Image Guidance in Telemanipulator Assisted Urology Surgery

Stephen Alexander Thompson

A dissertation submitted for the degree of
Doctor of Philosophy
of the
University of London.

Centre for Medical Image Computing
Department of Medical Physics
University College London

2010

I, Stephen Thompson confirm that the work presented in this thesis is my own. Where information has been derived from other sources, I confirm that this has been indicated in the thesis.

Abstract

This thesis describes an image guided surgery system for use during telemanipulator assisted radical prostatectomy. The thesis is primarily concerned with determining the system accuracy. We first defined a method to present the component errors as an on-screen projection error in pixels. This allowed the error due to each component to be compared and then summed to give a system error.

An MRI image of the patient is transformed into the intraoperative coordinate system, defined by the coordinate system of an optical tracking system. The endoscope is calibrated and tracked during surgery, defining a transformation to the endoscope screen. This transform is used to project the MRI image onto the endoscope video display.

A novel algorithm for registering MRI to ultrasound images of pelvic bone was used to transform the MRI image to the intraoperative coordinate system. This algorithm localises the prostate to within 7 mm, giving an on-screen error of 28 pixels.

The on-screen error due to endoscope tracking was found to be 65 pixels. The high tracking error is caused by a non-normally distributed marker tracking error, highlighting an important shortcoming in the bulk of the image guided surgery literature. Due to the high tracking errors we implemented a limited image guidance system that does not use endoscope tracking. The final part of the thesis details our experience in implementing this system on 5 patients.

The main contributions of this thesis are:

- A robust error analysis of an image guided endoscopic surgery system.
- A novel algorithm for fitting inter patient CT data to an MRI image. The algorithm compares well in with the state of the art for segmenting pelvic bone from MRI images.
- A method to analyse the endoscope tracking error that does not depend on the assumption of normally distributed, homogeneous marker tracking error.

Acknowledgements

My thanks go to my supervisors, David Hawkes, Graeme Penney, and Prokar Dasgupta. Also to Dean Barratt, Tim Carter and Damien Buie who helped greatly with the technical part of the work. The clinical work would not have been possible without the help of Oussama Elhage, Benjamin Challacombe, Declan Murphy, and Maria Nightingale. Thank you also to the patients who volunteered to participate in the study. The work was largely funded by the EPSRC. I am grateful to Prostate UK who funded the clinical implementation of the system.

I could not have completed this work without the support and encouragement of my wife Julia. Thank you also to my parents in law Janice and Peter Bramich, and my aunt Mary Thompson who helped look after our son Samuel on many occasions.

Contents

1 Introduction **21**
1.1 Aim 21
1.2 Contributions 21
1.2.1 Error Analysis 21
1.2.2 Ultrasound to MRI Bone Registration 21
1.2.3 MRI to CT Registration 22
1.2.4 Tracking Error Analysis 22
1.2.5 Clinical Use 22
1.3 Summary 22
1.4 System Description and Structure of the Thesis 23
1.5 Data 29
1.6 Measuring System Error and Summing Component Errors 30
1.6.1 Registration Error at Prostate 30
1.6.2 On Screen Error 30
1.7 Ethics 33

2 Literature Review **35**
2.1 Prostate Cancer and Radical Prostatectomy 35
2.1.1 Description of robot assisted radical prostatectomy. 37
2.1.2 Drawbacks and Opportunities of Robot Assisted Radical Prostatectomy 42
2.2 Image Guided Surgery 42
2.3 Preoperative Knowledge 45
2.3.1 Surrounding Anatomy 46
2.3.2 Prostate Pathology 46
2.4 Create Model ($T_{SM \Rightarrow M}$) 47
2.4.1 Segmentation of Bone from MRI ($T_{SM \Rightarrow M}$) 47
2.4.2 Statistical Shape Modelling in Medical Image Segmentation 52
2.4.3 Segmentation of Soft Tissue Targets from MRI 56
2.5 Preoperative Planning 56
2.6 Registration ($T_{M \Rightarrow ES}$) 57

