPS+PU registration approach can be used, where the information from the first plane and the information from the second plane are used at the same time. The results obtained in this chapter have shown that the registrations can be improved when using approximating thin-plate splines with anisotropic errors in a TPS+PU registration framework instead of interpolating thin-plate splines in a sequential warp framework.

Compared results of a sequential warp with interpolating TPS and of a TPS+PU registration show a clear improvement in accuracy when using TPS+PU registration. A recent study [12] reported a range of up to 11 *mm* (mean 4.5 *mm*) for aortic movement. Despite applying deformations of 10 *mm*, almost at the top of this reported range, the required registration accuracy better than 3 *mm* was achieved (as proposed in [12]) for half the TPS+PU registrations. No sequential warp registration achieved < 3 *mm* accuracy (see figure 4.15(b)).

The single-warp method presented in this chapter fits well within the standard clinical workflow. It requires knowledgeable manual input for point picking, but only a few points are needed, which are usually located in standard clinical images within a few minutes (compared to the procedure time of 4+ hours). Automated landmark identification would be preferred, however this is very difficult to achieve with guaranteed 100% robustness, and subsequent required checks on automated landmark selection are likely to require knowledgeable manual visual inspection anyway. As opposed to many registration methods (see the literature review of this thesis) the framework presented requires no optimisation strategy to obtain the deformation field that is applied to the aorta surface. Non-rigid registration can involve a high number of degrees of freedom, making the search of an optimal solution a lengthy process. In addition, the problems that are optimised are rarely convex, which means that the algorithms used to maximise a similarity value can return a local maximum instead of a global maximum. The approach presented is both time-efficient and robust to local extrema.

The experiments carried out in this chapter describe 2D-3D non-rigid registration using two DSA images, though one is only assumed to show information on the aortic outline. As mentioned in the clinical background chapter of this thesis, contrast agents can induce nephropathy and their introduction should therefore be

limited during the course of the surgical procedure. Bi-planar X-ray systems have been developed and can allow acquisition of two X-ray images at the same time while contrast is being injected inside the vessels of interest. The apparatus relying on simultaneous bi-plane fluoroscopy acquisition tend to add disruption to the clinical setting due to unwieldiness and sterility of the operating-room issues. For these reasons, the acquisition of two X-ray images from a different angle is generally restricted to one single DSA image and a subsequent X-ray image (without dye injection) with a different view direction. Different possibilities can be explored to obtain information on the aorta without the use of contrast and are depicted in figure 4.17. First, the 3D position of wires inside the aorta can easily be (manually or even automatically) retrieved using two X-ray images acquired at a similar point in time from different views. 3D source and target points can be manually picked on the aorta surface and in the vicinity of the 3D wire respectively to force the aorta surface to encompass the instruments visible in the fluoroscopy image. This requires making assumptions on the positions of target points, which can result in erroneous correspondences. Another approach that makes use of the 3D position of the instruments is to extract the continuous temporal positions in space of instruments from video sequences so as to map out areas of the patent lumen of the aorta. Lastly, some imaging methods can be used to enhance the fluoroscopy images and increase soft-tissue contrast. In particular, interventional digital tomosynthesis has been proposed (see [1]) to give information on the intra-operative geometry of blood vessels. Briefly, the method described in [1] uses a set of 2D intra-operative X-ray images (acquired at different positions in a ‘sweep’ in a limited angle range) that are rigidly 2D-3D registered to a pre-operative scan. A reconstruction surface is defined along the aorta centreline, and rays are back-projected from each view to the surface. For each point of the reconstruction surface, intersecting rays intensities are averaged and projected into a target image. The target image contains information specific to the aorta (i.e. soft-tissue and calcium deposits) and shows less clutter than conventional digital tomosynthesis. This approach is in its infancy and could be coupled to single-plane angiography to provide a second-plane enhanced fluoroscopy image at little cost to the patient.

The thin-plate splines plus projection uncertainties method presented here assumes that only one major non-rigid deformation (which occurs after inserting a




Figure 4.17: (a) wire information: points can be selected on the aorta surface and the X-ray image to force the surface to encompass intra-operative wires (left picture). (b) instrument motion: the position of the wires can be viewed within a video sequence. Using all the 2D positions, a rough outline of the aorta can be drawn on the fluoroscopy plane. (c) digital tomosynthesis: a ‘sweep’ of the X-ray system arm around the patient provides the surgeons with a series of low soft-tissue contrast X-ray images acquired at different angles. All these X-ray images can be back-projected on a user-defined reconstruction surface tailored to the geometry of the aorta. For each point of the reconstruction surface, an intensity I is computed as the average of each back-projection ray intensities (I_1 to I_n). This intensity I is projected on a target image. The resulting target image shows a ‘de-cluttered’ X-ray image with information (soft-tissue and calcium deposits) on the structures of the aorta delimited by the reconstruction-surface.

stiff wire) takes place during the procedure, therefore, the non-rigid registration algorithm would only need to be applied once per procedure. This assumption remains to be clinically verified.

4.6 Conclusion

This chapter describes a mathematical framework to enable accurate non-rigid registration of an aorta during minimally invasive surgery. The framework uses anisotropic error ellipsoids to loosely match 3D points manually-picked on an aorta surface to corresponding 2D points picked on X-ray images using approximating thin-plate splines. Error-ellipsoids are tailored in accordance with the X-ray projection directions to account for the lack of information perpendicular to the screen. Experiments on both simulated and interventional data show that this framework improves on current rigid registration strategies and on non-rigid thin-plate spline registration approaches (which do not take into consideration the 3D positional uncertainties resulting from the multi-dimensionality 2D-3D nature of the registration problem). A registration error of 3 *mm* has been proposed as a target for a registration system to assist image-guided endovascular abdominal aneurysm repair interventions [12]. The method proposed in this chapter is a step towards this target and its use inside the operation theatre can realistically be envisioned as it does not significantly disrupt the clinical workflow.

The work presented here assumes that a mathematical approach can be adjusted using clinical constraints (e.g. fixations of bony anatomy, deformations based upon intra-operative imaging) to produce physically-realistic displacement fields. Ensuring that displacement fields adhere to the laws of the physics is a hard task and cannot necessarily be done using the limited information present in theatre. The next chapter of this thesis investigates the use of bio-mechanical models to solve this problem and discusses how to use finite-element models in the context of image-guided surgery.

Chapter 5

Non-Rigid 2D/3D Registration using Finite Element Methods

Contents

5.1	Overview	102
5.2	Method	104
5.2.1	Registration Pipeline	104
5.2.2	Finite-Element Methods for Aorta Registration	110
5.3	Experiment 1: Using Physically Based Mechanical Parameters	114
5.3.1	Data	115
5.3.2	Method	116
5.3.3	Validation	117
5.3.4	Results	121
5.3.5	Discussion	122
5.4	Experiment 2: synthetic data	123
5.4.1	Data	123
5.4.2	Method	124
5.4.3	Results	124
5.4.4	Discussion	126
5.5	Experiment 3: Optimal Parameter Values	127

5.5.1	Data	127
5.5.2	Method	127
5.5.3	Validation	128
5.5.4	Results	128
5.5.5	Discussion	131
5.6	Overall Discussion	133
5.7	Conclusion	137

5.1 Overview

The previous chapter described a method to warp the aorta using a global spline-based interpolation function. This approach used a mathematical model to calculate a deformation field from points that had been manually selected on a 2D intra-operative fluoroscopy plane and to match these 2D points with 3D points picked on a pre-operative aorta surface. The physical realism of the deformations was tailored by a combination of two factors: 1) deformations were chosen to produce limited localised bending of the aorta (through a minimisation of a bending energy function) and 2) anatomical constraints were added, based on the positions of rigid structures close to the aorta (e.g. the vertebrae).

This chapter proposes to take into account the biomechanics of soft-tissues to improve the physical realism of the deformations. It investigates a very different method that is based upon the intra-operative use of finite-element simulations. As explained in the clinical background of this thesis, limited information on the patient’s anatomy is available inside the operation theatre: the soft-tissues cannot be seen at all times (the injection of contrast agents is kept to a minimum owing to their nephrotoxicity), therefore surgeons mainly have visual access to the bony anatomy of the patient and the instruments present inside them. While the surgery is being performed, the main source of aortic deformation is the insertion of those instruments inside the patient’s vessels. In particular, the introduction of a stiff wire prior to stent-graft deployment leads to the largest displacements of the aorta, and it is this deformation which the approach discussed in this chapter intends to model.




Figure 5.1: 1) Wire initialisation: a wire is initialised at the centreline of the pre-operative aorta 2) Intra-operative wire generation: the stiff guide-wire is visible inside two X-ray images acquired from different angles at similar points in time after its insertion. The 3D position of the wire is computed from the two views using triangulation techniques. 3) Wire displacement: the wire is moved from its initial position to its intra-operative position 4) Aorta deformation: the pre-operative aorta is deformed to its intra-operative shape so as to encompass the intra-operative wire. Steps 3) and 4) are performed in a step-wise iterative fashion.

Fully modelling the introduction of a stiff wire inside the aorta would not be feasible, because there are constraints inherent to the surgical procedure (overall computation times should not exceed a certain duration, e.g. 10 minutes, so as to avoid disruption of the clinical workflow) and the forces that are exerted by the surgeons cannot be easily retrieved. Instead, the method followed in this chapter consists of sequentially deforming the surface of an aorta to encompass a wire that is being moved from an idealised position (at the centrelines of the aorta surface extracted from the pre-operative CT scan) to its positions after it has been inserted during the intervention. This is illustrated in figure 5.1. The position of the intra-operative wire is the only intra-operative information required as an input of the method described above and this process could be automated as per automated methods for instrument segmentation provided in the literature (see [4] for an automated reconstruction method).

This chapter is divided as follows: the first section of this chapter describes the methods that have been devised to extract information from the data acquired before and during the EVAR surgery and the ways the finite element simulations were carried out. Next, a first set of experiments on four patient datasets is presented, where mechanical parameters obtained from the literature were input into




Figure 5.2: Input data: 3D pre-operative CT scan and two 2D intra-operative X-ray images. 1) A volume of the aorta is segmented from the CT scan. 2) A wire is initialised at the centrelines of the aorta. 3) The intra-operative position of the wire is extracted from the X-ray images. 4) A deformation field from the initial wire to its intra-operative position is derived. 5) The volume of the aorta is deformed in accordance with the deformation field.

the algorithm. Failure of the registration led to an investigation of the role the main parameters play using synthetic data. This is presented in section 5.4. The last section describes how information derived from the experiments on synthetic data was applied to real clinical data. Optimal parameters were computed on one patient dataset and used on three other patients datasets.

5.2 Method

This section describes both the registration pipeline that was introduced in the previous section and the use of finite-element methods to deform an aorta so that it encompasses a moving wire.

5.2.1 Registration Pipeline

The input data, each step of the registration pipeline, the output data are detailed below and an overview is provided in figure 5.2.

Input Data

pre-operative 3D CT scan A pre-operative 3D contrast-enhanced CT scan is used, which gives information on the geometry of the patent lumen of the aorta prior to the intervention.

intra-operative 2D X-ray images Two intra-operative 2D X-ray images acquired from different angles are used, which give information on the bony anatomy of the patient, and the position of instruments during the intervention. The two X-ray images cover the abdominal region and are acquired at similar points in time, after the insertion of a stiff wire and prior to the deployment of a stent-graft. Accurate alignment of the 2D images with the 3D pre-operative CT scan is obtained using the 2D/3D rigid registration method detailed in [69].

Pipeline Steps

volume extraction Following segmentation from a pre-operative CT scan, the aorta surface of a patient is meshed into tetrahedral volumetric elements. The aorta surface is used to position a wire ('initial wire') inside the pre-operative aorta in a physically realistic way (this step is referred to as 'wire initialisation' and described in the next paragraph), while the aorta mesh is to be used at a later stage to carry out the finite-element simulations.

wire initialisation A virtual wire is initialised within the pre-operative aorta. The aim is to position the virtual wire in as realistic a position as possible (for example, low curvature configurations should be preferred) where it is completely contained within the pre-operative aorta. An overview of this process using the centrelines of the pre-operative aorta is shown in figure 5.3. The initialisation is done as follows:

- the aorta is skeletonised by means of sequential morphological erosion. This step produces a thin shape that approximately represents the centrelines of the iliac arteries, of the aorta main segment, and of the surrounding branches.
- The user manually picks a set of points on the skeleton, excluding points that are located away from the aorta centrelines. While full automation of the




Figure 5.3: 1) A skeleton is extracted from the aorta volume 2) Points are manually picked on the skeleton at the centrelines of the aorta 3) A polynomial that is encompassed by the aorta is fitted to the points

point selection process is feasible, it requires additional pre-steps: the parts of the skeleton that correspond to the visceral arteries and to one of the iliac arteries (the iliac artery into which the guide-wire is not introduced during the intervention) must be removed prior to automatically selecting points. These pre-steps are harder to automate, which is the reason why a manual selection is preferred.

- An x, y, z set of least square polynomials is fitted to the points (along the X, Y and Z axes respectively). The order (starting from 0) of these polynomials is increased up until the moment the aorta encompasses the polynomial curve, along its entire length.

This wire represents the guide-wire used during aneurysm repair procedures, and is initialised inside the pre-operative aorta. This wire is to be dragged towards its intra-operative position.

intra-operative wire generation After having positioned the X-ray images in the same coordinate system, the 3D position of the guide-wire during the aneurysm repair procedure is computed from its 2D positions obtained from the two intra-operative X-ray images, as illustrated in Fig. 5.4. Putting the two X-ray views in the same coordinate system is done using the method presented in [69]. Briefly, [69] proposes to rigidly register a 2D X-ray image acquired during EVAR surgery to a CT scan obtained prior to the intervention. This is done by varying the position of the CT volume with respect to the X-ray image, projecting the CT volume onto DRRs and comparing the DRRs with the X-ray images using a gradient difference similarity measure. Six rigid body parameters (three translation X, Y, Z plus three




Figure 5.4: 1) Corresponding 2D points are picked on the wire inside the two X-ray images. Corresponding 3D point positions are computed using the 2D positions 2) The 3D points are joined with a smooth curve.

rotations $\theta_x, \theta_y, \theta_z$) are searched for using a coarse global optimisation followed by a hill climbing search. Using this algorithm, each of the two X-ray views are put in the CT coordinate system.

The computation of the 3D position of the guide-wire after it has been deployed is done in three steps. Firstly, an expert user picks a set of 2D points on the wire inside the first X-ray image. Because of the high contrast between the stiff wire and the surrounding structures in X-ray images, manual selection of points on the wire can be done reasonably precisely (e.g. with an error below 1 *mm*, as measured from the observable diameter of a renal ostium and of a guide wire visible in an X-ray image used to conduct the registration experiments in this chapter). For each of these points, the user selects the intersection between the wire in the second X-ray image and the epipolar line (i.e. the line from the first camera to the point selected on the first X-ray image projected on the second X-ray image). This yields a second set of corresponding 2D points on the wire inside the second X-ray image. Secondly, a 3D position is retrieved for each pair of corresponding 2D points. Errors in positioning of the 3D points are influenced by the errors in retrieving the 2D points and are therefore low (below 1 *mm*, which is lower than the clinically acceptable registration error of 3 *mm* mentioned in [12]). Thirdly, a x,y,z set of cubic splines is fitted to the 3D points along the X, Y and Z axes respectively. This yields a smooth curve that interpolates the 3D points.

computation of a deformation field A deformation field is computed that goes from a starting position (i.e. initialisation of the guide-wire), to an ending position (i.e. intra-operative position of the guide wire). Two approaches were followed.

The first approach is described below and non-rigidly registers the ‘initial wire’ to its intra-operative counterpart, using the non-rigid ICP surface registration method proposed by Amberg et al. in [2]. In this article, Amberg et al. sequentially registered a template surface to a target surface by: 1) Deforming the template surface using locally affine transformations for each sequence of the registration process. 2) Constraining the set of locally affine transformations using a stiffness parameter (higher stiffness parameter values result in smaller variations between each local transformation being accepted). The stiffness parameter was progressively decreased. The deformation fields obtained using this method were found to be unrealistic for this application, as the length of the wire was changed during the registration process: something which would not physically occur.

A second strategy was subsequently devised to produce more realistic deformation fields: this involved preserving the length of the wire, as shown in figure 5.5. The algorithm is described below:

- A series of equidistant points is selected on the intra-operative wire.
- Different series of points are automatically computed on the ‘initial’ wire, where the distances between consecutive points are equal to the distances between corresponding consecutive points on the intra-operative wire. Each series of starting wire points, together with corresponding intra-operative wire points, form a displacement field.
- For each displacement field (u, v or w in figure 5.5), individual distances between corresponding points on the intra-operative (I_i) and starting wires (P_i^u , P_i^v or P_i^w) are computed. Rather than calculating Euclidean distances between each point I_i and P_i^w , each vector (\vec{u} , \vec{v} or \vec{w}) is decomposed into a vector normal to the starting wire at the position (P_i^u , P_i^v or P_i^w), \vec{n} and a vector tangential to the wire, \vec{t} : $\vec{u} = \vec{n} + \vec{t}$. The length of the tangential vector is used as a tangential distance between I_i and (P_i^u , P_i^v or P_i^w). These individual tangential distances are summed to give a displacement field distance.




Figure 5.5: A series of equidistant points (I_1, I_2, I_3 , etc.) is generated on the intra-operative wire. The starting point I_1 is matched to different starting points (P_1^u, P_1^v or P_1^w) on the starting wire. Subsequent points (P_2, P_3 , etc.) are computed so that the length of the wire is preserved. For all points I_i , a distance ($d(u_i)$) to its counterpart is computed, and a deformation field distance (d_u, d_v, d_w , etc.) is computed as the sum of all ($d(u_i)$). The displacement field with the minimum distance is kept as the output displacement field.

- The displacement field with the minimum overall displacement field distance is kept as the output displacement field. The distance measure defined above is built to promote displacement fields with low up-and-down motion of the wire (which are less likely to cause larger bending of the aorta surface).

deformation of the aorta The guide-wire is moved from its starting (pre-operative) position to its ending (intra-operative) position. An aorta-wire collision detection algorithm is implemented, as depicted in figure 5.6. Initially the guide-wire is located inside the aorta. This guide-wire is moved towards its intra-operative position in a number of steps. At each step of the guide-wire motion, collisions between the guide-wire and the aorta are checked: if the wire is still inside the aorta, the aorta is left unaltered, if the wire is outside the aorta, the aorta is deformed so as to encompass the wire. Deformations of the aorta are performed using finite-element methods and are described in section 5.2.2.

Some finite-element software is equipped with collision detection capabilities, inclusive of the package used for the experiments presented in this chapter. The reasons why these capabilities were not exploited are detailed in section 5.2.2.




Figure 5.6: Start: A wire is initialised at the centreline of the initial (pre-operative) aorta. Step i-1: The wire is moved towards its intra-operative position. The aorta remains unchanged while the wire is still inside it. Step i: The wire is moved towards its intra-operative position. The wire is now outside the aorta. A collision between the aorta and the wire is detected. The aorta is deformed so as to encompass the wire. End: The aorta has been successively deformed to encompass the intra-operative wire.

Output Data

A deformed surface of the aorta is obtained, that encompasses the 3D guide-wire position retrieved from intra-operative X-ray images.

5.2.2 Finite-Element Methods for Aorta Registration

Finite-element software

Finite element simulations were done using an open-source nonlinear dynamic finite element solver, NiftySim (as described in [94]), version 2.3. NiftySim allows GPU execution, which makes it a sensible choice for use during an AAA repair procedure, where computational time cannot exceed a certain duration (e.g. 10 minutes).

NiftySim takes as input a tetrahedral mesh, boundary conditions (fixed nodes, displaced nodes and displacements/forces), material parameters (Young’s modulus, Poisson’s ratio, density, damping coefficient), and simulation parameters (simulation time, time step).

NiftySim allows specification of different material formulations. To model the hyperelastic behaviour of the aorta, a neo-Hookean formulation was used. Neo-Hookean models predict the non-linear stress-strain behaviour of materials under-

going large deformations, and can be expressed as:

$$\Psi = \frac{\mu}{2} (\bar{I}_1 - 3) + \frac{\kappa}{2} (J - 1)^2, \quad (5.1)$$

where $J = \det F$ (with F the deformation gradient), \bar{I}_1 is the first principal invariant of the modified right Cauchy-Green deformation tensor $\bar{\mathbf{C}}$ (with $\bar{\mathbf{C}} = J^{-\frac{2}{3}} \mathbf{C}$ and \mathbf{C} is the right Cauchy-Green deformation tensor), μ is the shear modulus and κ is the bulk modulus [46]. The experiments presented in this chapter intend to retrieve the deformation an aorta surface undergoes during an EVAR procedure but do not aim to do so by modelling the aorta wall and the blood flowing through the artery as precisely as possible. Such a physics-based simulation would require extensive knowledge of patient-specific parameters (which is hard to garner) and high simulation times, which would likely disrupt the clinician's work-flow. Instead, the method proposed here models the aorta as a solid structure made of a neo-Hookean material. Previous neo-Hookean representations have been successfully employed for the modelling of tissues of a different nature, e.g. breast tissues ([37]) or the prostate gland ([41]).

NiftySim allows adding a membrane. A membrane was added to the models of the aorta, as this was found to increase the stability of the finite-element simulations carried out for this chapter.

Meshing of the aorta

The aorta surface is extracted from the pre-operative scan, and a 3D tetrahedral mesh is produced from the 2D surface. Care is taken to minimise the total number of nodes of the mesh and to produce the largest tetrahedra that, taken as a whole, reproduce the aorta geometry with reasonable fidelity. This is done for two reasons: 1) so that forces are only computed on a few nodes during the finite-element simulation 2) the smallest element length (i.e. the length of the smallest tetrahedron edge where all tetrahedra of the mesh are considered) is increased, therefore the critical time step Δt_{cr} is increased, where Δt_{cr} is the limit above which all finite-element simulations are bound to break down (see [94]). This results in: 1) reduced computation times 2) increased stability of the simulations.