2.6.1 Direct Registration ($T_{M \Rightarrow ES}$) . . . 57
2.6.2 Registration $T_{M \Rightarrow O}$ . . . 58
2.6.3 Endoscope Tracking and Calibration ($T_{O \Rightarrow ES}$) . . . 64
2.7 Continuous Model Update . . . 65
2.8 Surgeon Interface / Surgeon . . . 66
2.9 Summary . . . 67

3 Segmentation of Bone from MRI 69
3.1 Introduction . . . 69
3.2 Segmentation Algorithm . . . 69
3.2.1 Shape Model Construction . . . 70
3.2.2 Model Fitting . . . 73
3.3 Effect of Number of Shape Model Modes on Error . . . 76
3.3.1 Aim . . . 76
3.3.2 Data . . . 78
3.3.3 Method . . . 78
3.3.4 Results . . . 80
3.4 Extrapolation of the Iliac Shape Using a Shape Model . . . 82
3.4.1 Aim . . . 82
3.4.2 Data . . . 84
3.4.3 Method . . . 84
3.4.4 Results . . . 85
3.5 Results as TRE at the Prostate . . . 86
3.5.1 Aim . . . 86
3.5.2 Method . . . 87
3.5.3 Results . . . 87
3.6 Differential Evolution Optimiser . . . 87
3.6.1 Aim . . . 87
3.6.2 Data . . . 88
3.6.3 Method . . . 88
3.6.4 Results . . . 88
3.7 Fitting the CT shape model to MRI Data . . . 90
3.7.1 Aim . . . 90
3.7.2 Data . . . 92
3.7.3 Method . . . 92
3.7.4 Results . . . 93
3.8 Projection of Prostate TRE to on Screen Error . . . 94
3.8.1 Aim . . . 94
3.8.2 Method . . . 94

3.8.3 Results . . . 94
3.9 Discussion . . . 96
3.9.1 Comparison of Pelvic Shape Model with Results from the Literature . . . 96
3.9.2 Correlation of Optimisation Metric with Boundary Error . . . 96
3.9.3 Result Summary . . . 99
3.10 Conclusion . . . 99

4 Registration of Ultrasound to Fitted Bone Image 101
4.1 Ultrasound Image Acquisition and Processing . . . 102
4.1.1 Ultrasound Image Acquisition . . . 102
4.1.2 Ultrasound Image Pre-Processing . . . 102
4.2 Ultrasound Image to 3D Points . . . 105
4.2.1 Ultrasound Probe Tracking . . . 105
4.2.2 Ultrasound Calibration . . . 111
4.2.3 Transforming the Ultrasound Data . . . 114
4.3 CT Image Pre-Processing . . . 114
4.4 Image to Image Registration . . . 115
4.4.1 Similarity Measure . . . 115
4.4.2 Optimisation Algorithm . . . 115
4.4.3 Summary of Registration Parameters . . . 117
4.5 Selection of Ultrasound Image Processing Parameters . . . 117
4.5.1 Aim . . . 117
4.5.2 Data . . . 117
4.5.3 Method . . . 120
4.5.4 Results . . . 121
4.6 Number and Distribution of Ultrasound Slices . . . 123
4.6.1 Aim . . . 123
4.6.2 Data . . . 123
4.6.3 Method . . . 124
4.6.4 Results . . . 124
4.7 Number of Step Size Reductions . . . 124
4.7.1 Aim . . . 124
4.7.2 Data . . . 126
4.7.3 Method . . . 126
4.7.4 Results . . . 126
4.8 Optimiser Capture Range . . . 127
4.8.1 Aim . . . 127
4.8.2 Data . . . 127
4.8.3 Method . . . 127

4.8.4 Results 127
4.9 Registration Using Real Ultrasound Data 128
4.9.1 Aim 128
4.9.2 Data 128
4.9.3 Method 128
4.9.4 Results 130
4.10 Discussion 132
4.10.1 Optimiser Performance and Registration Run Time 132
4.10.2 Registration Error 132
4.10.3 Conclusion 132