Instead of modelling the aorta as a hollow cylinder-like structure, the methods presented in this chapter represent it as a solid structure. This is because the experiments presented here do not aim at conducting finite-element simulations on the aorta in the most physically accurate way. This would be hard to perform inside an operation theatre (due to the high computation times such simulations would entail) and would also require inputting some patient-specific material parameters, the values of which are hard to clinically retrieve. The way the aorta was modelled for the experiments presented below allows an accurate-enough representation of the behaviour of a cylinder-shaped structure through which a high-pressure fluid is flowing, while allowing fast finite-element simulations to be performed on it.

Boundary conditions

At each step when a collision is detected as described above, displacements are computed between the current position of the wire and the position of the wire at the following step. These displacements are then applied to the aorta mesh, in order to get a new surface of the aorta that encompasses the wire. The procedure to apply the displacements to the mesh is as follows: for each displacement, the intersection with the aorta mesh is computed, and displacement contributions of the same magnitude and orientation are assigned to the mesh nodes that surround the intersection. A weight that is inversely proportional to the distance between the intersection point and each node is assigned to each displacement contribution. For each node, a weighted average of the displacement contributions is computed, using previously calculated weights. This step yields the displacements (boundary conditions) that are input into the finite-element simulations.

The computation of displacement contributions involves a selection of surrounding nodes within a certain radius of the intersection points. When using a small radius, fewer surrounding nodes are selected, and therefore the output mesh is less likely to encompass the intermediate wire. On the other hand, when a larger radius is used, more surrounding nodes are selected. This results in more input boundary conditions and thus the simulation is more computationally expensive. In addition (for high radii), displacements of large zones of the mesh are assigned, instead of being (physically) governed by the finite-element software, and the resulting solu-

tion is not likely to be as smooth as one with fewer manual constraints (this can result in artefacts, especially ridges, around the intra-operative wire).

Simulation convergence

The finite-element simulations described in this chapter use explicit time integration operators and are conditionally stable: the time step Δt must be smaller than a critical limit Δt_{cr} (as described above, see [94] for further explanations). Δt_{cr} is dependent on material parameters: the higher the Young's modulus and the higher the Poisson's ratio, the smaller the critical time step. This prevents the use of Young's Moduli higher than a certain threshold (e.g. $10^{11} Pa$) while carrying out simulations on a mesh which accurately represents the aorta surface.

Throughout the course of the experiments carried out for this chapter, it was observed that NiftySim simulations became unstable when large displacements were input into the finite-element solver: inputting large displacements would require very small time steps which would cause long simulation times impractical for clinical applications (e.g. more than 12 hours). Several strategies have been used to increase the stability of the simulations with respect to that issue: 1) using smaller simulation time steps 2) subdividing the displacements into several sets of displacements and apply one simulation per set of displacements.

Additional details for the second strategy are provided below. When a simulation \mathcal{S} involving a set of displacements (d) applied to a mesh M fails to converge, this set of displacements is broken down into a series of n sets of displacements (d_1, d_2, \dots, d_n where $d = d_1 + d_2 + \dots + d_n$). Starting from the initial mesh $M_1 = M$, each set of displacements d_i is applied to the mesh M_i to produce a mesh M_{i+1} . The end mesh M_{n+1} is then used as the result of the simulation \mathcal{S} . If one of the intermediate simulations fails to converge, the process of subdividing the set of displacements is applied recursively to the intermediate set of displacements corresponding to the failed simulation.

Contact modelling

NiftySim can simulate contact between soft-tissue and geometries frequently encountered in medical settings, by predicting the deformation of a master surface

(the aorta surface in the context of this thesis), searching for contacts between the master surface and a slave surface (the wire surface in the context of this thesis) and altering the master surface deformation based upon the detected contacts [46]. The method proposed in this chapter describes the displacement of a stiff wire inside the aorta, where the aorta is represented as a solid structure (as opposed to a hollow cylinder-like shape inside which a wire is being moved). The method described in [46] cannot be applied to this particular geometry, where the whole wire is already in contact with the aorta geometry. For this reason, a method to drag the wire from its starting position to its ending position and to deform the aorta alongside has been implemented (and is described in section 5.2.1).

Aorta deformation

When the finite-element simulations have been performed, a deformed version of the input aorta tetrahedral mesh is obtained. This deformed aorta mesh is used to deform the pre-operative aorta surface extracted from the input CT scan. After running the pipeline described in section 5.2.1 a representation of the intra-operative aorta surface is computed in several steps, as illustrated with figure 5.7: first, the pre-operative aorta is meshed using a tetrahedral finite-elements mesher. Second, the mesh is deformed towards its intra-operative shape using an aorta-wire collision detection algorithm, as shown in figure 5.6. Finally, to get a picture of the entire deformed aorta surface (as opposed to the simplified intra-operative mesh), a thin-plate spline deformation is applied to the pre-operative aorta, using the nodes of the initial mesh as source points and the corresponding nodes of the target mesh as target points. This last step yields a surface of the aorta as it would be presented to clinicians during the operation.

5.3 Experiment 1: Using Physically Based Mechanical Parameters

An initial experiment was conducted using physically-based mechanical parameters, which were obtained from the literature, to deform the aorta following the method proposed in this chapter.

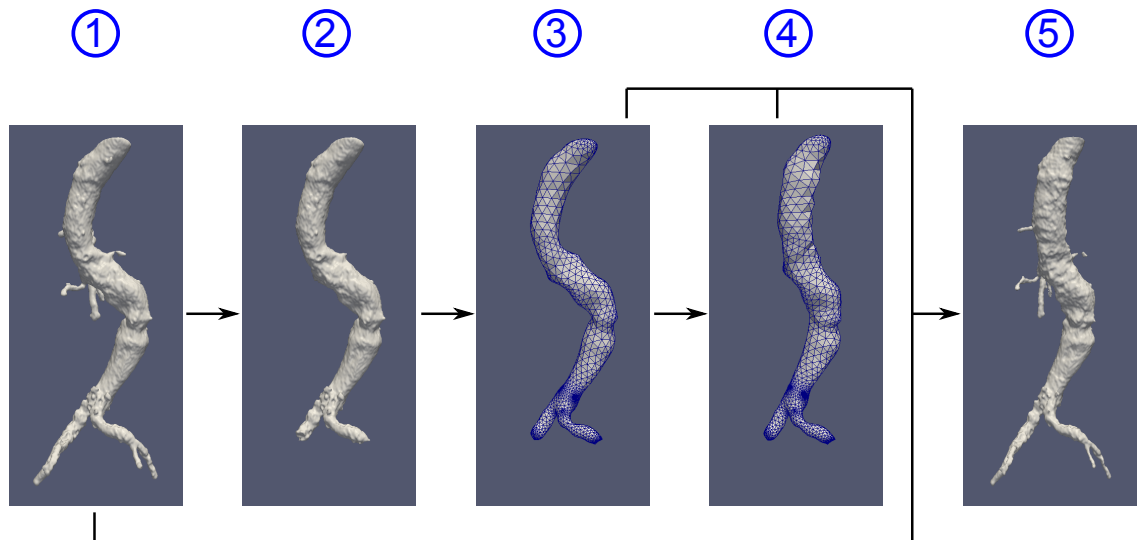


Figure 5.7: 1) The pre-operative aorta is segmented from the CT volume. 2) The pre-operative segmentation is ‘cleaned-up’: branches are removed and the top and bottom parts are cropped. 3) The pre-operative aorta is meshed. 4) The pre-operative aorta mesh is deformed using finite-element methods. 5) A displacement field is computed using node correspondences from the pre-operative aorta mesh (source nodes) to the deformed mesh (target nodes). The calculated displacement field is applied to the pre-operative aorta. An ‘intra-operative’ aorta is computed using a thin-plate spline deformation with the source and target points mentioned previously.

5.3.1 Data

Data from four patients undergoing EVAR were used. For each patient this was comprised of a pre-operative CT scan and two intra-operative X-ray images acquired at different views, shortly after the stiff guide-wire was inserted.

Narrow branches can cause 3D finite-element meshers to produce small tetrahedral elements (in comparison with those used to model the main segment of the aorta) and these small elements can make finite-element simulations less stable, for reasons reported in [94]: for the simulations to be stable, the critical time step must be smaller than a critical limit, which decreases as the smallest characteristic element length in the assembly decreases. The experiments carried out in this chapter aim at modelling the deformation that the main segment of the aorta undergoes during EVAR procedures and are particularly concerned with the position of the ostia of the visceral vessels. They do not intend to provide information on the branches of the aorta after a stiff-wire has been inserted into it. Because of the above, the 3D surface of the aorta was cleaned-up for all four patients: branches

were manually removed and the bottom and top sections were manually cropped. A surface mesh was extracted using ITKsnap version 2.4.0 (see [100] for a full description of this segmentation tool) with a semi-automated active contours (‘snake’) segmentation. From this surface mesh, a 3D volumetric mesh was generated using NETGEN version 5.1 (as described in [87]), an automatic 3D tetrahedral mesh generator. Parameters that control the granularity of the mesh were manually set to obtain as small a number of (large) tetrahedral elements as possible. The 3D volumetric mesh was input into the finite-element simulations.

5.3.2 Method

The pipeline described in section 5.2.1 was followed:

1. a ‘virtual’ wire was initialised at the centreline of the aorta
2. each of the 2D X-ray images was rigidly registered to the 3D CT-scan
3. the intra-operative position of the stiff wire was calculated from the two X-ray images
4. a displacement field was computed from the initial to the intra-operative position of the wire
5. using finite-element methods, the aorta was gradually warped to its position during the intervention as the initial wire was moved towards the intra-operative wire

The main source of deformation during EVAR procedures is the introduction of a stiff wire into the aorta. Other instruments inserted into the lumen of the aorta can cause limited additional motion of the aorta wall. The contrast-enhanced X-ray images used to validate the experiments presented in this chapter were acquired after the insertion of an undeployed stent graft (which was pulled-up after the introduction of the stiff guide-wire). This undeployed stent-graft is likely responsible for some additional deformation of the aorta surface, which is why a two-step deformation process was followed, as described in figure 5.8:

1. a wire (stiff guide-wire) was displaced from a starting position to an intra-operative position, and the aorta moved so that the wire remains inside the aorta.
2. when positioned at its intra-operative position, the wire was expanded so as to match the shape of an undeployed stent-graft, which has a much larger diameter than a stiff wire. To expand the wire, a displacement field was generated from the intra-operative wire towards the undeployed stent-graft as described below. A set of (uniformly distributed) points were selected on the intra-operative wire. For each $p_{intraop}$ of these points a plane \mathcal{P} normal to the intra-operative wire at the position $p_{intraop}$ was computed. The position p_{preop} on the initial wire corresponding to $p_{intraop}$ was retrieved. The ray originating at p_{preop} going towards $p_{intraop}$ was projected on \mathcal{P} and the intersection of \mathcal{P} with this projected ray returned a point p_{stent} . The displacement $\overrightarrow{p_{intraop}p_{stent}}$ was assigned to the point $p_{intraop}$ for the displacement field from intra-operative wire to undeployed stent-graft.

This stage returned the shape of the aorta right after an undeployed stent-graft has been inserted into it.

Finite-element simulations were carried out using the finite-element software described in 5.2.2. Mechanical parameters for the simulation were chosen according to the work done in [31]: the Young's modulus of the aorta was set to 4.66 MPa , the material density to 1000 kg/m^3 (close to the literature value of 1062 kg/m^3) while the Poisson's ratio was set to 0.45 to simulate the behaviour of nearly incompressible tissue.

5.3.3 Validation

Validation was performed in two ways:

- Two experts (expert 1, expert 2) picked the left and right renal ostia inside the pre-operative CT scan and in the DSA image. Reprojection distance errors, as described in [69], were computed for all the expert picked points.
- Expert 1 and expert 2 picked a few points on the edges of the aorta, at locations where contrast was sufficient to determine the edge of the aortic




Figure 5.8: 1) Top: the wire is displaced from a starting position to its intra-operative position. The pre-operative aorta is deformed accordingly. Bottom: the displacement field from initial wire to intra-operative wire minimises a distance criterion and preserves the length of the wire 2) Top: the intra-operative wire is expanded to match the undeployed stent-graft. The deformed aorta obtained at stage 1 is deformed accordingly. Bottom: displacements go from the intra-operative wire to the outer part of the stent-graft and are orthogonal to the intra-operative wire.




Figure 5.9: Calculation of point-to-surface distance errors 1) A distance map is computed around the deformed aorta 2) A point is picked on the intra-operative DSA image, on a visible edge of the aorta. That point is back-projected into a line in the 3D space. 3) For each point of the back-projected line, its distance to the aorta surface is computed using the distance map. The absolute value of the minimum distance (represented by 1 in this figure) is selected as the point-to-surface distance registration error.

lumen. An illustration of errors distances calculated from these edge points is given in figure 5.9. A distance map is computed from the deformed aorta volume: the further the points are from the walls, the larger their absolute distances are. Points inside the aorta are assigned a negative value and points outside the aorta have a positive value. For each point, a ray is back-projected to the X-ray source and distances on this ray are computed from the distance map. The point-to-surface distance error is defined as the absolute value of the smallest of the distances computed on the ray. For each patient, the mean and standard deviation of all the edge points were reported as an error criterion.

The points used for the computation of point-to-surface distance errors are shown in figure 5.10. Insufficient contrast propagation makes accurate selection of points on the edges of the aorta harder: some parts of the aorta surface (especially outside the direct vicinity of the renal arteries) are not as clearly delineated as others. As an example, only a small number of points were picked below the renal arteries of patient 4's aorta. A more conservative approach was taken by expert 1, who only picked points in areas with high visibility.




Figure 5.10: Validation: manual selection of edge points. Points were manually picked by two expert users on the visible edges of the aorta. Insufficient contrast propagation means some parts of the aorta could not be well delineated.

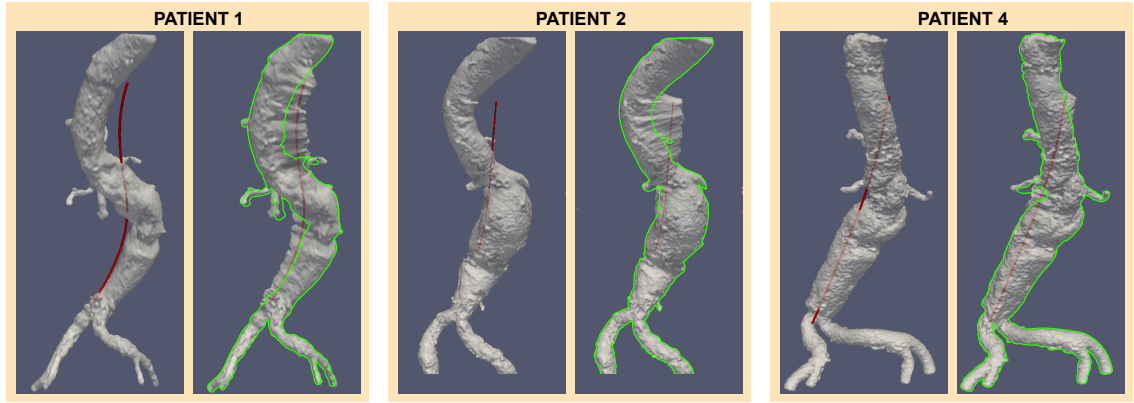


Figure 5.11: Deformation of the aorta for patient 1, patient 2 and patient 4. For each patient, the left picture shows the pre-operative aorta surface with the red intra-operative wire and the right picture shows the deformation of the aorta to encompass the wire, with the green outline representing the surface of the pre-operative aorta.

left renal	rigid registration		non-rigid registration	
	expert 1	expert 2	expert 1	expert 2
patient 1	8.2	6.2	4.7	4.8
patient 2	5.6	2.8	5.6	2.8
patient 4	9.6	7.5	9.6	7.5

Table 5.1: Left renal ostium: reprojection errors (in mm) computed on three patients (patient 1, patient 2 and patient 4) for a rigid registration and for the non-rigid registration using the finite-element method proposed in this chapter. Both rigid and non-rigid registration errors are given for two different experts. Non-rigid registration did not improve the reprojection errors on patient 1 and 4, indicating that for both these patients the left renal ostium was not displaced by the non-rigid registration.

5.3.4 Results

Figure 5.11 shows how each patient's aorta was deformed on three of the patient datasets (patient 1, patient 2 and patient 4). For all patients, most of the aorta tissue stayed in place, while only a small section was pulled to encompass the wire. The finite-element simulations would not complete for patient 3 as the portion of the aorta that was moved was too narrow to encompass the wire at a certain stage.

Reprojection errors for the left and right renal ostia are presented in table 5.1 and table 5.2 respectively.

Reprojection errors were not changed by the use of the non-rigid registration algorithm on patient 1 for the right renal ostium, on patient 2 for the left renal ostium and on patient 4 for the left renal ostium. This is due to the fact that these

right renal	rigid registration		non-rigid registration	
	expert 1	expert 2	expert 1	expert 2
patient 1	6.7	5.0	6.7	5.0
patient 2	3.4	4.6	2.6	4.4
patient 4	8.8	12.0	8.6	12.0

Table 5.2: Right renal ostium: reprojection errors (in mm) computed on three patients (patient 1, patient 2 and patient 4) for a rigid registration and for the non-rigid registration using the finite-element method proposed in this chapter. Both rigid and non-rigid registration errors are given for two different experts. Non-rigid registration did not improve the reprojection errors on patient 2, indicating that for this patient the right renal ostium was not displaced by the non-rigid registration.

	rigid registration		non-rigid registration	
	expert 1	expert 2	expert 1	expert 2
	mean (std)	mean (std)	mean (std)	mean (std)
patient 1	5.6 (3.0)	6.2 (3.3)	3.1 (2.7)	3.4 (3.3)
patient 2	2.2 (0.7)	2.1 (1.4)	1.6 (0.6)	1.9 (1.3)
patient 4	2.4 (1.8)	2.3 (1.8)	1.9 (1.6)	1.9 (1.7)

Table 5.3: Mean and std point-to-surface distance errors (in mm) computed on three patients (patient 1, patient 2 and patient 4) for a rigid registration and for the non-rigid registration using the finite-element method proposed in this chapter. Both rigid and non-rigid registration errors are given for two different experts.

portions of the aorta have not been displaced by the finite-element simulations. Other reprojection errors generally show some improvement when using non-rigid registration rather than rigid registration, with errors coming closer to the clinical target of 3 *mm* defined in [12].

Mean point-to-surface distance errors and their standard deviations are reported in table 5.3.

5.3.5 Discussion

As can be seen in figure 5.11, the results that were obtained show some unrealistic deformations of the aorta. This indicates that retrieving mechanical parameters from the aorta wall and applying those to a solid aorta structure (as opposed to a hollow cylinder-like shape, see section 5.2.2) cannot yield physically accurate results, for the reasons described below. The physical parameters applied to the aorta structure were taken from [31], which is primarily concerned with predicting aortic wall stress and uses a model of the aorta wall. What is proposed in this chapter is a method to deform the entire aorta surface during an EVAR procedure. Modelling

the aorta as a hollow surface would not be accurate for this particular application, as the blood that flows through it exerts pressure on the walls during the intervention, which affects the rigidity of the entire structure. In addition, meshing the wall of the aorta requires much smaller elements than those needed when generating a mesh of the entire volume. The time step used for finite element simulations is proportional to the smallest edge length in the mesh on which the simulation is run. Therefore, finite-element simulations on the aorta wall (as opposed to on the whole volume) would take much more time, which would be an issue for a system designed to run during interventions.

To investigate whether the value of material parameters (the most important ones being the Young's modulus and the Poisson's ratio, as described in [31]) used was the cause of the observed unrealistic deformations, a series of experiments on synthetic data were run.

5.4 Experiment 2: synthetic data

Experiments were carried out on a simple set of synthetic data and were designed to analyse how the finite-element methods performed over a range of material parameters. Although the synthetic data (a tetrahedral mesh of a cube) were not produced to resemble clinical data (by replicating the shape of an aorta for example), applying forces on them provided useful information about the behaviour of a structure for different sets of mechanical parameters. The aim of these experiments was to investigate which values of materials parameters produce deformations visually similar to those observed in vascular interventions.

5.4.1 Data

A cube volume was discretised into tetrahedral elements and finite-element methods were run on it for the same type of deformation (see figure 5.12), using different material parameters each time.

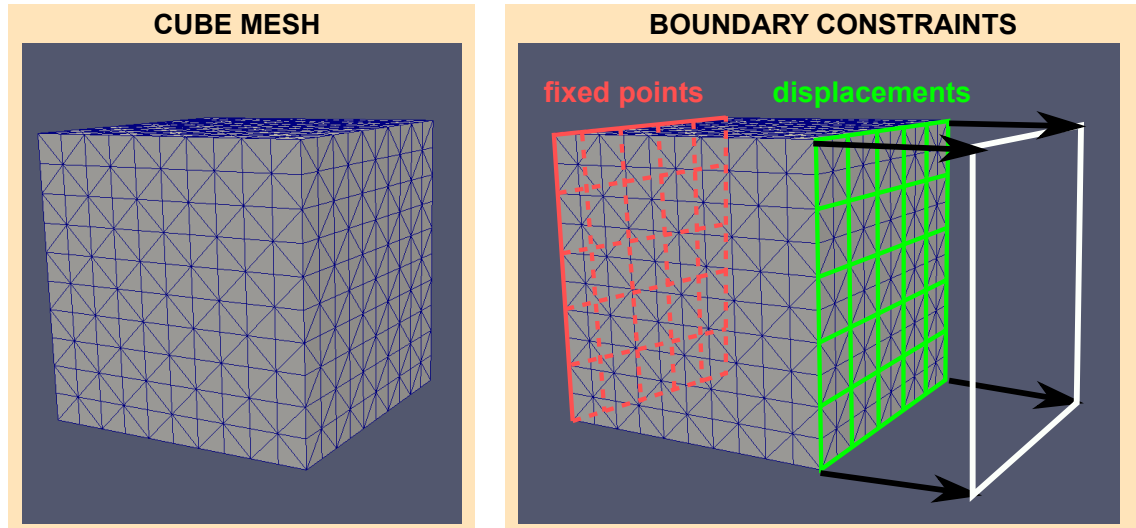


Figure 5.12: Finite-element simulation on a tetrahedral mesh of a cube. Cube mesh: a cube is subdivided into tetrahedral elements. Boundary constraints: the left face nodes are fixed, while the right face nodes are moved in the right direction (displacement lengths are set to 0.4 times the length of the cube edges).