5 Endoscope Calibration and Tracking using an Optical Tracker 135
5.1 Endoscope Tracking $T_{O \Rightarrow ET}$ 135
5.1.1 Calibration - Calculation of $T_{EndoTrkToEndoLens}$ and $T_{Intrinsic}$ 138
5.2 Determination of Calibration and Tracking Errors 140
5.2.1 Monte Carlo Modelling of Calibration 140
5.3 Determination of IRED Localisation Error 141
5.4 In Theatre Calibration of Endoscope 144
5.4.1 Aim 144
5.4.2 Method 144
5.4.3 Results 144
5.5 Estimation of Error Due to Calibration 148
5.5.1 Aim 148
5.5.2 Method 148
5.5.3 Data 151
5.5.4 Results 151
5.5.5 Comparison of Actual Versus Simulated Errors 151
5.6 Estimation of On Screen Error due to IRED Tracking Error 154
5.6.1 Aim 154
5.6.2 Method 154
5.6.3 Data 154
5.6.4 Results 154
5.7 Effect of Errors on Calibration Accuracy 154
5.7.1 Aim 154
5.7.2 Method 155
5.7.3 Results 155
5.8 Improvement Of Tracking Error Using Motion Constraints 156
5.8.1 Aim 156
5.8.2 Method 157

5.8.3 Results . . . 159
5.9 Using Normally Distributed IRED Tracking Error . . . 159
5.9.1 Aim . . . 159
5.9.2 Methods . . . 159
5.9.3 Results . . . 160
5.10 Discussion . . . 162
5.11 Error Summation . . . 163

6 Patient Trials **165**
6.1 Introduction . . . 165
6.2 MRI Data . . . 166
6.3 Ultrasound Data Capture . . . 166
6.3.1 Patient Positioning During Imaging . . . 167
6.3.2 Processing of Ultrasound Data . . . 168
6.4 Acquisition of Video Data . . . 168
6.5 Estimation of Tracking Error for a Moving Endoscope . . . 169
6.5.1 Aim . . . 169
6.5.2 Data . . . 169
6.5.3 Method . . . 171
6.5.4 Results . . . 171
6.5.5 Estimated Tracking Error . . . 171
6.5.6 Discussion . . . 173
6.6 Direct Alignment of the MRI Image to the Endoscope Image . . . 173
6.7 Display of MRI Overlaid on Endoscope Screen . . . 175
6.7.1 Projection Model . . . 176
6.7.2 Motion Tracking . . . 178
6.8 Clinical Experience . . . 179
6.9 Conclusions . . . 181

7 Conclusions and Further Work **183**
7.1 Summary of Results . . . 183
7.2 Future Work . . . 185
7.2.1 Shape Model Fitting . . . 185
7.2.2 Registration Using Ultrasound . . . 185
7.2.3 Endoscope Tracking . . . 185
7.2.4 Clinical Implementation . . . 186
7.3 Conclusions . . . 186
7.3.1 Defining Clinical Utility . . . 186
7.3.2 Defining And Improving Accuracy . . . 188

7.3.3 Combining Feasibility and Utility 188

List of Acronyms **190**

Bibliography **193**

Publications and Presentations **203**

A Details of Data used in Study **207**

A.1 Shape Model Patient A 208
A.2 Shape Model Patient B 209
A.3 Shape Model Patient C 210
A.4 Shape Model Patient D 211
A.5 Shape Model Patient E 212
A.6 Shape Model Patient F 213
A.7 Shape Model Patient G 214
A.8 Shape Model Patient H 215
A.9 Shape Model Patient I 216
A.10 Shape Model Patient J 217
A.11 Shape Model Patient K 218
A.12 Shape Model Patient L 219
A.13 Shape Model Patient M 220
A.14 Shape Model Patient N 221
A.15 Shape Model Patient O 222
A.16 Shape Model Patient P 223
A.17 Shape Model Patient Q 224
A.18 Shape Model Patient R 225
A.19 Shape Model Patient S 226
A.20 Shape Model Patient T 227
A.21 Shape Model Patient U 228
A.22 Prostatectomy Patient 01 229
A.23 Prostatectomy Patient 02 230
A.24 Prostatectomy Patient 03 231
A.25 Prostatectomy Patient 04 232
A.26 Prostatectomy Patient 05 233
A.27 Brachytherapy Patient 01 234
A.28 Brachytherapy Patient 02 235
A.29 Cadaver Patient 01 236
A.30 Cadaver Patient 02 237
A.31 Cadaver Patient 03 238