5.4.2 Method

The following constraints were applied to the cube mesh: one face (figure 5.12, left of the cube) was fixed, while all nodes on the opposite face (figure 5.12, right of the cube) were displaced using the same vector (normal to that face, going outwards, and the modulus of which was set to 0.4 times the length of the cube edges).

Using these boundary constraints, finite-element simulations were run for different sets of parameters which are shown in figure 5.13:

- a Young's Modulus of either $10^4 Pa$ or $10^{10} Pa$ for a fixed Poisson's ratio of 0.40 and a fixed material density of $10^3 kg/m^3$
- a Poisson's ratio of either 0.01 or 0.49 for a fixed Young's modulus of $10^8 Pa$ and a fixed material density of $10^3 kg/m^3$
- a material density of either $10 kg/m^3$ or $10^8 kg/m^3$ for a fixed Young's modulus of $10^8 Pa$ and a fixed Poisson's ratio of 0.40

5.4.3 Results

Figure 5.13 shows the results obtained using the different material parameters.




Figure 5.13: Young's Modulus: 1) input data. 2) deformation with low Young's modulus ($E = 10^4 \text{ Pa}$). 3) deformation with high Young's modulus ($E = 10^{10} \text{ Pa}$). Poisson's ratio: 1) input data. 2) deformation with low Poisson's ratio ($\nu = 0.01$). 3) deformation with high Poisson's ratio ($\nu = 0.49$). Density: 1) input data. 2) deformation with low density ($\text{density} = 10 \text{ kg/m}^3$). 3) deformation with high density ($\text{density} = 10^8 \text{ kg/m}^3$).

- an increase in the Young's modulus leads to a stiffer behaviour for the material and to less localised deformations.
- an increase in the Poisson's ratio leads to better preservation of the volume.
- an increase in the density leads to more localised deformations.

From the clinician's knowledge and data (contrast-enhanced X-ray images) collected during EVAR procedures the aorta appears to roughly maintain its cross sectional shape, but can be often translated and straightened when the surgeons exert forces on the aorta surface with their instruments. Within a cross-sectional area the deformation is relatively evenly spread across the aorta. This is contrary to what is shown in figure 5.11. Visual inspection of figure 5.13 shows different types of deformations. Low Young's Moduli or high material densities cause a very limited portion of the cube to be stretched to match boundary conditions. From the clinician's knowledge, the deformation of the aortic wall is much more evenly spread, meaning that the above type of deformation would be unrealistic for an aorta. On the other hand, high Young's moduli or low material densities cause a stretching of the cube lengthwise (and a prodding widthwise) and are much more representative of the kinds of deformation observed inside the operation theatre.

5.4.4 Discussion

The experiments on synthetic data illustrate the importance of selecting a good set of material parameters. Very low Young's moduli are responsible for highly localised deformations, as was also observed in section 5.3. The same effects can be observed for high material densities. High Young's moduli (and low material densities), on the other hand, lead to deformations where the entire aorta is affected by the boundary constraints. The Poisson's ratio parameter is related to a material's compressibility, and high values force the volume of the object to be preserved.

The clinical knowledge of surgeons and contrast-enhanced X-ray images acquired during the intervention provide a body of evidence that the entire cross section of the aorta is displaced by the instruments inserted inside it, as opposed to just a very localised area of the surface in contact with the instruments. For these reasons, the experiments to follow focus on high values of the Young's modulus. The material

density of the aorta has similar effect to the Young's modulus. To reduce the dimension of the parameter space studied in the next experiments, its value was set to 10^3 kg/m^3 to keep it close to values quoted in the literature. The work described in [31] assumes an incompressible aorta wall, however, this chapter uses the mesh of an entire aorta volume as opposed to just the aorta wall, and is therefore not subjected to the same incompressibility conditions. The Poisson's ratio used for the following experiments is allowed to vary from a nearly incompressible material (value of 0.45) to much more compressible materials (value of 0.15).

While the experiments on synthetic data provide useful information on how a simple structure (i.e. a cube) behaves when forces are applied onto it, the simple geometry makes it harder to know exactly the effect a set of parameters has on a more complex structure (such as the aorta described in the first experiment on clinical data) based solely on the knowledge garnered from synthetic data. The results obtained in this section are therefore used to give the directions to follow in the next set of experiments, which are concerned with finding optimal material parameters to use on the aorta volume, by varying Young's moduli and Poisson's ratios.

5.5 Experiment 3: Calculating Optimal Parameter Values

An experiment was conducted to deform the clinical data by taking account of the results obtained on synthetic data.

5.5.1 Data

The same data as those described in section 5.3 were used.

5.5.2 Method

Manual experimentation showed that a Young's modulus $E = 10^{10} \text{ Pa}$ resulted in physically realistic deformations of the aorta, and that the use of Young's moduli

above a certain value $E = 5 \times 10^{10} Pa$ led to greater instabilities in the finite-element simulations and to frequent failures of the non-rigid registration. Three Young's moduli were tested on patient 1, $E = 10^{10} Pa$, $E = 5 \times 10^{10} Pa$ and the last Young modulus $E = 5 \times 10^9 Pa$ which was chosen to be one order of magnitude lower. Along with these three Young's moduli, three Poisson's ratios $\nu = 0.15$, $\nu = 0.30$ and $\nu = 0.45$ were chosen to assess the influence of the volume preservation on the fidelity of the deformation to intra-operative measurements.

For one of the patient datasets every possible combination (E, ν) of Young's modulus and Poisson's ratio were tested. For each combination the deformed aorta surface was overlaid on a DSA image acquired during the intervention. The computation of mean point-to-surface distance errors was used to select the optimal (E, ν) , and the optimal deformation was further checked through visual inspection.

Finite-element simulations were run on the three other patients (patient 2, patient 3 and patient 4), using the optimal Young's modulus and Poisson's ratio found for patient 1.

5.5.3 Validation

An overlay of the deformed aorta surface on an intra-operative DSA image was produced for patient 2, patient 3 and patient 4 if the corresponding finite-element simulations had converged. Point-to-surface distance errors were computed to check whether the non-rigidly registered aorta surface agreed with intra-operative measurements.

5.5.4 Results

Figure 5.14 shows the registered aortas using different Young's moduli and Poisson's ratios in the finite-element simulations using the data from patient 1.

Mean point-to-surface distance errors for different Young's moduli and Poisson's ratio are given in table 5.4. Best results were obtained for a Young's modulus $E = 10^{10} Pa$. The finite-element simulations did not converge when a Young modulus $E = 5.0 \times 10^{10} Pa$ was used in combination with a Poisson's ratio $\nu = 0.45$.




Figure 5.14: Overlay of the aorta surface for deformations obtained with different Young's Moduli and Poisson's ratios. Top left: rigid registration overlay. Top right (Poisson's ratio 0.15), bottom left (Poisson's ratio 0.30) and bottom right (Poisson's ratio 0.45): deformations for different Young's moduli (5 GPa , 10 GPa and 50 GPa) The finite-element simulations failed for a 50 GPa Young's modulus coupled with a 0.45 Poisson's ratio.

	Young's modulus (Pa)	5.0×10^9	1.0×10^{10}	5.0×10^{10}
Poisson's ratio				
0.15		1.87	0.58	0.80
0.30		1.74	0.61	0.66
0.45		1.14	0.54	—

Table 5.4: Mean point-to-surface distance errors (in mm) for different combinations of Young's modulus and Poisson's ratios, calculated using data from patient 1.

The optimal mechanical parameters found for patient 1 (i.e. $E = 10^{10}$ Pa and $\nu = 0.45$) were used to run finite-element simulations on patient 2, patient 3 and patient 4. Deformations were obtained for patient 2 and patient 3. The finite-element simulations did not complete for patient 4 due to instabilities of the simulations. For the finite-element simulations to converge, the time step used in these simulations must be lower than a critical limit [94], and this critical limit for the time step decreases as the Young's modulus increases and/or the smallest characteristic element length decreases. The optimal mechanical parameters found for patient 1 involve a large Young's modulus, which in turn results in a small simulation critical time step for the experiments conducted using such a high Young's modulus. As mentioned in section 5.2.2, strategies were devised to improve the stability of the simulations: 1) small simulation steps were used (7×10^{-6} for a simulation starting at time 0 and finishing at time 1) 2) the displacements applied to the aorta mesh were subdivided if the finite-element simulation failed to converge at a certain step of the collision-detection algorithm. If any of the intermediate simulations failed to converge, the corresponding intermediate displacements were further subdivided into extra-intermediate displacements, and finite-element simulations were carried-out on these extra-intermediate displacements. The simulations carried-out for patient 4 did not complete even after displacements were recursively subdivided, where the large order of recursion (3) makes the use of the proposed method impractical in a clinical setting.

Figure 5.15 shows the deformed aorta surface (along with the rigid registration output) overlaid on an intra-operative DSA image acquired after insertion of a stiff guide-wire and before deployment of the stent-graft for patients 2 and 3. Both overlays show how the wire is encompassed by the deformed aorta surface after running the finite-element simulations and how the aorta edges of the surface are more closely aligned with those of the aorta edges visible inside the DSA image.

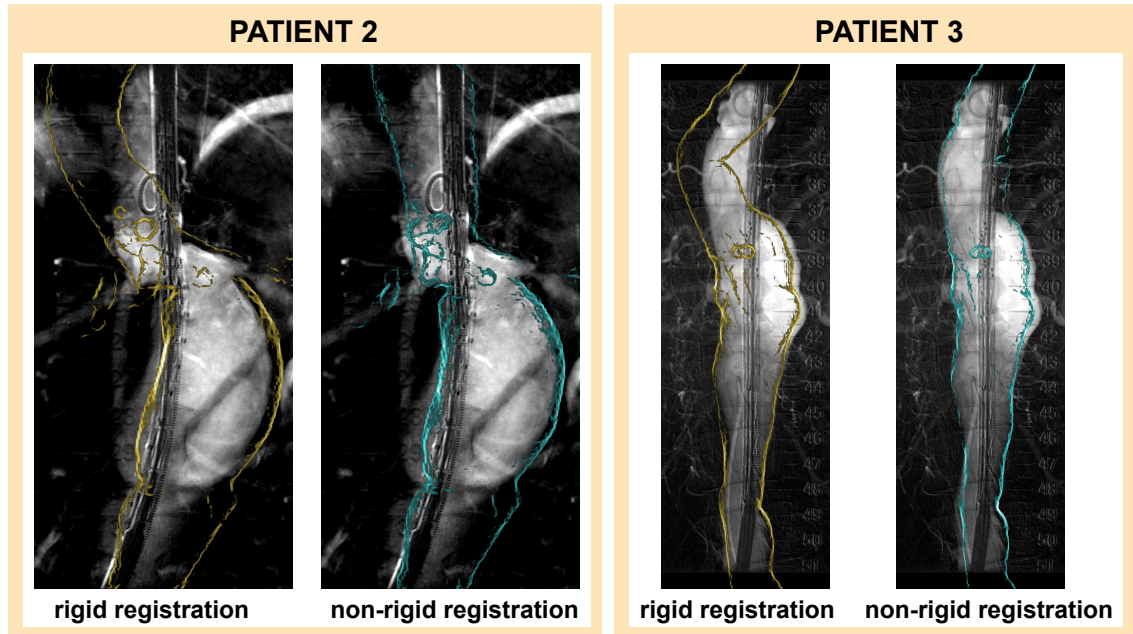


Figure 5.15: Overlay of the registration outcome (rigid and non-rigid) on patient 2 and patient 3. For both patients a Young’s modulus of $E = 1.0 \times 10^{10} Pa$ and a Poisson’s ratio of 0.45 were used in the finite element simulations. Visual inspection shows how non-rigid registration has brought the upper part of the aorta surface closer to the intra-operative ground truth for both patients.

Mean point-to-surface distance errors for patient 2 and 3 are presented inside table 5.5. This table shows how the mean point-to-surface distance errors were at least improved by a factor of 1.4 by using the non-rigid registration proposed in this chapter.

Simulations were run three times for each patient to compute average running times. Average running times of (1 hour, 43 minutes, 58 seconds), (31 minutes, 1 second) and (45 minutes, 27 seconds) were obtained for patient 1, 2 and 3 respectively using a computer powered with an Intel(R) Core(TM) i7-2600K CPU @ 3.40GHz processor and an NVIDIA Corporation GF110 [GeForce GTX 590] graphics card.

5.5.5 Discussion

The results obtained on patient 1 shows a very close agreement between the deformed aorta surface and the intra-operative data for a large set of Young’s moduli and Poisson’s ratios. Visual inspection shows that out of all the parameters that were tested, the optimal Young’s modulus was $E = 10^{10} Pa$. Higher Young’s mod-

left renal	rigid registration		non-rigid registration	
	expert 1 mean (std)	expert 2 mean (std)	expert 1 mean (std)	expert 2 mean (std)
patient 1	6.7 (2.7)	6.2 (3.3)	0.5 (0.4)	0.6 (0.6)
patient 2	2.2 (0.7)	2.1 (1.4)	1.1 (0.6)	1.5 (1.3)
patient 3	2.3 (2.1)	2.6 (2.5)	0.8 (0.5)	1.0 (1.0)

Table 5.5: point-to-surface mean distance and std errors (in mm) computed on patient 1, patient 2 and patient 3 for a rigid registration and for the non-rigid registration using the finite-element method proposed in this chapter. Mechanical parameters were tuned on patient 1. Both rigid and non-rigid registration errors are given for two different experts. Results show how large rigid registration errors are brought down (to less than 1.5 mm) when using the non-rigid registration proposed in this chapter.

uli ($E = 5 \times 10^{10} Pa$) seem to bring the top-left edge of the aorta too close to the stiff wire, while small Young’s moduli ($E = 5 \times 10^9 Pa$) have the opposite effect. In all cases, the deformed aorta surface matches the contrast-enhanced X-ray images more closely than the rigidly registered aorta does. This suggests that, provided that the Young’s modulus and Poisson’s ratio are kept within a range of $5 \times 10^9 Pa$ to $5 \times 10^{10} Pa$ and 0.15 to 0.45 respectively, the non-rigid registration proposed in this chapter will improve on the rigid registration.

The deformed aorta surfaces of patient 2 and patient 3 were more closely matched to the intra-operative DSA image than the rigidly registered aorta surface was. This indicates that the optimal parameters found on patient 1 can be used on a larger population of patients.

The finite-element simulations on patient 4 did not converge because of numerical instabilities. Manual experimentation showed that these numerical instabilities happen mostly for very high values of the Young’s modulus or of the Poisson’s ratio. This means that smaller values for both these parameters may be prudent choices if the non-rigid registration algorithm is used on a wider range of data sets.

To minimise disruption of the clinical setting, the delays required to provide computer assistance to the surgeons have to be kept to a minimum. The durations reported in section 5.5.4 exceed clinically acceptable times. Those times would be approximately lower than 10 minutes for example. As described in section 5.3.2, the experiments on anatomical data were carried out in two stages: 1) a wire was

displaced from an initial position to an intra-operative position (obtained after the stiff guide-wire has been introduced into the patient's aorta) and 2) after it had been positioned at its intra-operative position, the wire was expanded to match the shape of an undeployed stent-graft. With these two stages, the motion of the wire from starting to ending position was broken down into 10 steps. For each of these steps, a finite-element simulation was carried out in case the wire moved outside the 'current' aorta mesh. If the finite-element simulation did not complete successfully, the current step was further subdivided into 10 sub-steps.

The simulations carried out for the initial steps of the first stage (wire moved from a pre-operative position to an intra-operative position) took approximately 50, 75 and 85 seconds for patient 1, 2 and 3 respectively. The farther the wire was moved from its initial position in the first stage, the longer the simulations took. Finite-element simulations carried-out at the end of the second stage (wire expanded to match the shape of an undeployed stent-graft) took approximately 110, 135 and 100 seconds for patient 1, 2 and 3 respectively. It was also observed that simulations carried out for the final steps of the second stage were much more unstable than those performed at the start of the first stage. Those unstable simulations required subdividing corresponding steps into extra sub-steps, which was a reason why the second stage took much more time to complete for the three patients than the first stage.

5.6 Overall Discussion

The finite-element simulations that have been presented inside this chapter can suffer from instabilities which are due to a variety of factors. Careful attention has to be paid when generating a volumetric mesh of the aorta: a trade-off has to be found between meshing the aorta with a small number of tetrahedral elements, which improves the stability of the finite-element simulations but reduces the fidelity to the aorta surface extracted from the pre-operative CT scan, and meshing the aorta with a higher number of tetrahedral elements, which has the opposite effect. In addition, the discrete nature of a mesh, especially when it contains a limited number of nodes, makes it harder to impose boundary conditions, as it is not fully clear where forces should be applied to force the aorta mesh to encompass the moving

wire. In addition to this source of instabilities, the mechanical parameters of the aorta (e.g. Young's modulus and Poisson's ratio) can have a strong impact on the behaviour of the aorta, and conservative values of Young's moduli (i.e. values under a certain limit) should be used to improve the chances of simulations completing successfully.

In the experiments carried out in this chapter, once any finite-element simulation was completed, the next finite-element simulation took the output mesh from the previous simulation as an input. A general guideline for tetrahedral meshes is that elements should be reasonably regular (i.e. small and large angles should be avoided) [89]. As the initial meshes of the aorta (described in section 5.2.2) get deformed by the subsequent finite-element simulations, tetrahedra are likely to become less and less regular. In particular, the tetrahedral elements can become elongated or compressed in any direction. This compressing of tetrahedra can result in an overall smaller characteristic element length in the mesh assembly, which then translates as a smaller critical time limit for the simulation. Sequential updates of the aorta mesh are, as a consequence, the likely reason why simulations are taking more and more time and becoming less and less stable.

The experiments carried out for this chapter used validation data that had been acquired after an (undeployed) stent-graft was introduced inside the patient's aorta. During an EVAR procedure, the main source of motion can be ascribed to the insertion of a stiff guide-wire and the surgeons are primarily concerned with the changes of the aorta geometry that occur directly after this stiff guide-wire has been inserted. For use in an operation theatre, the method presented in this chapter would therefore not require a second stage (where the wire matches the shape of the undeployed stent-graft) and, as a consequence, would take much less time than the experiments reported above, as most instabilities of the finite-element simulations were observed during the second stage. To improve computation times, the methods proposed in this chapter could benefit from avoiding situations where highly 'irregular' meshes are input into finite-element simulations. This can be done by re-meshing the aorta surface after a few simulations (e.g. before carrying out the second stage of the experiments - matching the shape of the undeployed stent-graft). Such a procedure suffers from a few limitations:

- meshing an aorta surface is a complex process which in itself is time-consuming

(though time gains would have likely offset the costs for the experiments presented in this chapter)

- errors in the geometry of the aorta shape can appear if the aorta surface is converted to a mesh representation (and the mesh representation itself converted to a surface representation as depicted in figure 5.7) many times
- the finite-element mesher can fail to generate an accurate mesh. While meshing can easily be performed off-line, producing finite-element meshes inside the operation theatre requires carefully control of the quality of the output mesh, a time-consuming task that might not be achievable within the time frame of a surgical procedure.

Alternatively, strategies could be devised to reduce the number of times the aorta mesh is being updated. This would require moving the wire from its starting to its ending position in less steps than 10, the number that was chosen for these experiments. Such a number was used to bring down the magnitude of displacements applied onto the aorta mesh (and therefore increase the stability of the finite-element simulations), but could be reduced if a smaller Young's modulus and/or Poisson's ratio were to be used.

The methods that have been proposed inside this chapter only made use of the position of the stiff wire to guide the deformation. Additional information could potentially be acquired to further constrain the finite-element simulations, such as other instruments positions if they are visible inside the fluoroscopy images and the soft-tissue edges that can be retrieved through the injection of contrast medium, as long as those are kept to a minimum.

Additional boundary conditions can also be derived from anatomical knowledge. The aorta is connected to other parts of the anatomy: the heart, the lungs, the spine, the kidneys, other organs, etc. Of all those, the anatomical knowledge of clinicians tells us that the spine is the main source of fixations. The aorta is connected to the vertebrae of the patients by the lumbar arteries and ligaments. These connections are hard to represent accurately in a finite-element simulation. As illustrated in figure 5.16, the effect these connections have on the aorta motion could be modelled using a simplified representation, where the aorta is anchored to the vertebrae by cylinders. The shape of the bases of the cylindrical connections can be altered using




Figure 5.16: Anatomical constraints: the aorta is connected to surrounding vertebrae by ligaments and lumbar arteries. These connections are hard to represent in a finite-element simulation. Modelling the constraints: connections between the aorta and the vertebrae can be modelled with a cylinder joining the front of vertebrae to the closest point on the aorta. This simplified model replicates the behaviour of the aorta, with its vertical and rotational motions constrained. Improved modelling: the shape of the cylindric connections can be modified, using ellipses rather than circles as a basis for the cylinders. The main axis of the cylinders is set to be parallel to the centreline of the aorta. This reduces up-and-down motion of the aorta while allowing greater freedom in rotational motion of the aorta around the vertebrae, as would be expected from anatomical knowledge and observations during aneurysm repair procedures.

ellipses rather than circles, in order to further constrain up-and-down motion of the aorta while enabling looser rotational motion around the vertebrae. Adding this type of constraint could allow more physically-realistic deformations of the aorta, but also requires meshing surfaces (the aorta plus surrounding connections) with smaller tetrahedral elements (especially in the case of cylinders with a narrow basis). This requires using smaller time steps for the finite-element simulations, making the approach proposed in this chapter less suited for interventional use.