A.32 Plastic Anatomy Phantom 239

B Fitting a Male Shape Model to a Female Pelvis 241

C Estimating the Tracking Synchronisation Error Using Image Information 245

C.1 Aim 245

C.2 Method 245

C.3 Data 245

C.4 Results 246

D Error due to using a Nominal Prostate 249

List of Figures

1.1 Example of MRI Overlay 23

1.2 MRI Image Transformations 25

1.3 Ultrasound Image Transformations 26

1.4 Endoscope Transformations 27

1.5 Direct Registration Transformations 28

1.6 Nominal Prostate Points 31

1.7 Error Projection Example 32

2.1 Prostate and Surrounding Anatomy 38

2.2 RARP, Step 1 and 2. 38

2.3 RARP, Step 3 and 4 39

2.4 RARP, Step 5 and 6 39

2.5 RARP, Step 7 and 8 40

2.6 RARP, Step 9 and 10 40

2.7 RARP, Step 11 and 12 41

2.8 RARP, Step 13 and 14 41

2.9 RARP, Step 15 and 16 42

2.10 Example of Segmented Anatomy Overlay 44

2.11 Block diagram of proposed surgical guidance system. 45

2.12 Ultrasound, CT and MRI 48

2.13 Ultrasound, CT and MRI by Slice 49

2.14 T1 and T2 Weighted MRI 51

2.15 Direct Alignment Using Pubic Arch Surface 58

3.1 Some example MRI edge profiles 72

3.2 Slice Distribution for Registration 79

3.3 TRE vs Number of Model Modes 81

3.4 Example Ultrasound Slice Location 83

3.5 MRI Extents 84

3.6 TRE Results for Masked Data 85

3.7 TRE versus Dice and Boundary Error 87

3.8 Differential Evolution Evaluation . . . 89
3.9 Repeatability of Differential Evolution . . . 91
3.10 Example Segmentation Results . . . 93
3.11 Overlay of Segmentation Error on Surgeon Display . . . 95
3.12 Comparison of Normally Distributed and Actual Errors . . . 95
3.13 Shape Model Evolution . . . 98
3.14 Overlay of Segmentation Error on Surgeon Display . . . 100

4.1 Typical Unprocessed Ultrasound Image . . . 102
4.2 Ultrasound Tracking . . . 103
4.3 Ultrasound Image Processing . . . 106
4.4 Probability Look Up Tables . . . 107
4.5 Ultrasound Tracker Geometry . . . 107
4.6 TRE vs FLE for Different IRED Configurations . . . 110
4.7 Reference and tracker geometry . . . 111
4.8 TRE for Different IRED Configurations . . . 112
4.9 Pinhead Ultrasound Calibration Object . . . 112
4.10 Ultrasound Image of the Calibration Pin Head . . . 113
4.11 CT Preprocessing . . . 114
4.12 Pelvic Phantom for Registration Experiments . . . 119
4.13 Registration TRE vs Artifact Threshold and Blur Radius . . . 122
4.14 Registration TRE versus Image Similarity . . . 123
4.15 TRE versus Number of Slices . . . 125
4.16 TRE versus Terminal Step Size . . . 126
4.17 TRE Versus Start Distance . . . 128
4.18 Overlay of Registration Error on Surgeon Display . . . 133

5.1 Tracking Collar and Optotrak Cameras . . . 136
5.2 Tracking Collar Schematic . . . 137
5.3 TRE / FLE Ratio for Endoscope Tracker . . . 137
5.4 Endoscope Calibration Method . . . 139
5.5 Distribution of Real Calibration Views . . . 145
5.6 Endoscope Tracking Points . . . 145
5.7 Example Endoscope Calibration Images . . . 147
5.8 IRED Visibility . . . 148
5.9 ILE for Each Frame in the Calibration Data . . . 149
5.10 Linear Fitting To IRED Errors . . . 150
5.11 Histogram of Average Back Projection Errors . . . 151
5.12 Point Projection Error Due to Calibration . . . 152