The results that were obtained show how choosing the right mechanical parameters can lead to a close alignment between the non-rigidly registered aorta surface using the method proposed in this chapter and the digitally-subtracted angiographs acquired inside the operation theatre. Those mechanical parameters can be chosen from a reasonably large range. Preliminary results show how the same set of parameters can improve accuracy, compared to rigid registration, on more than a single dataset.

5.7 Conclusion

This chapter describes the use of finite-element methods to predict physically realistic deformations of the aorta during an EVAR procedure. The procedures that have been presented consist of displacing a ‘virtual’ wire from an idealised position (the centreline of the aorta) towards an intra-operative position that can be computed using intra-operative data and of deforming the aorta so that it encompasses the wire all the time. The input of these methods can be easily acquired during the operation: there is no need for contrast injections (which can induce nephropathy) and only two X-ray images acquired at a different angle (e.g. 20°) are required to compute the position of the intra-operative wire. This can be automatically done following approaches proposed in the literature.

Deformations of the aorta were carried out using a finite-element solver, and required generating a volumetric mesh of the pre-operative aorta and choosing several parameters (simulation parameters, such as time step of the simulation, or mechanical parameters) for the finite-element simulations. Producing a mesh that accurately corresponds to the aorta geometry, while not being overly-detailed, is important, as it ensures that the simulation can be done within time durations that are clinically acceptable and increases the stability of the methods presented in this chapter. Experiments on real data and synthetic data showed how crucial it is to choose a sound set of material parameters. Instabilities of the finite-element simulations occurred, in particular, when high Young’s moduli were used to ensure that the entire aorta is deformed. A good visual agreement was however found between the deformed aorta and the intra-operative digitally subtracted angiographs on three patients. This was corroborated by numerical results, where point-to-surface mean distance errors were found to be lower than 3.0 mm on the three same patients.

The work presented here relies on the computation of the position of a stiff wire inserted before the insertion of a stent-graft. This is only part of the information available during the intervention, where additional data can be gathered at little cost to the patient. The next chapter of this thesis describes the novel use of digital tomosynthesis as an interventional modality to retrieve the 3D positions of calcium

deposits during the EVAR procedures and use those to drive a non-rigid registration algorithm.

Chapter 6

Interventional Digital Tomosynthesis for Computer Assisted Aneurysm Repair

Contents

6.1	Introduction	140
6.2	Theory	140
6.2.1	Digital Tomosynthesis	141
6.2.2	Enhancement of 2D Projection Images during EVAR with Digital Tomosynthesis	142
6.2.3	Retrieving and Using 3D Information during EVAR with Digital Tomosynthesis	148
6.3	Interventional Tomosynthesis Registration Method	152
6.3.1	Automatic Registration of Manually Selected Patches	153
6.3.2	Automatic Registration of Automatically Selected Patches	156
6.3.3	Automatic Correction of Outliers in Registered Automat- ically Selected Patches	158
6.4	Experiments	163
6.4.1	Data	163
6.4.2	Method	164
6.4.3	Validation	166

6.5	Results	167
6.5.1	3D ground truth	167
6.5.2	2D ground truth	167
6.6	Discussion	169
6.7	Conclusion	174

6.1 Introduction

Digital tomosynthesis has been almost entirely used for diagnostic purposes to obtain 3D volumes at lower radiation doses than CT imaging. This chapter investigates interventional uses of digital tomosynthesis, and how the technique can be adapted to the specific application of endovascular abdominal aneurysm repair to drive a 3D-3D registration algorithm based on pre- and intra- operative information on the calcifications of a patient’s aorta.

The first section of this chapter is concerned with previous applications of digital tomosynthesis. Next, the steps taken so that digital tomosynthesis can be adapted to guide a non-rigid registration algorithm are discussed. Following this course of action, the aorta acquired pre-surgery is deformed to represent the intra-operative scene, using a method that is described in a third section. Experiments to validate these methods are then presented, and the results given and discussed.

6.2 Theory

This section first gives a background on digital tomosynthesis (DTS) and explains the main issues involved with the use of this modality. An approach that addresses these limitations is presented in following paragraphs and is concerned with the enhancement of 2D fluoroscopy images during EVAR procedures. A non-rigid registration pipeline which uses DTS images is proposed. This consists of developing a DTS methodology which enhances features useful to the registration algorithm (such as the calcium around the aorta) while suppressing features (like the vertebrae) that would confound the non-rigid registration. Last, 3D correspondences are established between a 3D CT scan obtained before surgery and a 3D digital

tomosynthesis volume acquired inside the operation theatre.

6.2.1 Digital Tomosynthesis

Digital X-ray tomosynthesis is a technique that acquires multiple X-ray images at different views to obtain sectional images of a patient. The mathematical and physical principles underlying this method and its clinical uses (from its inception to 2003) are described in [21]. A comprehensive review of the foundations and applications of digital tomosynthesis is outside the scope of this thesis, therefore only a brief overview of the content of [21] is given below.

Due to its 2D projective nature, conventional X-ray imaging suffers from several limitations: it lacks depth information and produces images where structures are overlaid along the ray projection directions. Some tomographic methods, such as computer tomography, overcome these obstacles by generating 3D volumes of the patient's anatomy using a 180° rotation of the X-ray tube around the patient. This results in a large radiation dose. Tomosynthesis imaging answers the concern of ionisation by capturing a much more limited number of projection images, exposing the patient to comparatively lower doses. However, the reduced angular range results in decreased spatial resolution in the direction approximately perpendicular to the central image plane (that is, the image acquired at the mean position from the angular sweep). Clinical diagnostic uses of the technique include angiography, chest imaging, breast imaging, dental imaging and joint imaging.

Tomosynthesis imaging produces a volumetric image of the patient by reconstructing 2D slice images at different heights using a set of X-ray projection images acquired at different angles (where the motion of the X-ray tube, e.g. a line or an arc, depends on the clinical application). The different structures from the patient anatomy appear at different positions in space in every projection image (see figure 6.1, left). Projection planes can be shifted, as shown in figure 6.1 (right) to line up any particular structures. This is followed by a summation of the information contained in all the projections. The resulting plane contains the structure of interest brought into focus and artefacts are caused by high-contrast structures above and below the focus plane. Out-of-plane clutter (also described as tomographic blur) is superimposed on the objects in sharp focus and therefore obscure details.

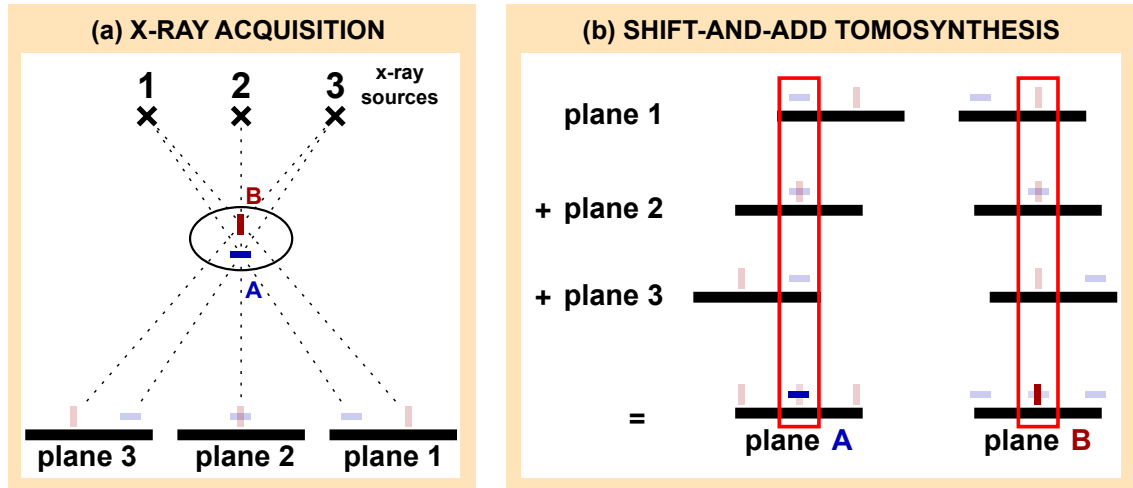


Figure 6.1: (a) X-ray acquisition: X-ray images of two objects A (horizontal stroke) and B (vertical stroke) are acquired at different angles (1, 2 and 3). For each plane the objects A and B are projected at different positions. On the second plane A and B overlap. (b) Shift-and-add tomosynthesis: Shifts are appropriately applied to planes to line up the object A (left) or the object B (right). After the planes have been shifted, the information carried by each projection is summed (plane A, left, plane B, right). This results in the object A (B) being brought into focus on plane A (plane B), while artefacts are visible on the left and right of the plane caused by out-of plane structure B (A).

Tomographic blur reduction is an essential factor to improve the quality of digital tomosynthesis images.

6.2.2 Enhancement of 2D Projection Images during EVAR with Digital Tomosynthesis

Digital tomosynthesis has been proposed in [1] as an interventional modality for EVAR to improve the visualisation of structures of interest inside 2D fluoroscopy images. As explained in the clinical background of this thesis, endovascular repair procedures are performed under X-ray image guidance. While motion of the C-arm is possible during the course of the intervention, large rotations are avoided to limit interruptions of the clinical workflow. Most clinical uses of digital tomosynthesis involve dedicated apparatus (applications can require high frame rate and acquisition of large sets of 2D projection images), where the positions of 2D planes in space are mechanically tracked. The approach presented in [1], however, uses a set of X-ray images obtained from a conventional C-arm X-ray unit (see figure 6.2, (a)).

A 2D/3D registration algorithm can be used to get all the projection images into




Figure 6.2: (a) Digital tomosynthesis clinical setting: A C-arm is rotated around a patient undergoing an EVAR procedure to acquire a series of 2D X-ray images. (b) Interventional digital tomosynthesis: a series of projection images (1 to n) are acquired and show information on the instruments (wire - black) inserted inside the patient, calcified tissues (calcium deposits - green) and bony anatomy (vertebrae - grey). The projection images are all 2D/3D registered to the pre-operative CT scan and put in the same coordinate system. Digital tomosynthesis allows 3D reconstruction of the instruments (white/grey) and the deposits of calcium (green). Artefacts caused by out-of-plane vertebrae (beige) make interpretation of the 3D scene harder.

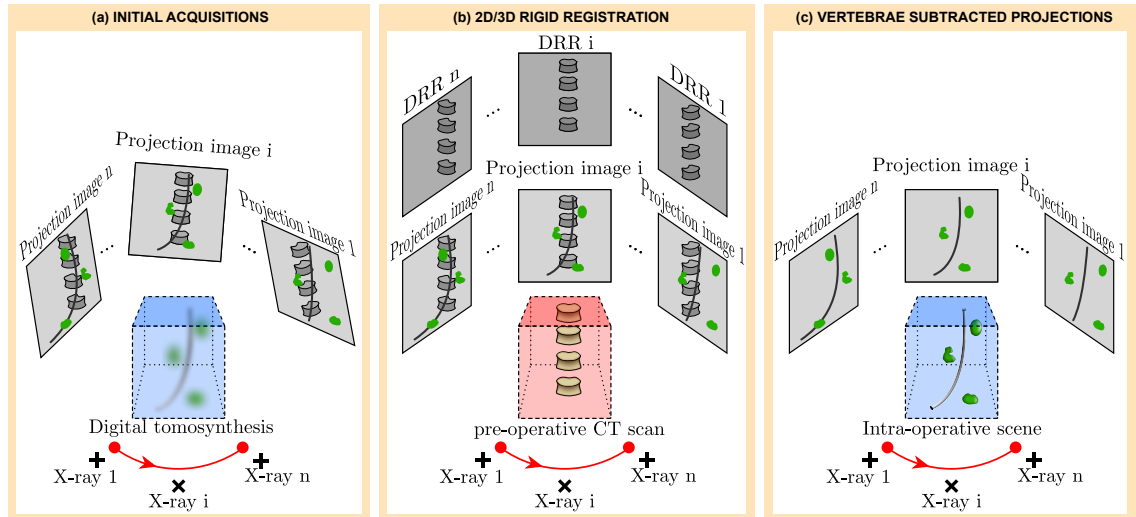


Figure 6.3: (a) Initial acquisitions: a series of projection images (1 to n) are acquired at different views, that contain primarily information on the bony anatomy of the patient (vertebrae - grey), the instruments inside them (wires - black) and calcium deposits (green). If patient movement occurs mechanical tracking is not sufficient to retrieve the exact position of each X-ray images. Generation of a digital tomosynthesis using the 2D images leads to blurring artifacts, as shown in the 3D digital tomosynthesis volume, with blurred vertebrae (beige), calcium deposits (green) and wire (grey/black) (b) 2D/3D rigid registration: each projection image is registered to the CT scan by matching the vertebrae. Once the projection are put into the same coordinate system, a set of digitally reconstructed radiographs (DRR) that only show the vertebrae of the patient are projected from the CT scan at the same positions in space as corresponding interventional X-ray images. (c) Projections with vertebrae subtracted: each DRR image is subtracted from corresponding interventional X-ray projection. Vertebrae are removed as they are the main source of out-of-plane clutter artefacts, obscuring the useful information on the aorta. This step yields a set of projection images that only contain information on calcified soft-tissues and on the instruments inserted inside the patient.

the same coordinate system as depicted in figure 6.3 (a) and (b). The work conducted in [1] makes use of the 2D/3D registration proposed in [69], which is based on matching rigid structures from the pre-operative to intra-operative data and is succinctly explained below. Prior to the intervention, the vertebrae of the patient are automatically segmented. During the surgery, the algorithm searches for the rigid body parameters that register the CT scan to the 3D fluoroscopy coordinate system. An initial estimate of these parameters is obtained via mechanical tracking of the C-arm. A gradient similarity measure is defined to compare one vertebra between the 2D fluoroscopy image and 2D digitally reconstructed radiographs from the CT scan. The similarity measure is maximised using a coarse global optimisa-

tion strategy followed by a hill climbing search. Full automation of the procedure was presented in [97] along with the use of graphics processor units (GPUs) to bring down computation times.

Clutter reduces the ability to extract meaningful information from a 3D digital tomosynthesis volume, as explained in section 6.2.1. In particular, out-of-plane vertebrae make the interpretation of the 3D scene much harder because of the high contrast of bony structures compared to soft tissues, as shown in figure 6.2 (b). The method in [1] improves the quality of the tomosynthesis by: 1) putting all projection images in the same coordinate system (described above) to eliminate motion-induced artefacts, 2) removing the bones from these projection images (described below), 3) fitting a reconstruction surface to the structures of clinical interest (described in a subsequent paragraph). As mentioned above, the first step relies on a rigid 2D/3D registration of the 2D projection images acquired inside the operation theatre with a 3D pre-operative CT scan. This 2D/3D registration is conducted using the vertebrae, consequently their position inside the 2D projection images is known with confidence after the registration took place. Subtraction of a digitally reconstructed radiograph (projected from the CT scan) from a projection image with a certain scaling factor (computed to maximise the pattern intensity of the difference image) yields a ‘de-boned’ 2D image where most of the bones has been removed (see figure 6.3 (b) and (c)).

Instead of reconstructing a full 3D volume, [1] enhances a 2D target image by highlighting the structures of interest on the aorta. Those structures are defined prior to the intervention, using a curved reconstruction surface that has been fitted to the aorta shape and that encompasses landmarks such as the ostia of visceral vessels (see figure 6.4 (a)). After the curved registration surface has been fitted to the pre-operative aorta shape, a ‘sweep’ of projection images are acquired during the surgery to produce the digital tomosynthesis image. The method in [1] proposes the novel use of DTS to ‘automatically reconstruct the information on a clinically relevant section or slice of interest’ and makes use of the reconstruction surface (section of interest) with the following steps:

- A target image is selected, either as one of the projections from the sweep or another image acquired subsequently within the range of that sweep (figure 6.4 (b)). This target image is the image which the surgeons are using to guide




Figure 6.4: From Top to Bottom, Left to Right. (a) Fit reconstruction surface: a reconstruction surface is fitted to the aorta surface segmented from the pre-operative CT volume. This reconstruction surface takes account of the curved geometry of the aorta. In addition, the curvature of the reconstruction surface is adjusted so that it passes through structures of interest (ostia of the visceral vessels - coeliac artery and superior mesenteric artery in the picture). (b) Back-project target pixels: a target image is selected to be enhanced after the digital tomosynthesis sweep has been acquired during the intervention. Each pixel of this target image is back-projected and its 3D intersection with the reconstruction surface computed. (c) Average projection intensities: the 3D intersection point is projected on all the projection images from the ‘sweep’ and all corresponding intensities (I_1 to I_n) are retrieved. (d) Enhance target: the initial pixel from the target image is assigned the average value of (I_1 to I_n). The resulting image is an enhanced target image with better visualisation of the structures of interest (outline of the aorta, calcified areas and ostia of visceral vessels).

their instruments inside the patient.

- All the projections from the sweep are rigidly 2D/3D registered to the pre-operative CT scan to bring them all in the same coordinate system, using the rigid 2D/3D registration method mentioned above.
- Each pixel of the target image is back-projected to the corresponding X-ray source and the 3D intersection of the reconstruction surface to that ray is computed (figure 6.4 (c)).
- The 3D intersection is projected into each 2D image from the sweep. Corresponding pixel intensities are retrieved.
- The 2D pixel intensities are averaged and this averaged value is assigned to the intensity of the 3D intersection point on the reconstruction surface. These steps create the patient-anatomy-specific digital tomosynthesis slice.
- Last, the patient-anatomy-specific digital tomosynthesis slice is projected onto the target image mentioned above. This allows better visualisation of structures of interest such as calcifications and ostia of visceral vessels.

The approach described in this section enables enhancement of 2D images at a low ionisation cost to the patient. In particular, the outline of the aorta can be more easily seen than in standard fluoroscopy, which suffers from relatively poor soft-tissue discrimination. This provides an alternative to the acquisition of X-ray images enhanced with nephrotoxic contrast agents, and therefore opens the door to non-rigid 2D/3D registration methods that use soft-tissue information from two different views (as described in the first experimental chapter of this thesis). However promising this new avenue is, it is not the one this chapter envisions to tread. Rather, the strategies presently devised intend to make use of the 3D nature of digital tomosynthesis volume to drive a 3D-3D non-rigid registration and are presented in the next section.




Figure 6.5: (a) ‘Standard’ digital tomosynthesis: a sweep is acquired of 2D projection images that contain information on the instruments inserted inside the patient (wires), calcifications (calcium deposits) and bony anatomy (vertebrae). The 3D digital tomosynthesis volume shows an overlap between sharply-delineated structures (calcium deposits, wires) and blurred vertebrae. (b) ‘de-boned’ digital tomosynthesis: de-boning of the X-ray images is performed after the acquisition of the sweep. This yields a set of projection images that only contain information on the instruments and calcifications. The resulting digital tomosynthesis volume only shows remaining high-contrast structures.

6.2.3 Retrieving and Using 3D Information during EVAR with Digital Tomosynthesis

Although the algorithms designed in this chapter do not employ patient anatomy-specific 2D reconstruction surfaces, they share two elements with the method detailed in section 6.2.2: the 2D projection images acquired from a sweep during the EVAR procedure are 2D/3D rigidly registered with a pre-operative scan and ‘de-boning’ is used for clutter-reduction, as can be seen in figure 6.5.

Carrying out these two steps, digital tomosynthesis may be used to get better soft-tissues differentiation than that of standard fluoroscopy (especially calcifications of the aorta, see figure 6.6) and to visualise EVAR instruments. As opposed to the 2D images suggested in [1], three-dimensional volumes are retrieved during the intervention, allowing correspondences to be drawn between 3D pre-operative data and 3D intra-operative data. Digital tomosynthesis generates 3D images where the resolution parallel to the projection direction from the centre of the sweep to corresponding X-ray source is lower and highly dependent on the angle between the outermost 2D images from the sweep (see figure 6.7). CT imaging, as per-




Figure 6.6: Comparison of a close-up of an enhanced DTS slice (the volume has been ‘de-boned’) and a close-up of standard fluoroscopy image (at a corresponding location). Calcium deposits are enclosed with light blue squares in the DTS slice. Calcifications are practically invisible in the fluoroscopy image at the same positions (red squares).




Figure 6.7: Example of DTS image showing slices of a digital tomosynthesis volume in three orthogonal directions. The coronal view shows information on the wires inside the patient and on the calcifications of the aorta while the two other views have a much lower resolution.

formed pre-operatively, uses X-ray images acquired from a much larger range of angles and thus the 3D volumes it produces cannot be easily aligned with those obtained through digital tomosynthesis. This chapter proposes the use of the 3D pre-operative CT scan to create a synthetic 3D pre-operative digital tomosynthesis volume to be matched with the real interventional 3D digital tomosynthesis volume (using same-dimensionality registration algorithms).

Several stages are followed to compute the synthetic pre-operative digital tomosynthesis image and these are depicted in figure 6.8. First, the calcified sections of the aorta are manually segmented as those are the parts of the anatomy which move with the clinical anatomy of interest (the aorta) and as they are visible during the repair procedure (although they can hardly be seen using standard fluoroscopy imaging due to very low contrast, see figure 6.6). A sweep of digitally reconstructed radiographs is then simulated, where the view positions and angles mirror those of the images acquired during the intervention. Last, the 3D synthetic digital tomosynthesis is produced from those images.

Contrary to images acquired during surgery, the data gathered before the intervention do not show any information on the instruments inserted inside the patient. A similarity criterion that is robust to wires is therefore required to align pre-operative and intra-operative volumes. The method discussed here uses pattern intensity, as defined in [98], which measures the number of structures/patterns in a difference image (I_{dif}) of the two images to be registered. Formally, this is expressed




Figure 6.8: Registration pipeline. Three steps are performed before the intervention: information from a pre-operative CT is segmented, i.e. 1) the vertebrae of the patient 2) the calcifications of the aorta 3) the aorta surface. During the intervention, the registration process is divided into the following the stages: 4) a sweep of X-ray images is registered to the pre-operative CT scan using information from the vertebrae 5) the registered sweep is used to generate and enhance a DTS image, with no clutter from the bones 6) a sweep of DRR images is generated from a volume of the calcifications of the patient prior to the intervention, at the same views of those of the registered interventional sweep 7) a synthetic DTS volume that only contains 3D information on calcifications is produced from the sweep of DRR 8) the synthetic pre-operative DTS and real intra-operative DTS volumes are registered together. This yields a displacement field that aligns pre- and intra- operative calcifications 9) the pre-operative aorta surface is deformed to its intra-operative representation using the displacement field.

as:

$$P_{r,\sigma} = \sum_{i,j} \sum_{d^2 \leq r^2} \frac{\sigma^2}{\sigma^2 + (I_{dif}(i,j) - I_{dif}(v,w))^2} \quad (6.1)$$

where $d^2 = (i - v)^2 + (j - w)^2$ and σ is a constant that weighs the function so that small deviations in intensity (like those caused by noise) do not have a significant influence on the measure. The difference term $(I_{dif}(i,j) - I_{dif}(v,w))$ measures the difference between one pixel (i,j) and its neighbours (v,w) within a certain radius r . The same value of λ was chosen for the experiments described in section 6.4 as the one used in [98]. The radius $r = 2$ was chosen so that the number of voxels (v,w) satisfying $d^2 \leq r^2$ would be approximately the same as the number of pixels used for the computation of the pattern intensity in [98].

The registration of the pre-operative to the intra-operative volume is done by aligning individual patches of calcium that can be observed prior to and during the intervention. Patches are found either manually (through visual inspection of salient features) or automatically (by detecting areas with a high calcium content, through a process described in section 6.3.2 and in figure 6.12). An optimisation is carried over six degrees of freedom (three translations and three rotations) to match each pre-operative patch to its intra-operative counterpart. Once all the patches are aligned, a set of displacements can be derived from the centres of patches acquired before the intervention to the centres of the patches acquired during it. The use of a transformation model allows interpolation of these displacements to non-rigidly transform the entire pre-operative image. Detailed information is given in the next section on how the registration of individual patches was carried out to produce an overall displacement field for each of the experiments conducted within this chapter.

6.3 Interventional Tomosynthesis Registration Method

Displacement fields that align the aorta calcifications in pre-operative and intra-operative digital tomosynthesis were generated using an automated registration method. Patch selection was carried out using both manually initiated algorithms and fully automated procedures which are described in detail below. Manual interaction was investigated in the first experiments, while full automation is discussed

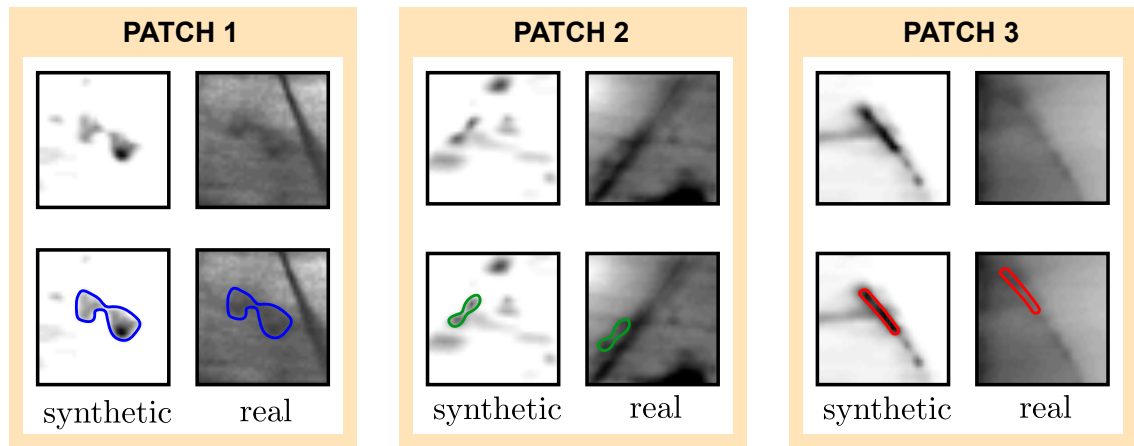


Figure 6.9: Alignment of patches extracted from synthetic and real DTS. The bottom row shows the structures of interest that should be matched between the synthetic and real patch. Patch 1: the structure is visible within both synthetic and real patches. Patch 2: the structure extracted from the synthetic image overlaps with a wire in the real DTS image, making alignment more complicated. Patch 3: the structure extracted from the synthetic image can potentially be matched with different positions in the real DTS image (where the high calcium content makes it harder to register elongated features).

in the two subsequent studies.

6.3.1 Automatic Registration of Manually Selected Patches

Highly calcified areas were retrieved through visual inspection in the 3D pre- and intra- operative DTS volumes. Wires in the interventional DTS, along with imaging artefacts make pre-operative and intra-operative alignment of calcium deposits harder to achieve, as figure 6.9 (patches 2 and 3) shows. Manual patches were selected to encompass features which were visible in both pre- and intra- operative DTS volume (like the structure in the first patch in figure 6.9).

The process to register corresponding pre- to intra- operative patches of calcium is illustrated in figure 6.10. After visual identification of the structures of interest, the centre of each patch was manually selected inside the intra-operative real DTS volume. A fixed (cubic) region of interest (patch) of $40 \times 40 \times 40 \text{ mm}^3$ was extracted inside the real DTS volume around the centre (figure 6.10, step 1), while a cubic search area of $80 \times 80 \times 80 \text{ mm}^3$ was extracted inside the synthetic DTS volume (figure 6.10, step 2). Each (real) region of interest was rigidly registered to its counterpart inside the (synthetic) search area. This involved looking for six parameters




Figure 6.10: Example of patches alignment in 2D following 4 steps: 1) Region of Interest (ROI) selection: The centre of a structure of interest is identified inside the intra-operative DTS volume. A small ROI is defined around that centre. 2) Search area definition: inside the pre-operative DTS volume, a search area is defined around the ROI previously defined. Only transformations that keep the blue object inside the search area are considered. 3) Input Transformations generation: A set of input transformations (1 to 8) that maps out the entire search area space is defined. This set of transformation is used as an input to an optimisation procedure (gradient ascent) that looks for the best alignment between the source and the target features. Only transformation 2 and 3 are close to the solution of the registration problem. 4) Gradient Ascent optimisation of transformations: the optimisation procedure yields a new set of transformations (1 to 8). Only transformations 2 and 3 converged towards the actual alignment transformation. The transformation (2) that returns the highest pattern intensity value of all is kept as the best transformation.

(three translations plus three rotations) using the pattern intensity similarity value (as described in section 6.2.3). Different transformations were used as an input to an optimisation algorithm, which were chosen to be the set of all $t_x, t_y, t_z, \theta_x, \theta_y, \theta_z$, with the rotation θ_x, θ_y and θ_z set to 0, and t_x, t_y and t_z each going from -15 to 15 mm in increments of 2 mm (figure 6.10, step 3). The search area of ± 15 mm was chosen based on reports of the maximum aorta displacements due to interventional instruments of 14.5 mm and 15 mm in [55] and [12] respectively. Transformations where the overall translation (t_x, t_y, t_z) had a modulus greater than 15 mm were therefore automatically rejected. A gradient ascent maximisation was carried out for all these input transformations (figure 6.10, step 4). The rigid patch displacement (denoted by D_i for the i^{th} patch) was selected as the transformation (translation and rotation) that yielded the highest pattern intensity.

As mentioned above, great care was taken to manually select patches where features were visible in both the pre- and intra- operative DTS volumes. For this reason, the registration approach mentioned above is expected to succeed for most of the manually-selected patches. Patches that have been mis-aligned can be eliminated following visual inspection, as described later in section 6.4.2. To correct the mis-registered patches, more sophisticated optimisation procedures can be carried out (like those mentioned in [17]). Such methods are likely to be required in the event that calcium features are not highly visible within either pre- or intra-operative patches and are therefore investigated in section 6.3.2, where the automated selection of patches makes it harder to ensure all patches contain visible calcifications.

As shown in figure 6.11, a displacement field was defined as the set of optimal translations $(t_x, t_y$ and $t_z)$ from the centre of each synthetic (pre-operative) patch to the centre of each real (intra-operative) patch. This displacement field was used to deform the aorta using interpolating thin-plate splines: fixed points were selected at the centre of each vertebra, while source points and target points were chosen as the centres of synthetic (pre-operative) patches and the centres of real (intra-operative) patches respectively.

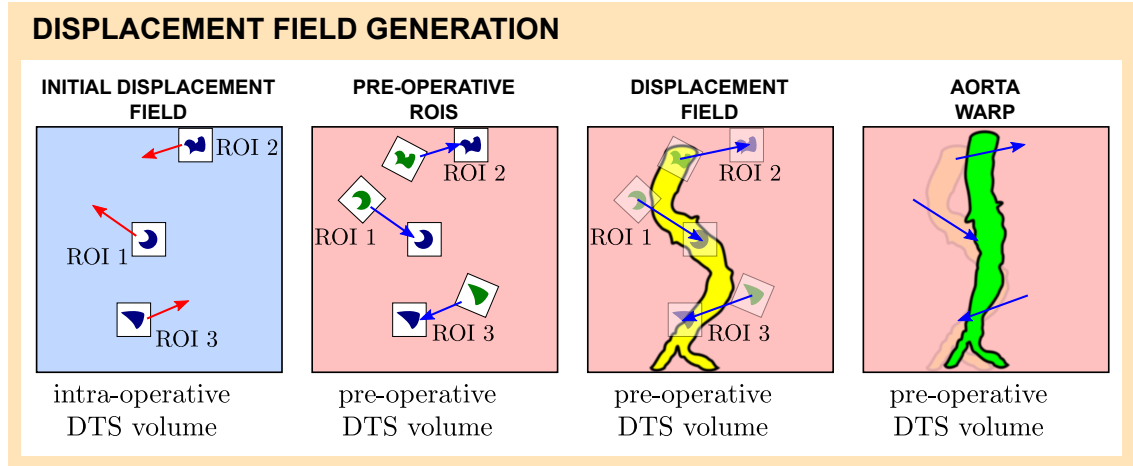


Figure 6.11: Initial displacement field: for each region of interest (ROI) inside the intra-operative DTS volume a transformation is computed as indicated previously. Pre-operative ROIs: corresponding translated and rotated ROIs are computed inside the pre-operative DTS volume. Displacement field: a displacement field is generated that goes from the centre of each ROI inside the pre-operative DTS volume to the centre of corresponding ROI inside the intra-operative DTS volume. Aorta warp: the pre-operative aorta is deformed to its intra-operative shape using the above deformation field.

6.3.2 Automatic Registration of Automatically Selected Patches

In order to further automate the methods patches were also automatically localised in calcified areas surrounding the aorta. This was achieved by applying the following steps (as depicted in figure 6.12):

1. the centrelines of the aorta and iliac arteries were retrieved by skeletonising the three vessels (using sequential morphological erosion), then manually picking points on the skeleton and fitting a polynomial curve to those manually picked points (figure 6.12, 1).
2. the (x_z, y_z) coordinate was computed for a list of z positions going from the bottom of the iliac arteries to the top of the aorta. This resulted in a list of (x_z, y_z, z) coordinates.
3. for all these (x_z, y_z, z) coordinates, a plane was defined perpendicular to the Z axis. On this plane, a grid of points was generated, that was computed with starting x, y positions $x_{min} = x_z - maxDistance$, $y_{min} = y_z - maxDistance$ and ending positions $x_{max} = x_z + maxDistance$, $y_{max} = y_z + maxDistance$,

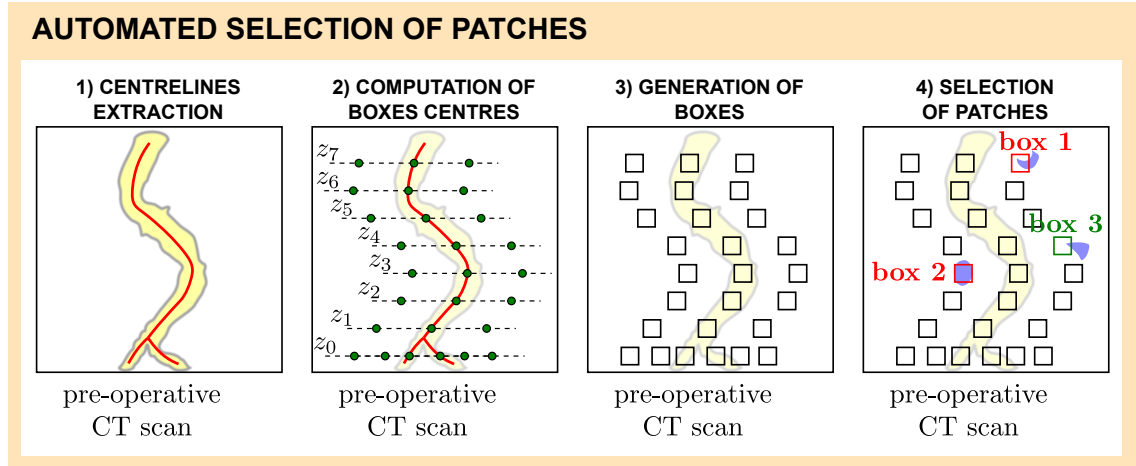


Figure 6.12: 2D example of automated selection of calcium patches. 1) Centrelines extraction: centrelines are extracted from the aorta and the iliac arteries. 2) Computation of boxes centres: points are selected on the aorta centrelines at different z positions (z_0 to z_7). For each of these points, a plane is defined perpendicular to the Z axis, upon which points are gridded. 3) Generation of boxes: Boxes are generated around all the above points. 4) Selection of patches: the amount of calcium is measured within each box. If the amount is too small (e.g. box 3), the box is rejected as not being a patch. If the amount is higher than a certain threshold (e.g. boxes 1 and 2), the box is accepted as a patch.

where $maxDistance$ is a maximum distance parameter computed as a function of the maximum aorta diameter (90 mm in these series of experiments) (figure 6.12, 2).

4. A 40 mm^3 box was centred on each of those points (figure 6.12, 3). For each box, a collection of corresponding CT scan voxels were retrieved. The ratio between the number of calcium voxels (as per the segmentation of calcifications) and the overall number of voxels was retrieved. If this ratio was greater than a user-defined threshold (chosen so that boxes of calcium are approximately visually aligned with the segmented calcifications of the aorta, 0.007 in these experiments) the box was kept in a list of highly calcified boxes (figure 6.12, 4).
5. For each highly calcified box, a $40 \times 40 \times 40\text{ mm}^3$ ROI was defined around the centre of the box (figure 6.12, 4) and extracted from the real interventional DTS volume (see figure 6.10 1).
6. Likewise, for each highly calcified box, a large $80 \times 80 \times 80\text{ mm}^3$ search area was defined around the centre of the box and extracted from the synthetic

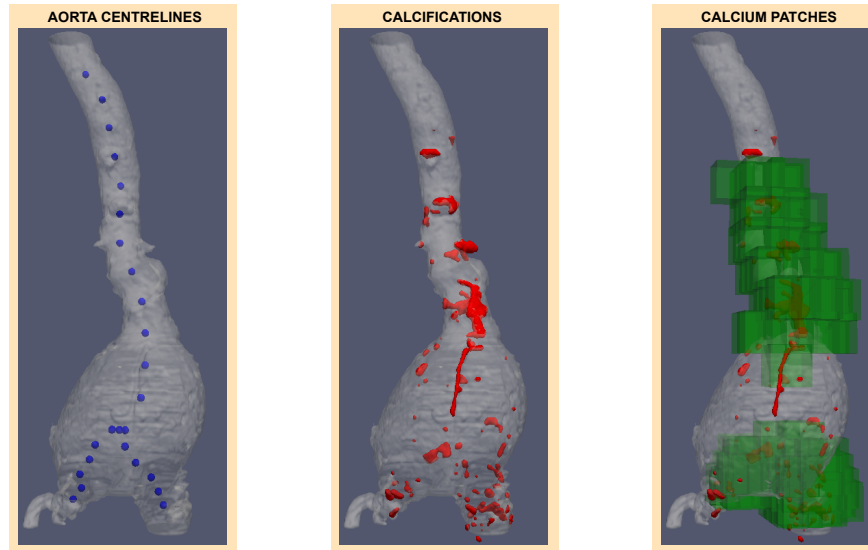


Figure 6.13: Automated selection of patches of calcium: patches of calcium (green boxes, right) are found using the centrelines of the aorta surface (left) and positions of the calcifications (centre) segmented from the CT scan.

pre-operative DTS volume (see figure 6.10 2).

Figure 6.13 shows the patches that were found using this method.

The generation of a displacement field (consisting of all D_i , where D_i is the rigid displacement corresponding to the i^{th} pre- and intra- operative patches of calcium) using the automatically found patches of calcium was conducted as in section 6.3.1. Deforming the aorta using this displacement field was done using the same method as was described with the manually found patches.

6.3.3 Automatic Correction of Outliers in Registered Automatically Selected Patches

The alignment of patches of calcium inside interventional DTS images with their counterparts in pre-operative synthetic DTS images can be hard for the reasons given in 6.3.1: wires in the intra-operative DTS image might obfuscate calcium structures which are visible in the pre-operative DTS image and highly calcified areas can yield several transformation candidates. The automated selection of patches mentioned in section 6.3.2 only considers the calcifications present inside the pre-operative CT scan and therefore does not take wire-induced artefacts inside the




Figure 6.14: Displacement field correction. 1) Initial displacement field: a few displacements (D_2 and D_5) are erroneous. 2) Initial warp: these result in a localised inaccurate warp. The overall deformation is wrong. 3) Corrected displacement fields. The overall displacement field is altered. In particular, D_2 and D_5 are forced to agree with surrounding displacements. 4) The result is a physically realistic deformation.

intra-operative DTS volume into account. Automatic registration of patches following the automated patch selection is therefore likely to result in an erroneous displacement field, where a portion of the displacements are inaccurate. An example is drawn in figure 6.14. Inaccurate displacements cause large localised inaccurate bending of the aorta, which often is very unrepresentative of the type of deformations occurring during surgical procedures.

This section proposes using a non-convex optimisation algorithm on the displacement field to force the bending energy of the overall deformation to a low value.

Similarity value

A similarity value was defined, combining a data term and a bending energy term, which is formalised as follows:

$$SV(D) = \sum_i w_i \cdot PI_i(D) - \lambda \cdot BE(D) \quad (6.2)$$

where SV stands for similarity value, $D = \{D_i\}$ is the displacement field, $PI_i(D)$ is the pattern intensity (as defined in section 6.2.3) between the intra-operative i^{th} patch and its pre-operative counterpart using the transformation $T = D_i$, w_i a weight proportional to the maximum pattern intensity value for these pre-operative

and intra-operative patches (as found above), with $\sum_i w_i = 1$ and $BE(D)$ is the bending energy of the thin-plate spline that is used to deform the aorta (with fixed, source and target points defined as in section 6.3.1). Erroneous displacement fields (i.e. displacement fields where at least one of the pre-operative patches of calcium is matched to a wrong intra-operative location) can still return high values for the term $\sum_i w_i \cdot PI_i(D)$ when pre-operative calcium deposits are matched with intra-operative structures (such as wires and stent-grafts). A misalignment of at least one of the patches can cause very unnatural localised bending of the aorta, resulting in a large bending energy. The bending energy can therefore be used as a regularising term to ensure that a low similarity value (as computed in equation 6.2) is given to the kind of erroneous displacement fields described above.

λ is a weighting parameter that allows the optimisation process to favour deformations that match closely pre- and intra- operative patches of calcium (for small values of λ), or to prefer deformations that do not strongly alter the shape of the pre-operative aorta locally (for high values of λ). A value of $\lambda = 0.1$ was chosen for the experiments presented in this chapter. Such a value was found through manual experimentation to be the highest power of 10 with which pre- and intra-operative patches of calcium were well aligned using the output displacement field from the simulated annealing procedure. For some higher values of λ (e.g. $\lambda = 1.0$) displacement fields were almost translational and failed to account for the non-rigid nature of the transformation taking place during the surgery.

The weights w_i are used to increase the robustness of the optimisation, as they reduce the influence on the overall maximisation process of patches where no correspondences could be found between the pre-operative data and intra-operative data (i.e. a small pattern intensity was obtained).

Simulated annealing optimisation

Correcting inaccurate transformations is done by optimising the displacement field D with respect to the similarity value SV defined above, so that the final displacement field correctly aligns pre- and intra- operative patches of calcium while smoothly deforming the pre-operative aorta shape. This is a computationally costly process, as such an optimisation procedure involves looking at $6 \cdot n$ parameters, i.e. 6 parameters (3 translations, 3 rotations) corresponding to each patch, where there

are n patches. Individual erroneous displacements can result in the initial displacement field (as found in section 6.3.1) being a long way from the optimal solution in the sense of the similarity measure defined above and this similarity measure is likely to have a high number of local extrema (because of imaging artefacts). A non-convex optimisation method is thus required to solve the maximisation problem. A simulated annealing strategy was used to optimise the displacement field from the pre-operative to the intra-operative space. The algorithm followed the procedure described in [17] and sequentially updates the displacement field D applied to the aorta towards the global extremum of the similarity value SV defined above.

The simulated annealing algorithm is briefly summarised below. The ‘current’ displacement field is set as the initial displacement field. The similarity value of the current displacement field is computed (and is referred to as current similarity value in what follows). A random scalar (random perturbation comprised within a certain range, as described in the next paragraph) is added to each translation (i.e. t_x^i , t_y^i and t_z^i) parameter for each i^{th} patch of calcium of the current displacement field in turn, producing intermediate displacement fields. In order to reduce computational costs, the rotation parameters (θ_x^i , θ_y^i and θ_z^i) were not modified in this study. Every time an intermediate displacement is produced, its corresponding similarity value is computed (this similarity value is referred to as intermediate similarity value). If the intermediate similarity value improves on the current similarity value, then the intermediate displacement field is kept (i.e. the current displacement field is updated to the intermediate displacement field and the current similarity value is updated to the intermediate similarity value). If the intermediate similarity value is lower than the current one, then the intermediate displacement field can either be kept or dismissed depending on a temperature parameter that is decreased throughout the course of the algorithm. The process of adding a random perturbation to each parameter of the displacement field (and updating the displacement field as described above) is denoted here by ‘perturbation loop’.