5.13 Scatter Plot of Back Projection Error Versus Actual Error . . . 152
5.14 Comparison of Actual and Simulated Errors . . . 153
5.15 Endoscope Tracking Error Overlay . . . 154
5.16 Calibration Projection Error Versus Grid Extraction Error . . . 156
5.17 Calibration Projection Error Versus Grid Scaling Error . . . 156
5.18 Calibration Projection Error Versus IRED Tracking Error . . . 157
5.19 Example Endoscope Motion Constraints . . . 158
5.20 Tracking Error with Constraints . . . 159
5.21 Histogram of Average Back Projection Errors . . . 161
5.22 Endoscope Tracking Error Overlay . . . 161
5.23 Total System Error Overlay . . . 164

6.1 Ultrasound Acquisition . . . 168
6.2 Sample Range for Patient 02 . . . 170
6.3 Optotrak Camera View of the Tracking Collar in Theatre . . . 170
6.4 ILE for Each Frame in the Patient 02 Tracking Data . . . 172
6.5 Patient 02 Endoscope Tracking Error Overlay . . . 173
6.6 Point Picking of Inner Pubic Arch Points . . . 174
6.7 Direct Alignment of the Ultrasound Image to the Patient . . . 174
6.8 Alignment at Different Frames . . . 175
6.9 Effect of Image Interpolation . . . 177
6.10 Effect of Image Interpolation Method . . . 177
6.11 Effect of Overlay Opacity . . . 178
6.12 Slice Projection Orientation . . . 178
6.13 Slice Selection . . . 179
6.14 Motion Tracking . . . 180

7.1 Clinical Utility Versus System Error . . . 187
7.2 System Feasibility Versus System Error . . . 189

B.1 Component coefficients for Cadaver and Phantom data . . . 242

C.1 Data Acquisition from the Workstation Clock . . . 246
C.2 Histogram of Recorded Frame Lag . . . 246
C.3 Timing Correlation . . . 247

D.1 Error in Calculated TRE due to prostate position . . . 250

List of Tables

1.1 A summary of the data used in this thesis. 29
1.2 Camera Intrinsic Parameters 31

2.1 Treatment Options for Localised Prostate Cancer 36
2.2 Parameters for Measuring the Success of a Prostatectomy 36
2.3 Core transforms in literature review 46
2.4 Published Bone Segmentation from MRI 53
2.5 Published Pelvis Shape Models 53
2.6 Selected results for published ultrasound calibrations 60
2.7 Published tracking accuracy for NDI Optotrak 62
2.8 Published Registrations using Ultrasound 64

3.1 Differential Evolution Parameters 74
3.2 Differential Evolution Strategies 76
3.3 Differential Evolution Constraint Strategies 77
3.4 Model on MRI for 12 and 20 Modes 82
3.5 Average model segmentation boundary errors 86
3.6 Differential Evolution Strategies Tested 88
3.7 Optimal set of parameters used by DE Algorithm 90
3.8 Affine alignment of MRI to CT data 92
3.9 Average Boundary Errors 93
3.10 Predicted TRE for MRI data 94
3.11 Comparison of Shape Model Results 97
3.12 Correlation Coefficient and Average Boundary Error 97
3.13 MRI Segmentation Result Summary 99

4.1 Registration Parameters 118
4.2 Distribution of Slices in each Data Set 129
4.3 Registration parameters for validation experiment 131
4.4 Experiment Results 131

5.1 Endoscope intrinsic parameters 146

5.2 ILE as Normal Distributions 160
5.3 On Screen Error for Different Error Models 162
5.4 Summary of System Errors 164

6.1 MRI Data Used for In Theatre Overlay 166
6.2 The Quantity of Ultrasound Images Collected 167
6.3 The Quantity of Endoscope Video Frames Collected 169
6.4 Expected Values of Tracking Errors for Patient 02 171
6.5 Clinical Data for the five Patients 181