High temperatures allow greater acceptance of intermediate displacement fields that produce a similarity value smaller than that of the current displacement field, while low temperatures tend to only accept intermediate displacement fields that improve on the current one. The simulated annealing procedure starts at a high

temperature, and this temperature is progressively decreased to a low temperature. Practically speaking, high temperatures used at the start of the simulated annealing allow the displacement field parameters to ‘jump’ from local maxima, while the low temperatures at the end of the procedure force the displacement fields to converge towards an optimum. For each intermediate temperature, the perturbation loop is repeated a certain number of times, which is given in the next paragraph for the experiments carried out in this chapter. The process of repeating the perturbation loop (for any temperature) a certain number of times is denoted here by ‘temperature loop’ (where there are as many temperature loops as there are times the temperature is updated).

When very low temperatures are used, intermediate displacements field that are far from the direct vicinity of the current displacement field tend to almost always get rejected. For this reason, the method described in [17] proposes to ensure that, as the temperature decreases, intermediate displacement fields are generated closer and closer to current displacement fields. This is done by diminishing the moduli of the random perturbations throughout the algorithm so that new displacement fields are not automatically rejected.

For the reasons given above, the experiments conducted in this chapter used a high temperature (value of 1.0) at the beginning of the procedure, which was progressively reduced to a low temperature (value of 0.0001), at which stage the procedure was ended and the last displacement field kept as the solution of the optimisation problem. For each temperature, the temperature loop consisted of repeating the perturbation loop 100 times. An annealing rate (value of 0.9) was applied throughout the course of the procedure to reduce the temperature of the simulated annealing: once the temperature loop had been completed the current temperature was multiplied by this annealing rate.

The moduli of the random perturbations were controlled as described below. An initial large value of 100 was assigned to the modulus (m_x^i , m_y^i or m_z^i) of each parameter (t_x^i , t_y^i or t_z^i) of the displacement field at the start of the simulated annealing procedure. In addition, at the start of each temperature loop a counter variable (c_x^i , c_y^i , c_z^i) was assigned to each translation parameter (t_x^i , t_y^i or t_z^i) of the displacement field, with an initial value of 0. While the perturbation loops

were executed within the temperature loop, for each parameter (t_x^i , t_y^i or t_z^i) the corresponding counter value was updated as follows: if adding a perturbation to t_x^i (or t_y^i or t_z^i) produced an intermediate displacement field with a higher similarity value than the current one, then the corresponding counter c_x^i (or c_y^i or c_z^i) was increased by one. At the end of the temperature loop, for each parameter that had a counter lower than 50 (half the number of iterations of the perturbation loop) the corresponding modulus of the random perturbation was divided by two (so that in the next temperature loop, the intermediate displacement fields be closer to the current displacement fields and therefore more displacement fields get accepted). Conversely, for each parameter that had a counter higher than 50 the corresponding modulus of the random perturbation was multiplied by two (reducing the number of displacement fields getting accepted at the next temperature loop). As 1) lower temperatures almost only accept perturbations that improve the current similarity value and 2) those perturbations can only be found in the direct vicinity of local minima, the moduli decreased as the temperature decreased.

Aorta deformation

Deforming the aorta using the displacement field found with simulated annealing optimisation was done as in section 6.3.1.

6.4 Experiments

This section describes the data that were obtained before and during minimally invasive surgery of the aorta and how they were used to drive a registration algorithm to deform a pre-operative surface. Three strategies were devised, with various degrees of automation, to generate a displacement field from calcifications visible in 3D digital tomosynthesis generated either from an intra-operative sweep of fluoroscopy images or from a pre-surgery CT scan. These three methods are presented below, alongside a discussion of the results achieved for each one of them.

6.4.1 Data

As illustrated in figure 6.15, experiments were carried out using data from a patient undergoing an EVAR procedure, during which digital tomosynthesis was performed

prior to the insertion of a stent-graft. The aorta of the patient was segmented from a CT scan (figure 6.15, top left) acquired before the intervention took place. Two surfaces were extracted: a surface of the aorta lumen and a surface of the aorta that includes the thrombus (figure 6.15, top right). The visceral vessels were visible in both representations. Other parts of the anatomy which were segmented from the CT scan include the calcifications (figure 6.15, top right) of the aorta and vertebrae L1 to L5 (figure 6.15, top right).

During the operation, a sweep of X-ray images (figure 6.15, bottom left) were acquired (without contrast-enhancement) previous to introduction of the undeployed prosthesis graft. The sweep comprised 232 images acquired at a 1047 mm source-to-detector distance, with an approximate difference of 55 degrees in angle between the first and the last image acquisitions. The method described in section 6.3 was used and steps 4 to 7 in figure 6.8 were followed to produce a real intra-operative and a synthetic pre-operative DTS volumes of the patient using the pre-operative scan and the above sweep. Figure 6.15, bottom right shows that the intra-operative volume presents information on calcifications and wires, while the pre-operative volume only displays the calcifications.

The methods described in this section are not concerned with the generation of the synthetic pre-operative and real intra-operative DTS volumes. Instead, what follows is the description of the measures that have been taken to register those two volumes together using calcium information.

6.4.2 Method

Following the approach proposed in section 6.3, four methods were used to non-rigidly register the aorta:

- deformation using the displacement field from manual selection of patches (manual selection)
- deformation using the displacement field from manual selection of patches followed by manual correction (manual selection + manual correction): wrongly registered patches were eliminated following visual inspection.
- deformation using the displacement field from an automated selection of patches (automatic selection)

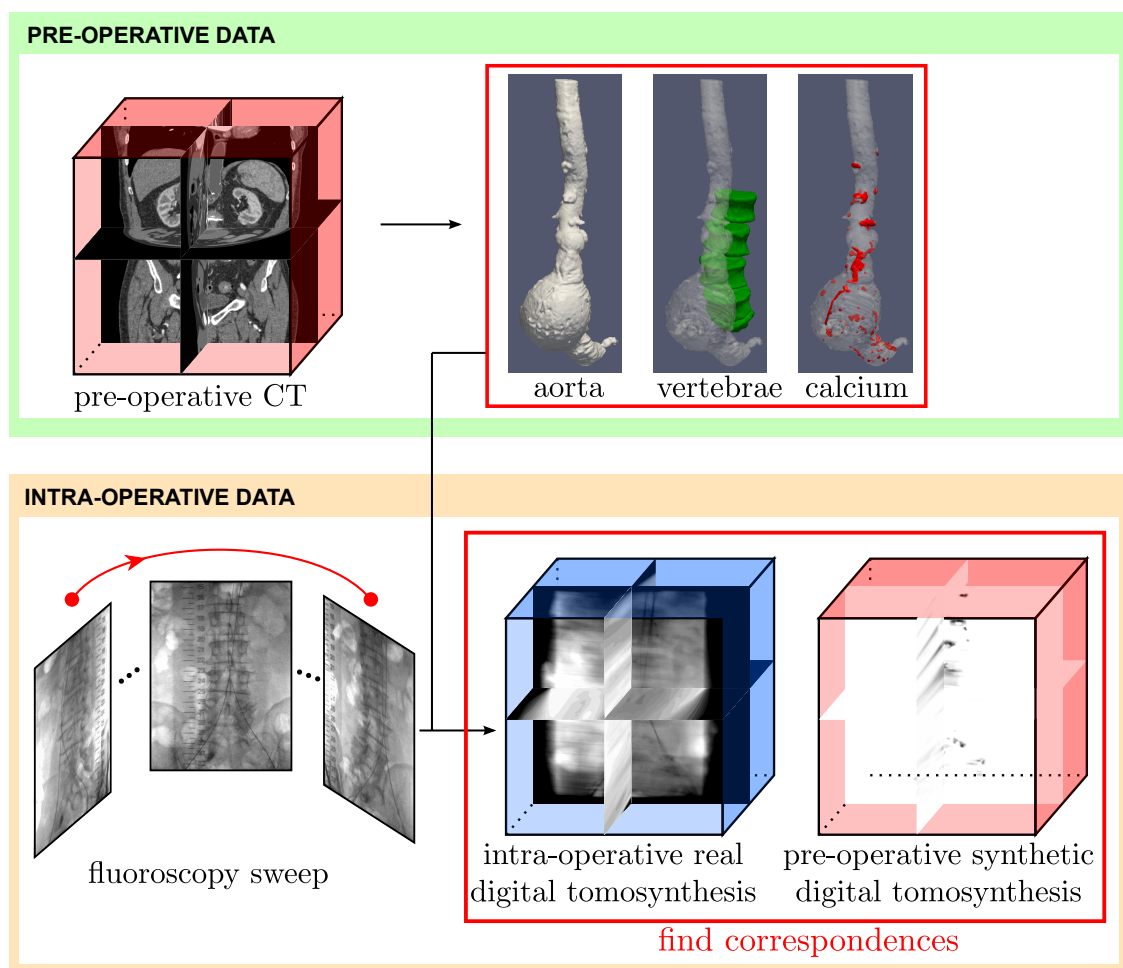


Figure 6.15: Data used in this chapter. The surface of the aorta was segmented from a pre-operative CT scan, along with its calcifications and surrounding vertebrae. During the operation, a sweep of fluoroscopy images were used to produce one synthetic and one real digital tomosynthesis volumes using the surfaces mentioned above. The experimental method presented in this section is not concerned with the steps necessary to produce the data enclosed with the red rectangles, which were described in section 6.2.3. Instead, it only aims at registering the aorta using the two DTS volumes.

- deformation using the displacement field from an automated selection of patches followed by automated correction (automatic selection + automatic correction). Displacements were corrected using a stochastic optimisation algorithm, simulated annealing. This step was repeated 50 times to check for the robustness of the method and resulted in 50 aorta deformations.

6.4.3 Validation

Results were validated using a combination of 3D and 2D ground truth.

3D ground truth

An expert user picked corresponding salient points on calcium deposits in the pre- and intra- operative DTS volumes. Using these points and the positions of the vertebrae from the pre-operative CT scan, a ground truth TPS transformation was computed as in section 6.3.1.

Two experts (expert 1 and expert 2) picked the position of the left and right renal ostia inside the CT scan.

The 3D distances between the left (right) renal ostium inside the 3D ground truth and the left (right) renal ostium inside the deformed aorta volume were computed for a non-rigid registration (i.e. non-deformed pre-operative aorta) and the four different types of deformation given in section 6.4.2.

2D ground truth

A contrast-enhanced X-ray image was acquired after the DTS sweep was performed, post insertion of the undeployed stent-graft. This 2D image was used as ground truth, where two experts (expert 1 and expert 2) picked the 2D position of the left and right renal ostium. Reprojection distance errors (as described in [69]) were computed on the ostia of the renal arteries, and visual inspection of the overlay of the deformed aorta surface on the X-ray image was conducted. Validation was carried out for the same types of deformation as for the 3D ground truth.

expert 1	rigid registration	manual selection	manual selection + correction	automatic selection	automatic selection + correction (mean , std)
left renal	6.2	5.6	1.2	6.6	(2.4 , 0.2)
right renal	4.1	15.9	1.5	16.3	(2.4 , 0.1)

Table 6.1: 3D distances between the left (right) renal ostium position of the 3D ground truth (expert 1) and corresponding ostium position after registration using displacement fields obtained following different methods)

expert 2	rigid registration	manual selection	manual selection + correction	automatic selection	automatic selection + correction (mean , std)
left renal	6.2	5.9	1.1	7.1	(2.3 , 0.2)
right renal	4.3	16.8	1.5	16.8	(2.4 , 0.1)

Table 6.2: 3D distances between the left (right) renal ostium position of the 3D ground truth (expert 2) and corresponding ostium position after registration using displacement fields obtained following different methods)

6.5 Results

6.5.1 3D ground truth

Results are presented in table 6.1 (expert 1) and table 6.2 (expert 2).

Figure 6.16 shows the rigidly registered aorta, the 3D ground truth and the deformed aorta surfaces that were obtained for the displacement fields generated using the methods described in section 6.4.3. As can be seen, only the methods that used a (manual or automated) correction step performed well. Erroneous displacements can be visualised in both the top-right and the bottom-middle picture. While the manual patch selection method only generates a few displacements, the automated method produces a high number of misregistered patches, especially around the renal arteries and near the iliac arteries. The automatic correction of the automated patches selection produces a smooth displacement field, with much more accurately positioned renal ostia.

6.5.2 2D ground truth

Reprojection distance errors obtained for the deformations mentioned above and the 3D ground truth are given in table 6.3 (expert 1) and table 6.4 (expert2).

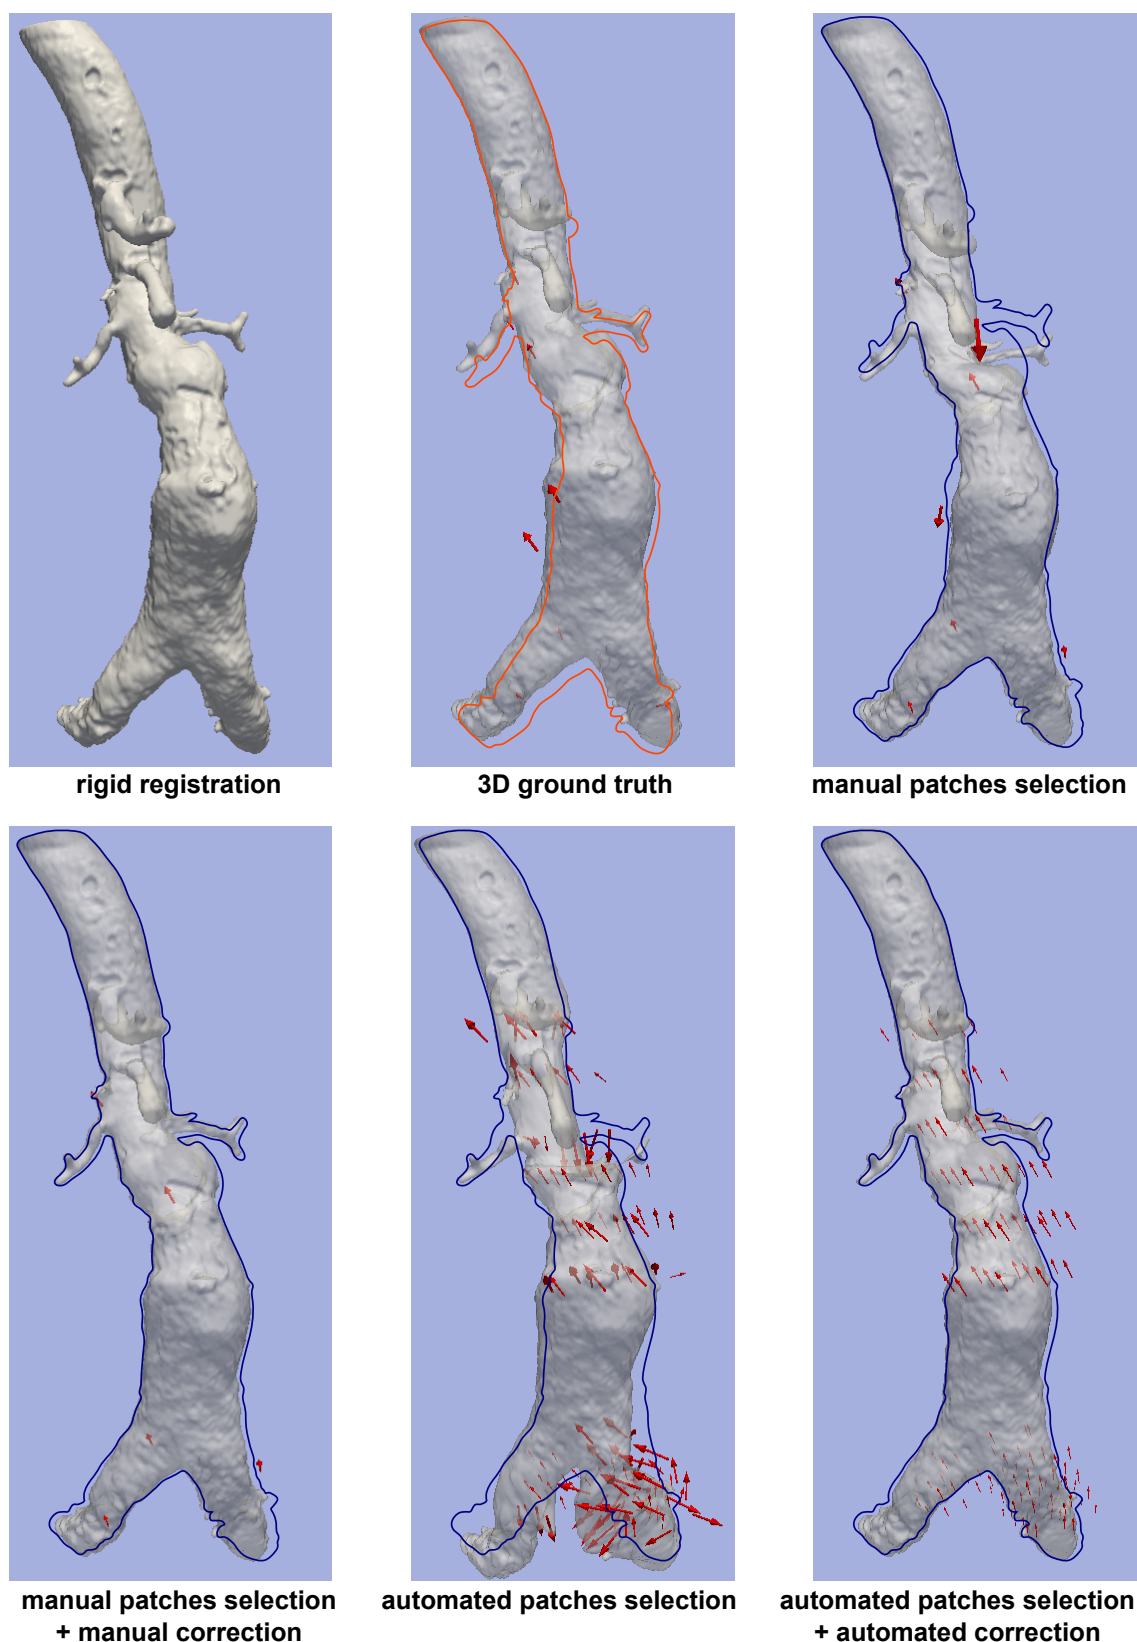


Figure 6.16: Top to bottom, left to right: rigid registration output, 3D ground truth, manual patches selection method output, [manual patches selection method + manual correction] output, automated patches selection method output, [automated patches selection method output + automated correction]. Displacement fields are represented by the red arrows. Top middle picture: the orange outline represents the rigidly registered aorta outline. Top right and bottom pictures: the dark blue outlines represent the 3D ground truth outline.

expert 1	rigid registration	3D ground truth	manual selection	manual selection + correction	automatic selection	automatic selection + correction (mean , std)
left renal	5.1	0.6	4.7	1.3	5.9	(1.4 , 0.1)
right renal	7.4	3.7	17.8	2.2	19.1	(1.6 , 0.1)

Table 6.3: 2D reprojection distances error between the left (right) renal ostium position of the 2D DSA ground truth (expert 1) and displacement fields obtained following different methods)

expert 2	rigid registration	3D ground truth	manual selection	manual selection + correction	automatic selection	automatic selection + correction (mean , std)
left renal	4.4	1.3	4.3	2.0	5.8	(2.1 , 0.1)
right renal	5.8	2.0	17.1	0.8	18.0	(1.0 , 0.1)

Table 6.4: 2D reprojection distances error between the left (right) renal ostium position of the 2D DSA ground truth (expert 2) and displacement fields obtained following different methods)

Figure 6.17 shows the overlay of the deformed aorta onto the DSA ground truth image, along with the positions of the renal ostia picked by expert 1 (the first of the 50 simulated annealing correction was used to produce the image at the bottom right). A good agreement is found between the 3D ground truth and the 2D DSA image in the region of the renal arteries. Discrepancies can be observed at the top of the aorta, above the coeliac artery, which are partly due to the fact that there is limited calcium information in this area (see figure 6.15, top-right) and could be due to further deformation caused by the instruments inserted after the digital tomosynthesis data were acquired. Apart from the 3D ground truth, the positions of the renal arteries ostia only match well for the manual patches selection method followed by a manual correction step and the automated patches selection method followed by an automated correction.

6.6 Discussion

Commercial digital tomosynthesis systems are almost exclusively dedicated to diagnose particular pathologies and are therefore not suited for a surgical environment, where the equipment must not obstruct access for the instruments inserted inside the patient. The method proposed in this chapter uses a standard C-arm X-ray machine to produce a sweep of fluoroscopy images and adapts the standard method for producing digital tomosynthesis volumes to this interventional modality. During




Figure 6.17: Top to bottom, left to right: overlay of rigid registration output, 3D ground truth, manual patches selection method output, [manual patches selection method + manual correction] output, automated patches selection method output, [automated patches selection method output + automated correction] on the DSA validation image. Renal ostia on the DSA image and on the overlay are indicated by yellow rings and red circles respectively.

the operation, X-ray images can only easily be acquired with a limited arc (larger angles require centering the patient), resulting in lower quality volumes. Artefacts must therefore be reduced before the DTS volumes can be input into a registration algorithm. Aligning the X-ray images with the CT scan has been proposed to ensure the position accuracy of the projection planes and to decrease image artefacts. This was further improved by ‘de-boning’ the fluoroscopy images prior to reconstruction. The two steps above result in 3D DTS volumes containing primarily information on the instruments inside the patient and the calcifications of the aorta with considerably reduced bone-caused artefacts.

A sweep of X-ray images is required to produce digital tomosynthesis images. This can be obtained at little cost (in terms of radiation dose) to the patient and only requires minor disruption of the clinical setting. This disruption can be afforded in the EVAR clinical practice, as the method proposed in this chapter has been devised to be used only once, after the introduction of a stiff wire (prior to stent-graft deployment) has caused the main deformation occurring during the operation. Other potential clinical benefits include the enhancement of 2D projection images as proposed in [1], where soft-tissues can be more easily seen, providing an alternative to DSA imaging that requires nephrotoxic contrast injections.

Manually-picked points or automatically-selected boxes of calcium have been proposed in section 6.3 to initiate the registration algorithm. Manual interaction can take a long time and is prone to errors. Automated matching of calcium patches has been suggested as an answer to the above issues. For each patch, the registration is done by using gradient ascent maximisation from a global set of starting positions covering the limited range of magnitude of instrument-induced aortic deformation, and the optimisation procedure is repeated multiple times with a different input transformation to avoid retrieving a local maximum. This procedure is computationally expensive, but can however be partly parallelised to alleviate the computations using multi-core processing, GPU programming and/or distributing the code over several machines (where each gradient ascent procedure can be performed independently from the other gradient ascents). Added to these costs is the correction step that was proposed at the end of the experiments section of this

chapter. Look-up tables can be generated prior to simulated annealing, at the start of the gradient ascent maximisation, so as to reduce the computation power used during the correction step.

Simulated annealing maximisations are stochastic. The parameter λ that weights the contribution of data (i.e. pattern intensities for all patches) and bending energy was chosen so that the optimisations are robust. The displacement correction algorithm devised in this chapter was shown to be robust after being applied on the same data 50 times. This was achieved by designing a similarity value that 1) privileges low bending energies and 2) gives more influence to patches that matched well prior to the optimisation than to those which returned low pattern intensities. As explained in section 6.3.3, the value of λ was chosen (through manual experimentation) to be the highest power of 10 with which pre- and intra- operative patches of calcium were well aligned. Smaller values of λ did not fully correct mis-aligned patches, while higher values of λ resulted in the simulated annealing algorithm returning almost rigid transformations. The use of digital tomosynthesis in EVAR procedures is very recent and only a few operations have been carried out using interventional DTS. For this reason, DTS volumes acquired right after the insertion of a stiff wire are not accessible for a wide population of patients, which limited the experiments conducted in this chapter to one patient. Whether the values of λ (and of other parameters of the simulated annealing optimisation) chosen in this chapter are valid for other patients should be investigated when more validation data is available.

The methods presented in this chapter rely on the calcifications of the patient's aorta. The amount of calcium that can be found varies between patients. Non-calcified areas (such as the top of the aorta in figure 6.15) are responsible for incomplete deformation fields that do not account for the whole deformation that the instruments induced during the intervention. Conversely, highly calcified regions make it harder for an optimisation procedure using pattern intensity to identify the correct transformation matching these regions between the pre- and intra- operative space. A correction step, such as the one proposed in section 6.3.3, is likely to be necessary in the latter case.

Matching calcium patches between images acquired prior to surgery and those acquired during it produced a sparse displacement field. Registration of the pre- and intra- operative data is done by using a transformation model that interpolates the above known displacements. Several models have been proposed in the literature, some (such as splines) providing a mathematical solution to the registration problem, some others taking a more physically-based approach to it (e.g. finite element models). The precise rigid transformation (translation plus rotation) found between each pre-operative patch and its corresponding intra-operative patch is generally not known with absolute certainty, either because the registration of the two patches failed, or because of the limited information contained within the intra-operative DTS volume (especially along the direction parallel to the projection from the central image of the sweep to its corresponding X-ray source). Applying biomechanical models to displacements which are not known with complete certainty is an open problem. Mathematical models such as thin-plate splines can better accommodate for uncertainties, as seen in the first experimental chapter of this thesis.

Validation of non-rigid registration applied to clinical data is hard. The 3D ground truth used for these experiments is derived from 3D interventional digital tomosynthesis. Only limited information can be visually obtained from these data: the volumes suffer from imaging artefacts, wires present inside them obfuscate the calcifications and only a subset of the visually available calcification shows salient features that can be matched to corresponding structures inside the pre-operative images. A gold standard thin-plate spline transformation was derived from the manually selected points on salient calcium features. This gold standard transformation was therefore obtained with a procedure resembling part of the experimental method described in this chapter, while not identical: all experiments use a pattern intensity similarity measure to match corresponding block as opposed to points within the pre- and intra- operative images and manual interaction to select the centre of the blocks only appears in the first experiment. Two-dimensional ground truth was used in addition to 3D ground truth, and did not require computing a gold standard transformation (with all the errors this step is likely to result in). Contrast was not injected simultaneously to the digital tomosynthesis acquisition

and therefore the DSA image used for validation was only available later in time, when deformation might have occurred because of insertion of the stent delivery device. The 2D projected nature of the contrast-enhanced X-ray image means that all 3D structures are projected onto the fluoroscopy plane. This loss of information along the X-ray projection direction means that the differences in position of points from the registered aorta surface and corresponding points from the DSA aorta outline cannot be computed in the 3D space as accurately as if 3D ground truth were available.

Results, using a combination of 3D and 2D ground truth, showed a large improvement of the methods using (1) manually and 2) automatically) selected and corrected displacement fields over a rigid registration: the left renal registration error is more than halved in both validation methods and by both expert ground truth picked points, while the right renal registration error is decreased by a factor of at least 1.7. All registration errors for the methods using correction are below 3 *mm*, the clinical target proposed in [12], meaning that accurate deployment of the stent-graft branches would be possible on the patient whose data these algorithms were applied to.

6.7 Conclusion

A novel use of digital tomosynthesis inside an EVAR operating theatre has been described in this chapter. DTS volumes can be produced and calcium on the aorta enhanced using several techniques. Specifically, registration of the X-ray projection images to the CT scan acquired prior to surgery and ‘de-boning’ of these images can dramatically increase the quality of the 3D DTS images for the purpose of non-rigid registration of the aorta. Using these improved images, strategies were devised to register patches of calcium visible within both pre- and intra- operative data, so as to obtain a displacement field that can be applied to the aorta. Experiments have been performed, with different levels of automation, and show how the method described in this chapter can retrieve the deformation occurring during an intervention. Using an automatically generated and corrected displacement field, registration errors (3D and 2D) lower than the clinical target of 3 *mm* proposed in [12] were obtained,

making Interventional Digital Tomosynthesis a promising modality for use to guide a non-rigid 2D-3D registration during EVAR procedures.

Chapter 7

Conclusion

Contents

7.1 Summary	176
7.2 Computer Assistance and Clinical Constraints	180
7.3 Future Work	181
7.4 Overall Summary	182

This thesis investigated the non-rigid registration of 3D CT scan data acquired prior to an EVAR procedure to 2D X-ray images obtained inside the operating theatre. What follows is a brief summary of the background of the research and the strategies that were devised to tackle the above problem. The next section investigates how well the current methods fit with the constraints inherent to a surgical environment. Future areas of research are proposed in a third section and an overall summary concludes this document.

7.1 Summary

Surgeons get access to the diseased anatomy of a patient by incising their body. Small cuts are preferred to minimise the risk of infection of the site of incision and to decrease the recovery time and the patient's discomfort. With minimally invasive surgery, the branch of surgery where large cuts are avoided, surgeons do not have direct access to the patient's diseased anatomy and rely on intra-operative imaging to guide their instruments. Only limited information is available from the

images acquired during the surgery, as the imaging apparatus must fit with the clinical workflow. Data collected before the intervention takes place (for diagnosis purposes or to plan the operation) are not subject to the same constraints. As a result, information is accessible from pre-operative images that cannot be obtained during the intervention. Computer assistance has been suggested to provide the surgeons with extra-information during surgery by bringing pre-operative images to the intra-operative scene.

Such computer assistance has been proposed for use in EVAR (see [69]), a particular type of minimally-invasive surgery that treats aneurysm of the abdominal aorta by inserting a prosthesis-graft inside the aneurysm to alleviate the pressure on the aneurysm walls: a 3D CT volume acquired prior to the surgery is rigidly registered to 2D fluoroscopy images collected during the intervention so that the pre-operative aorta surface (extracted from the CT data) can be overlaid on intra-operative X-ray images. During EVAR procedures, instruments are inserted inside the aorta and its branches to position the (undeployed) graft inside the aneurysm and eventually deploy it. For a large number of patients, the delivery system is responsible for deformations of the soft-tissues it has been introduced into [55]. As a consequence, image-guidance systems that align the CT-scan and the X-ray images using a rigid registration fail to account for the changes in the shape of the aorta and surrounding soft-tissues that occur during the surgery. The object of this thesis was to recover the aorta deformations using non-rigid registration of the 3D pre-operative CT scan to the 2D intra-operative X-ray images.

Non-rigid registration of medical images involves finding a transformation that accurately deforms a source to a target image. In the context of image guided surgery, most non-rigid registration methods have been proposed for same-dimensionality data, that is 2D-2D or 3D-3D, whereas the pre- and intra-operative data acquired before and during EVAR procedures have a different dimensionality: those are of a 3D and 2D nature respectively. To date, non-rigid registration of 2D to 3D data for image-guided surgery is in its infancy. Many approaches have been presented, that use techniques as varied as algebraic reconstruction technique to measure the distance of a 3D pre-operative image to 2D intra-operative images, statistical shape

modelling, using graphs to represent blood-vessels or manually picking points on an aorta surface and on X-ray images. The previous 2D-3D non-rigid registration approaches that were presented in chapter 3 make assumptions that do not apply to the data obtained from EVAR procedures: for example, [59] assumes that biomechanical models can be used prior to surgery to predict intra-operative deformations, [51] assumes that two contrast-enhanced X-ray images acquired at a different view can be used for registration purposes and [33] assumes that the deformations blood vessels undergo are not strongly influenced by surrounding anatomy. Therefore, these approaches cannot be directly translated to this thesis.

The first experimental chapter of this thesis investigated the use of approximating thin-plate splines for 2D-3D registration purposes. Thin-plate splines can be used to interpolate displacement fields and thus to manually register same-dimensionality data by manually selecting corresponding points inside a source and a target image. Loose matching of those points along specific directions can be done using approximating thin-plate splines coupled with anisotropic error ellipsoids. The method described in chapter 4 uses a 3D aorta surface extracted from a pre-operative CT scan, a 2D contrast-enhanced X-ray image and another 2D X-ray image acquired at a different angle. It allows 3D points picked on the aorta surface to be loosely matched with 2D points picked on either of the two X-ray images, by tailoring an error ellipsoid in accordance with the X-ray projection direction to account for the lack of information perpendicular to the screen. Experiments were carried out on simulated synthetic and real data and on real clinical data from an EVAR procedure. Results show how the tailored approximating thin-plate method yielded registration errors of $2.0\text{ mm}(0.3)$ and $3.8\text{ mm}(0.6)$ (mean and standard deviation for reprojection errors computed using points picked by four different experts) for the left and right renal ostia respectively, less than half of corresponding rigid registration errors which were $4.1\text{ mm}(0.1)$ and $8.7(0.4)$ respectively. As can be seen above, non-rigid registration results for the left renal ostium were below the clinical target of 3 mm proposed in [12], while the right renal ostium was brought closer to this target.

The second experiment presented inside this thesis bypasses the need to inject nephrotoxic contrast into the patient by calculating the 3D positions of instruments during the EVAR procedure, from two X-ray images, and using these positions to drive finite-element simulations on the aorta. An idealised wire is initialised at the centreline of the 3D pre-operative surface and moved towards the intra-operative wire, the positions of which have been computed from two different X-ray views. An algorithm has been implemented to detect collisions between the idealised wire and the aorta. Every time a collision is detected between the moving wire and the aorta, finite-elements simulations are carried out on the aorta surface to force it to encompass the wire. The influence of mechanical parameters were assessed using synthetic data, and experiments on clinical data from four patients undergoing an EVAR procedure showed how the finite-element simulations yielded registrations close to the intra-operative ground truth (contrast-enhanced X-ray images): mean non-rigid point-to-surface distance errors of 0.5 *mm*, 1.1 *mm* and 0.8 *mm* were obtained on three patients while corresponding mean rigid registration errors were 6.7 *mm*, 2.2 *mm* and 2.3 *mm*.

Lastly, a strategy was devised to obtain 3D intra-operative information on the calcifications of the aorta to drive a 3D-3D non-rigid registration of the aorta. Standard digital tomosynthesis (DTS) reconstruction methods have been altered to enhance the calcium deposits on the aorta wall by making use of pre-operative information (the bony-anatomy of the patient). The pre-operative CT scan is used to produce a ‘synthetic’ pre-operative DTS volume with enhanced calcifications. A method has been proposed to automatically register corresponding patches of calcium visible within the ‘synthetic’ pre-operative DTS volume and the real intra-operative DTS volume, which yields an overall displacement field. This displacement field is applied to the pre-operative surface. Using clinical data from one patient undergoing an EVAR procedure, non-rigid reprojection registration errors on the left and right renal ostia of 2.4 *mm* and 2.4 *mm* were obtained, while corresponding rigid registration errors were 6.2 *mm* and 4.1 *mm*, well above the clinical target of 3 *mm* proposed in [12].

7.2 Computer Assistance and Clinical Constraints

Inside the operation theatre computer assistance systems are subject to a number of constraints. How the methods proposed in this thesis perform with respect to these constraints is the object of the next paragraphs.

Automation of the Methods The first approach (thin-plate splines + projection uncertainties) requires manual interaction, with only a few points (less than 10) that need to be picked for the registration to be carried out. Automatic selection of corresponding landmarks would be preferred but is hard to achieve robustly. Manual extraction of an intra-operative stiff wire was conducted for the second experimental chapter (finite-element simulations for use in EVAR). This, if achieved automatically (using the work conducted in [4] for example), would make the second method fully automated. The last method (interventional DTS to drive a 3D-3D non-rigid registration of the aorta) gets a displacement field by automatically registering patches that have been manually or automatically selected.

Disruption of the Clinical Setting The first method presented in this thesis uses manual selection of a limited number of points (which takes less than 10 minutes), the method proposed in the second chapter requires the acquisition of two X-ray images at a different angle, and the third method requires the acquisition of multiple views in a ‘sweep’ over a limited angle (e.g. 40°). Both of these processes could be achieved with only a short (1-2 minutes) interruption to the standard clinical workflow. The introduction of a stiff wire inside the patient’s aorta is the main source of deformation from the clinicians’ knowledge and only occurs once during EVAR. Minor disruption of the clinical setting that only happens once during the intervention is acceptable if it provides the additional accuracy of overlaying a non-rigidly deformed aorta surface on the interventional X-ray images.

Registration Accuracy For all three methods, experiments were carried out on synthetic and real data. Validation is harder to conduct for clinical data acquired prior and during surgery, as 3D ground truth is not readily available during EVAR procedures and contrast-enhanced X-ray images are not available at all time (to

limit nephrotoxicity). Routinely acquired contrast-enhanced 2D X-ray images acquired after the insertion of a stiff guide-wire were used to assess the quality of the registrations performed in this thesis using visual assessment and for calculation of quantitative error measurements. For the first experimental method, reprojection errors of points picked on the left and right ostia of the renal artery came close to the clinical target of 3 *mm* (left renal ostium: mean 2.0 *mm* and right renal ostium 3.8 *mm*) and improve on the rigid registration accuracy (mean 4.1 *mm* and 8.6 *mm* respectively). Mean point-to-surface distance errors were lower than 1.1 *mm* using the proposed method, while the averaged corresponding rigid registration errors equated to 3.7 *mm*. Reprojection errors of 2.4 *mm* were obtained for the left and right renal ostium (lower than the 3 *mm* clinical target) using the non-rigid registration proposed in the third experimental chapter with either a partially or a fully automated registration method. Corresponding left and right rigid registration reprojection errors were 6.2 *mm* and 4.1 *mm* respectively.

7.3 Future Work

A few avenues can be explored to improve the work described in this thesis.

Validation on Clinical Data The methods that have been proposed rely heavily on the use of clinical data to check for their accuracy. Building a database of patient datasets where X-ray images are acquired in a ‘sweep’ fashion after the insertion of the stiff guide-wire is a preliminary to further investigation of the registration accuracy.

Anatomical Constraints Anatomical constraints (i.e. the positions of the centre of vertebrae L1 to L5) were input into the first and third method. Constraints that allow rotational motion of the aorta around surrounding vertebrae while restraining up-and-down motion could be added to the finite-element simulation method to reflect the clinicians’ experience on how the aorta is constrained. Other sources of constraint, such as the calcifications of the aorta, are accessible prior to EVAR procedures and to a certain extent during interventions (using the digital tomosynthesis method proposed in the third experimental chapter for instance) and could be used to further constrain the aorta deformations.

Computation times Fast non-rigid 2D-3D registration is another area requiring further investigation. To avoid significant disruption of the clinical workflow, registration times must be lower than a certain threshold (10 minutes, for example). Although the results for the first method presented in this thesis can be generated almost instantaneously after points have been manually picked, the two other methods currently require long computation times. Improving the aorta mesh generation, along with a better subdivision of the wire displacements could lead to improvements of the methods that rely on the use of finite-element simulations, while the approach described in the third chapter could benefit from parallel programming.

Computer Assistance Care has been taken to make the work proposed in this thesis usable in a clinical setting, as explained in section 7.2. The three approaches that have been discussed inside this thesis have all been shown to be accurate. Validation should be conducted on a larger cohort of patients' data to assess the robustness of the experimental methods. If a registration should fail during an EVAR procedure, it is crucial that strategies be devised to detect its failure, so that no false information is displayed to the surgeon. Last, this thesis assumes that the main source of deformation during EVAR surgery is the stiff wire introduced prior to stent deployment. Carrying-out 2D-3D non-rigid registration at other times of the procedure may also be worthwhile, even though instrument-induced motion is expected to be not as significant as in the case of the deformations induced by the initial introduction of a stiff guide-wire.

7.4 Overall Summary

This thesis has investigated non-rigid registration of the aorta surface for use in an EVAR procedure. The data acquired prior to (a 3D CT scan) and during (2D fluoroscopy images) surgery differ in their dimensionality and on the information they show: soft-tissues are visible inside the pre-operative CT scan and cannot be seen in the intra-operative X-ray images, unless nephrotoxic contrast agents are used. Taking into account clinical constraints, three 2D-3D non-rigid registration methods (that use the intra-operative positions of soft-tissues, instruments and calcium deposits respectively) have been developed, implemented and tested on

real data. The numerical results that have been obtained for these three methods are close to the clinical target of 3 *mm* proposed in [12]. I hope that the generation of computer-assistance systems currently in use during EVAR procedures could, in the future, be improved by the findings of this thesis.

Bibliography

- [1] Mazen Alhrishy, Andreas Varnavas, Tom Carrell, Andrew King, and Graeme Penney. Interventional digital tomosynthesis from a standard fluoroscopy system using 2D-3D registration. *Medical image analysis*, 19(1):137–148, 2015.
- [2] Brian Amberg, Sami Romdhani, and Thomas Vetter. Optimal step nonrigid ICP algorithms for surface registration. In *Computer Vision and Pattern Recognition, 2007. CVPR'07. IEEE Conference on*, pages 1–8. IEEE, 2007.
- [3] Michel A Audette, Frank P Ferrie, and Terry M Peters. An algorithmic overview of surface registration techniques for medical imaging. *Medical image analysis*, 4(3):201–217, 2000.
- [4] Shirley AM Baert, Everine B Van De Kraats, Theo van Walsum, Max Viergever, Wiro J Niessen, et al. Three-dimensional guide-wire reconstruction from biplane image sequences for integrated display in 3-D vasculature. *Medical Imaging, IEEE Transactions on*, 22(10):1252–1258, 2003.
- [5] Shirley Antoinette Maurice Baert. *Guide wire tracking and reconstruction in interventional radiology*. PhD thesis, University Medical Center Utrecht, 2003.
- [6] R. J. Baird, J. F. Gurry, J. F. Kellam, and D. R. Wilson. Abdominal aortic aneurysms: recent experience with 210 patients. *CMA*, 118, 1978.
- [7] B. Timothy Baxter, Michael C. Terrin, and Ronald L. Dalman. Medical management of small abdominal aortic aneurysms. *Circulation*, 117(14):1883–1889, 2008.
- [8] G Berci and KA Forde. History of endoscopy. *Surgical endoscopy*, 14(1):5–15, 2000.

- [9] Abi Berger. How does it work?: Magnetic Resonance Imaging. *BMJ: British Medical Journal*, 324(7328):35, 2002.
- [10] Alexandre Bilger, Eric Bardinet, Sara Fernández-Vidal, Christian Duriez, Pierre Jannin, and Stéphane Cotin. Intra-operative registration for deep brain stimulation procedures based on a full physics head model. In *MICCAI 2014 Workshop on Deep Brain Stimulation Methodological Challenges-2nd edition*, 2014.
- [11] Michael Bock and Frank K Wacker. MR-guided intravascular interventions: Techniques and applications. *Journal of Magnetic Resonance Imaging*, 27(2):326–338, 2008.
- [12] T.W.G. Carrell, B. Modarai, J.R.I. Brown, and G.P. Penney. Feasibility and limitations of an automated 2D-3D rigid image registration system for complex endovascular aortic procedures. *Journal of Endovascular Therapy*, 17(4):527–533, 2010.
- [13] Timothy J Carter, Maxime Sermesant, David M Cash, Dean C Barratt, Christine Tanner, and David J Hawkes. Application of soft tissue modelling to image-guided surgery. *Medical engineering & physics*, 27(10):893–909, 2005.
- [14] Vincent Chan and Anahi Perlas. Basics of ultrasound imaging. In *Atlas of Ultrasound-Guided Procedures in Interventional Pain Management*, pages 13–19. Springer, 2011.
- [15] Kevin Cleary and Terry M Peters. Image-guided interventions: technology review and clinical applications. *Annual review of biomedical engineering*, 12:119–142, 2010.
- [16] Timothy F Cootes, Christopher J Taylor, David H Cooper, and Jim Graham. Active shape models-their training and application. *Computer vision and image understanding*, 61(1):38–59, 1995.
- [17] Angelo Corana, Michele Marchesi, Claudio Martini, and Sandro Ridella. Minimizing multimodal functions of continuous variables with the simulated annealing algorithm. *ACM Transactions on Mathematical Software (TOMS)*, 13(3):262–280, 1987.

- [18] William R Crum, Thomas Hartkens, and DLG Hill. Non-rigid image registration: theory and practice. *The British Journal of Radiology*, 2014.
- [19] Stefanie Demirci, Frode Manstad-Hulaas, and Nassir Navab. Quantification of abdominal aortic deformation after EVAR. In *SPIE Medical Imaging*, pages 72611U–72611U. International Society for Optics and Photonics, 2009.
- [20] S Dindyal, S Rahman, and C Kyriakides. Review of the use of ionizing radiation in endovascular aneurysm repair. *Angiology*, page 0003319714548564, 2014.
- [21] James T Dobbins III and Devon J Godfrey. Digital x-ray tomosynthesis: current state of the art and clinical potential. *Physics in medicine and biology*, 48(19):R65, 2003.
- [22] Mary T Draney, Christopher K Zarins, and Charles A Taylor. Three-dimensional analysis of renal artery bending motion during respiration. *Journal of Endovascular Therapy*, 12(3):380–386, 2005.
- [23] Jean Duchon. Interpolation des fonctions de deux variables suivant le principe de la flexion des plaques minces. *Revue française d’automatique, informatique, recherche opérationnelle. Analyse numérique*, 10(3):5–12, 1976.
- [24] Calvin B. Ernst. Abdominal aortic aneurysm. *New England Journal of Medicine*, 328(16):1167–1172, 1993.
- [25] J Earle Estes. Abdominal aortic aneurysm: a study of one hundred and two cases. *Circulation*, 2(2):258–264, 1950.
- [26] Mark F Fillinger, Steven P Marra, ML Raghavan, and Francis E Kennedy. Prediction of rupture risk in abdominal aortic aneurysm during observation: wall stress versus diameter. *Journal of Vascular Surgery*, 37(4):724–732, 2003.
- [27] J Michael Fitzpatrick, Jay B West, and Calvin R Maurer Jr. Predicting error in rigid-body point-based registration. *Medical Imaging, IEEE Transactions on*, 17(5):694–702, 1998.

- [28] Markus Fleute and Stéphane Lavallée. Nonrigid 3-D/2-D registration of images using statistical models. In *Medical Image Computing and Computer-Assisted Intervention–MICCAI 99*, pages 138–147. Springer, 1999.
- [29] Nicolas Foroglou, Amir Zamani, and Peter Black. Intra-operative MRI (iop-MR) for brain tumour surgery. *British journal of neurosurgery*, 23(1):14–22, 2009.
- [30] F Gebhard, M Kraus, E Schneider, M Arand, L Kinzl, A Hebecker, and L Bätz. Radiation dose in the operating room—a comparison of computer-assisted procedures. *Unfallchirurg*, 106:492–497, 2003.
- [31] George Giannakoulas, George Giannoglou, Johannes Soulis, Thomas Farmakis, Stella Papadopoulou, George Parcharidis, and George Louridas. A computational model to predict aortic wall stresses in patients with systolic arterial hypertension. *Medical hypotheses*, 65(6):1191–1195, 2005.
- [32] Crawford F Gray. Practice-based cone-beam computed tomography: a review. *Primary Dental Care*, 17(4):161–167, 2010.
- [33] Martin Groher, Maximilian Baust, Darko Zikic, and Nassir Navab. Monocular deformable model-to-image registration of vascular structures. In *Biomedical Image Registration*, pages 37–47. Springer, 2010.
- [34] Martin Groher, Darko Zikic, and Nassir Navab. Deformable 2D-3D registration of vascular structures in a one view scenario. *Medical Imaging, IEEE Transactions on*, 28(6):847–860, 2009.
- [35] Pascal Haigron, Aurélien Dumenil, Adrien Kaladji, Michel Rochette, B Bou Said, Simon Esneault, Hélène Walter-Le Berre, Ghizlane Mouktadiri, Miguel Castro, Pierre Louat, et al. Angiovision: Aortic stent-graft placement by augmented angionavigation. *IRBM*, 34(2):167–175, 2013.
- [36] Giselle G Hamad and Myriam Curet. Minimally invasive surgery. *The American Journal of Surgery*, 199(2):263–265, 2010.
- [37] Lianghao Han, John H Hipwell, Christine Tanner, Zeike Taylor, Thomy Mertzaniidou, Jorge Cardoso, Sebastien Ourselin, and David J Hawkes. De-

- velopment of patient-specific biomechanical models for predicting large breast deformation. *Physics in Medicine and Biology*, 57(2):455, 2011.
- [38] George G Hartnell. Imaging of aortic aneurysms and dissection: CT and MRI. *Journal of thoracic imaging*, 16(1):35–46, 2001.
- [39] David J Hawkes, D Barratt, Jane M Blackall, C Chan, Philip J Edwards, K Rhode, Graeme P Penney, Jamie McClelland, and Derek LG Hill. Tissue deformation and shape models in image-guided interventions: a discussion paper. *Medical Image Analysis*, 9(2):163–175, 2005.
- [40] Derek LG Hill, Philipp G Batchelor, Mark Holden, and David J Hawkes. Medical image registration. *Physics in medicine and biology*, 46(3):R1, 2001.
- [41] Yipeng Hu, Timothy J Carter, Hashim Uddin Ahmed, Mark Emberton, Clare Allen, David J Hawkes, and Dean C Barratt. Modelling prostate motion for data fusion during image-guided interventions. *Medical Imaging, IEEE Transactions on*, 30(11):1887–1900, 2011.
- [42] JD Humphrey. Review paper: Continuum biomechanics of soft biological tissues. In *Proceedings of the Royal Society of London A: Mathematical, Physical and Engineering Sciences*, volume 459, pages 3–46. The Royal Society, 2003.
- [43] H Humphreys, JE Coia, A Stacey, M Thomas, A-M Belli, P Hoffman, P Jenks, and CA Mackintosh. Guidelines on the facilities required for minor surgical procedures and minimal access interventions. *Journal of Hospital Infection*, 80(2):103–109, 2012.
- [44] B Jaffray. Minimally invasive surgery. *Archives of disease in childhood*, 90(5):537–542, 2005.
- [45] Kaj Johansen and Thomas Koepsell. Familial tendency for abdominal aortic aneurysms. *Jama*, 256(14):1934–1936, 1986.
- [46] Stian F Johnsen, Zeike A Taylor, Matthew J Clarkson, John Hipwell, Marc Modat, Bjoern Eiben, Lianghao Han, Yipeng Hu, Thomy Mertzaniidou, David J Hawkes, et al. NiftySim: A GPU-based nonlinear finite element

- package for simulation of soft tissue biomechanics. *International journal of computer assisted radiology and surgery*, 10(7):1077–1095, 2015.
- [47] Adrien Kaladji, Aurélien Dumenil, Miguel Castro, Alain Cardon, Jean-Pierre Becquemin, Benyebka Bou-Saïd, Antoine Lucas, and Pascal Haigron. Prediction of deformations during endovascular aortic aneurysm repair using finite element simulation. *Computerized medical imaging and graphics*, 37(2):142–149, 2013.
- [48] Ben Kenwright. A beginners guide to dual-quaternions. *WSCG'2012*, 2012.
- [49] Ali Khojenezhad, Rajeev Rao, and W Anthony Lee. Endovascular repair of abdominal aortic aneurysm. In *Catheter-Based Cardiovascular Interventions*, pages 971–983. Springer, 2013.
- [50] Th M Koller, Guido Gerig, Gabor Szekely, and Daniel Dettwiler. Multiscale detection of curvilinear structures in 2-D and 3-D image data. In *Computer Vision, 1995. Proceedings., Fifth International Conference on*, pages 864–869. IEEE, 1995.
- [51] Rui Liao, Yunhao Tan, Hari Sundar, Marcus Pfister, and Ali Kamen. An efficient graph-based deformable 2D/3D registration algorithm with applications for abdominal aortic aneurysm interventions. In *Medical Imaging and Augmented Reality*, pages 561–570. Springer, 2010.
- [52] Michael J Mack. Minimally invasive and robotic surgery. *Jama*, 285(5):568–572, 2001.
- [53] K Marias, CP Behrenbruch, M Brady, S Parbhoo, and A Seifalian. Multi-scale landmark selection for improved registration of temporal mammograms. In *Proc. IWDM*, pages 580–586, 2000.
- [54] Primoz Markelj, D Tomaževič, Bostjan Likar, and F Pernuš. A review of 3D/2D registration methods for image-guided interventions. *Medical image analysis*, 16(3):642–661, 2012.
- [55] Blandine Maurel, Adrien Hertault, Teresa Martin Gonzalez, Jonathan Sobocinski, Marielle Le Roux, Jessica Delaplace, Richard Azzaoui, Marco

- Midulla, and Stéphan Haulon. Evaluation of visceral artery displacement by endograft delivery system insertion. *Journal of Endovascular Therapy*, 21(2):339–347, 2014.
- [56] Calvin R Maurer and J Michael Fitzpatrick. A review of medical image registration. *Interactive image-guided neurosurgery*, 17, 1993.
- [57] A Melzer, A Schmidt, K Kipfmüller, D Grönemeyer, and R Seibel. Technology and principles of tomographic image-guided interventions and surgery. *Surgical endoscopy*, 11(9):946–956, 1997.
- [58] Jhimli Mitra, Arnau Oliver, Robert Martí, Xavier Lladó, Joan C Vilanova, and Fabrice Meriaudeau. A thin-plate spline based multimodal prostate registration with optimal correspondences. In *Signal-Image Technology and Internet-Based Systems (SITIS), 2010 Sixth International Conference on*, pages 7–11. IEEE, 2010.
- [59] Ashraf Mohamed, Christos Davatzikos, and Russell Taylor. A combined statistical and biomechanical model for estimation of intra-operative prostate deformation. In *Medical Image Computing and Computer-Assisted Intervention MICCAI 2002*, pages 452–460. Springer, 2002.
- [60] Sascha EA Muenzing, Bram van Ginneken, and Josien PW Pluim. On combining algorithms for deformable image registration. In *Biomedical Image Registration*, pages 256–265. Springer, 2012.
- [61] Bart E Muhs, Koen L Vincken, Arno Teutelink, Eric LG Verhoeven, Matthias Prokop, Frans L Moll, and Hence JM Verhagen. Dynamic cine-computed tomography angiography imaging of standard and fenestrated endografts: differing effects on renal artery motion. *Vascular and endovascular surgery*, 42(1):25–31, 2008.
- [62] Stefano Nazari. Expandable device type III for easy and reliable approximation of dissection layers in sutureless aortic anastomosis. ex vivo experimental study. *Interactive CardioVascular and Thoracic Surgery*, 10(2):161–164, 2010.

- [63] Bradley J Nelson, Ioannis K Kaliakatsos, and Jake J Abbott. Microrobots for minimally invasive medicine. *Annual review of biomedical engineering*, 12:55–85, 2010.
- [64] Edward L Nickoloff and Philip O Alderson. Radiation exposures to patients from CT: reality, public perception, and policy. *American Journal of Roentgenology*, 177(2):285–287, 2001.
- [65] Juan C Parodi, JC Palmaz, and HD Barone. Transfemoral intraluminal graft implantation for abdominal aortic aneurysms. *Annals of vascular surgery*, 5(6):491–499, 1991.
- [66] Juan Carlos Parodi et al. Endovascular repair of abdominal aortic aneurysms and other arterial lesions. *Journal of vascular surgery*, 21(4):549–557, 1995.
- [67] Y Pauchard, MR Smith, and MP Mintchev. Modeling susceptibility difference artifacts produced by metallic implants in magnetic resonance imaging with point-based thin-plate spline image registration. In *Engineering in Medicine and Biology Society, 2004. IEMBS'04. 26th Annual International Conference of the IEEE*, volume 1, pages 1766–1769. IEEE, 2004.
- [68] GP Penney, JA Little, J Weese, DLG Hill, and DJ Hawkes. Deforming a preoperative volume to represent the intraoperative scene. *Computer Aided Surgery*, 7(2):63–73, 2002.
- [69] Graeme Penney, Andreas Varnavas, Neville Dastur, and Tom Carrell. An image-guided surgery system to aid endovascular treatment of complex aortic aneurysms: Description and initial clinical experience. In Russell Taylor and Guang-Zhong Yang, editors, *Information Processing in Computer-Assisted Interventions*, volume 6689 of *Lecture Notes in Computer Science*, pages 13–24. Springer Berlin / Heidelberg, 2011. 10.1007/978-3-642-21504-9_2.
- [70] Terry M Peters. Image-guided surgery: from X-rays to virtual reality. *Computer methods in biomechanics and biomedical engineering*, 4(1):27–57, 2001.
- [71] M Prümmer, J Hornegger, M Pfister, and A Dörfler. Multi-modal 2D-3D non-rigid registration. In *Medical Imaging*, pages 61440X–61440X. International Society for Optics and Photonics, 2006.

- [72] A Raheem, T Carrell, B Modarai, and G Penney. Non-rigid 2D-3D image registration for use in endovascular repair of abdominal aortic aneurysms. In *Medical Image Understanding and Analysis*, volume 15, pages 3–157, 2010.
- [73] Joseph J Ricotta and Gustavo S Oderich. Fenestrated and branched stent grafts. *Perspectives in vascular surgery and endovascular therapy*, 20(2):174–187, 2008.
- [74] Torsten Rohlfing, Calvin R Maurer Jr, Walter G O’Dell, and Jianhui Zhong. Modeling liver motion and deformation during the respiratory cycle using intensity-based nonrigid registration of gated MR images. *Medical physics*, 31(3):427–432, 2004.
- [75] Karl Rohr. Elastic registration of multimodal medical images: A survey. *KI*, 14(3):11–17, 2000.
- [76] Karl Rohr, Mike Fornefett, and H Siegfried Stiehl. Spline-based elastic image registration: integration of landmark errors and orientation attributes. *Computer Vision and Image Understanding*, 90(2):153–168, 2003.
- [77] Karl Rohr, H Siegfried Stiehl, Rainer Sprengel, Thorsten M Buzug, Jürgen Weese, and MH Kuhn. Landmark-based elastic registration using approximating thin-plate splines. *Medical Imaging, IEEE Transactions on*, 20(6):526–534, 2001.
- [78] Karl Rohr and Stefan Wörz. An extension of thin-plate splines for image registration with radial basis functions. In *Biomedical Imaging (ISBI), 2012 9th IEEE International Symposium on*, pages 442–445. IEEE, 2012.
- [79] S Rossitti and M Pfister. 3D road-mapping in the endovascular treatment of cerebral aneurysms and arteriovenous malformations. *Interventional Neuro-radiology*, 15(3):283–290, 2009.
- [80] Grace S Rozycki. Surgeon-performed ultrasound: its use in clinical practice. *Annals of surgery*, 228(1):16, 1998.
- [81] Daniel Rueckert, Alejandro F Frangi, and Julia A Schnabel. Automatic construction of 3D statistical deformation models using non-rigid registration.

- In *Medical Image Computing and Computer-Assisted Intervention–MICCAI 2001*, pages 77–84. Springer, 2001.
- [82] Daniel Rueckert, Luke I Sonoda, Carmel Hayes, Derek LG Hill, Martin O Leach, and David J Hawkes. Nonrigid registration using free-form deformations: application to breast MR images. *Medical Imaging, IEEE Transactions on*, 18(8):712–721, 1999.
- [83] Daniel Ruijters, Robert Homan, Peter Mielekamp, Peter Van de Haar, and Drazenko Babic. Validation of 3D multimodality roadmapping in interventional neuroradiology. *Physics in medicine and biology*, 56(16):5335, 2011.
- [84] V Sadick, W Reed, L Collins, N Sadick, R Heard, and J Robinson. Impact of biplane versus single-plane imaging on radiation dose, contrast load and procedural time in coronary angioplasty. *The British journal of radiology*, 2014.
- [85] Natzi Sakalihasan, Raymond Limet, and OD Defawe. Abdominal aortic aneurysm. *The Lancet*, 365(9470):1577–1589, 2005.
- [86] David Sarrut. Deformable registration for image-guided radiation therapy. *Zeitschrift für medizinische Physik*, 16(4):285–297, 2006.
- [87] Joachim Schöberl. NETGEN an advancing front 2D/3D-mesh generator based on abstract rules. *Computing and visualization in science*, 1(1):41–52, 1997.
- [88] Paul Schoenhagen, Uma Numburi, Sandra S Halliburton, Peter Aulbach, Martin von Roden, Milind Y Desai, Leonardo L Rodriguez, Samir R Kapadia, E Murat Tuzcu, and Bruce W Lytle. Three-dimensional imaging in the context of minimally invasive and transcatheter cardiovascular interventions using multi-detector computed tomography: from pre-operative planning to intra-operative guidance. *European heart journal*, 31(22):2727–2740, 2010.
- [89] Hang Si. TetGen, a Delaunay-based quality tetrahedral mesh generator. *ACM Transactions on Mathematical Software (TOMS)*, 41(2):11, 2015.
- [90] S Speidel, S Roehl, S Suwelack, R Dillmann, H Kenngott, and B Mueller-Stich. Intraoperative surface reconstruction and biomechanical modeling for

- soft tissue registration. In *Proc. Joint Workshop on New Technologies for Computer/Robot Assisted Surgery*, 2011.
- [91] Wenhong Sun, Weidong Zhou, and Mingqiang Yang. Medical image registration using thin-plate spline for automatically detecting and matching of point sets. In *Bioinformatics and Biomedical Engineering,(iCBBE) 2011 5th International Conference on*, pages 1–4. IEEE, 2011.
- [92] Vania Tacher, MingDe Lin, Pascal Desgranges, Jean-Francois Deux, Thijs Grünhagen, Jean-Pierre Becquemin, Alain Luciani, Alain Rahmouni, and Hicham Kobeiter. Image guidance for endovascular repair of complex aortic aneurysms: comparison of two-dimensional and three-dimensional angiography and image fusion. *Journal of Vascular and Interventional Radiology*, 24(11):1698–1706, 2013.
- [93] Christine Tanner, Julia A Schnabel, Daniel Chung, Matthew J Clarkson, Daniel Rueckert, Derek LG Hill, and David J Hawkes. Volume and shape preservation of enhancing lesions when applying non-rigid registration to a time series of contrast enhancing MR breast images. In *Medical Image Computing and Computer-Assisted Intervention–MICCAI 2000*, pages 327–337. Springer, 2000.
- [94] Zeike Taylor, Mario Cheng, Sébastien Ourselin, et al. High-speed nonlinear finite element analysis for surgical simulation using graphics processing units. *Medical Imaging, IEEE Transactions on*, 27(5):650–663, 2008.
- [95] Zeike A Taylor, Mario Cheng, and Sébastien Ourselin. Real-time nonlinear finite element analysis for surgical simulation using graphics processing units. In *Medical Image Computing and Computer-Assisted Intervention–MICCAI 2007*, pages 701–708. Springer, 2007.
- [96] Michael A Unser. Splines: a perfect fit for medical imaging. In *Medical Imaging 2002*, pages 225–236. International Society for Optics and Photonics, 2002.

-
- [97] Andreas Varnavas, Tom Carrell, and Graeme Penney. Fully automated initialisation of 2D–3D image registration. In *Biomedical Imaging (ISBI), 2013 IEEE 10th International Symposium on*, pages 568–571. IEEE, 2013.
- [98] Jürgen Weese, Thorsten M Buzug, Cristian Lorenz, and Carola Fassnacht. An approach to 2D/3D registration of a vertebra in 2D X-ray fluoroscopies with 3D CT images. In *CVRMed-MRCAS'97*, pages 119–128. Springer, 1997.
- [99] JE Wickham. The new surgery. *BMJ*, 295(6613):1581–1582, 1987.
- [100] Paul A Yushkevich, Joseph Piven, Heather Cody Hazlett, Rachel Gimpel Smith, Sean Ho, James C Gee, and Guido Gerig. User-guided 3D active contour segmentation of anatomical structures: significantly improved efficiency and reliability. *Neuroimage*, 31(3):1116–1128, 2006.
- [101] Li Zhang. A globally convergent BFGS method for nonconvex minimization without line searches. *Optimization Methods and Software*, 20(6):737–747, 2005.
- [102] Barbara Zitova and Jan Flusser. Image registration methods: a survey. *Image and vision computing*, 21(11):977–1000, 2003.

Publications

Articles in Journal

- AJ Coleman, GP Penney, TJ Richardson, **A Guyot**, MJ Choi, N Sheth, E Craythorne, A Robson, and R Mallipeddi. Automated registration of optical coherence tomography and dermoscopy in the assessment of sub-clinical spread in basal cell carcinoma. *Computer Aided Surgery*, 19(1-3):1–12, 2014.

Articles in Conference Proceedings

- Mazen G Alhrishy, Andreas Varnavas, **Alexis Guyot**, Tom Carrell, Andrew King, and Graeme Penney. Remapping of digital subtraction angiography on a standard fluoroscopy system using 2D-3D registration. In *SPIE Medical Imaging*, pages 94131O–94131O. International Society for Optics and Photonics, 2015.
- **Alexis Guyot**, Andreas Varnavas, Tom Carrell, and Graeme Penney. Non-rigid 2D-3D registration using anisotropic error ellipsoids to account for projection uncertainties during aortic surgery. In *Medical Image Computing and Computer-Assisted Intervention–MICCAI 2013*, pages 179–186. Springer, 2013.

Ethics Statement

The clinical data that were used for the experiments conducted for this thesis were collected with National Research Ethics Service approval 09/H0707/064.