7.1 Summary of System Errors 187

B.1 Cadaver Coefficients Test 243

Chapter 1

Introduction

1.1 Aim

Image guided surgery has long been proposed as a method to improve patient outcomes for various surgical procedures. Systems are in place for surgery where the anatomy is largely rigid, for example neurosurgery and orthopaedics. Image guided surgery enables the surgeon to refer to preoperative images of the patient in an intuitive way during surgery. Using the example of neurosurgery, a preoperative image of a sub surface tumour may be projected onto the surface of the brain. In recent years there has been a growing focus on applying image guided surgery to thoracic and abdominal surgery. This thesis is based on an attempt to design and implement an image guided surgical system for robot assisted radical prostatectomy (RARP). The work has two key goals.

1. Develop and test in theatre an image guidance system for RARP.
2. Determine the accuracy that the system can achieve.

Successful completion of the first goal will give an enhanced understanding of the clinical requirements. Successful completion of the second goal will give an understanding of the system's capability. Combining the two outputs will allow us to answer the question of whether the system is usable in theatre and what areas to target for improvement.

1.2 Contributions

The thesis makes several important contributions to the field of study.

1.2.1 Error Analysis

Understanding of the accuracy of an image guided surgery system is necessary so the surgeon can make informed decisions about the reliability of the information presented. The main thrust of this thesis has been not only the development of such a system but furthermore a robust analysis of its errors. This has shown that at present the accuracy is limited by the accuracy with which the endoscope can be tracked.

1.2.2 Ultrasound to MRI Bone Registration

To enable registration of the preoperative magnetic resonance imaging (MRI) image to the patient in theatre we developed a novel algorithm to register two images of the pelvic bone. The first image is the MRI collected preoperatively. The second image is an ultrasound image collected intra operatively. The method combines a novel MRI to computed tomography (CT) registration algorithm, see next point, and an existing ultrasound to CT registration algorithm. The combined algorithm forms a novel and interesting approach to the problem of registering preoperative images to the intraoperative scene.

1.2.3 MRI to CT Registration

As mentioned in the previous point we developed a novel algorithm for fitting inter patient CT data to a MRI image. This algorithm is functionally equivalent to a MRI bone segmentation algorithm. Bone segmentation from MRI is of great interest for orthopaedic studies. The proposed algorithm compares well with the state of the art in this area. The algorithm may also prove useful in allowing radiotherapy planning using MRI images instead of CT.

1.2.4 Tracking Error Analysis

We propose tracking the endoscope with a marker based tracking system. Such systems are well established in the literature and used in commercial image guided surgery systems. Therefore the study of the accuracy of these systems is of great importance to the proper understanding of the accuracy of image guided surgery systems. A common assumption used in the literature is that the error at each marker is independent, homogeneous, and normally distributed. We found that these assumptions do not hold in our case. Therefore we proposed and tested a novel method to correctly model the error. We show that the commonly help error assumptions significantly underestimate the tracking error in our case.

1.2.5 Clinical Use

The system has been used in theatre on five occasions. We have shown that it is a practical system. At present its utility is limited by a high endoscope tracking error. Nonetheless the feed back from the surgeon has been positive.

1.3 Summary

This thesis describes the development, analysis and implementation of an image guided surgery platform for RARP. The system implemented takes voxels (X_M, in 3D) of preoperative MRI data and projects them, using a series of transformations, to points on the endoscope screen (X_{ES}, in 2D) so that they appear to be coincident with the corresponding points of the visible patient anatomy. At present the system is only implemented on one of the *daVinci*^{TM} video channels, so is monocular rather than stereoscopic. Figure 1.1 shows an example of the output of our overlay program, as seen by the surgeon intraoperatively. A transverse MRI plane is shown overlaid on a typical surgical scene. The surgeon has the option of changing the opacity, slice, or slice direction of the overlay. The interface and display have been kept very simple. This thesis is primarily concerned with describing the methods used to define the transforms and robustly determining the errors at each stage. Whilst the implementation is specific to RARP using a *daVinci*^{TM} telemanipulator we envisage that many of the components can be applied to other guidance applications.

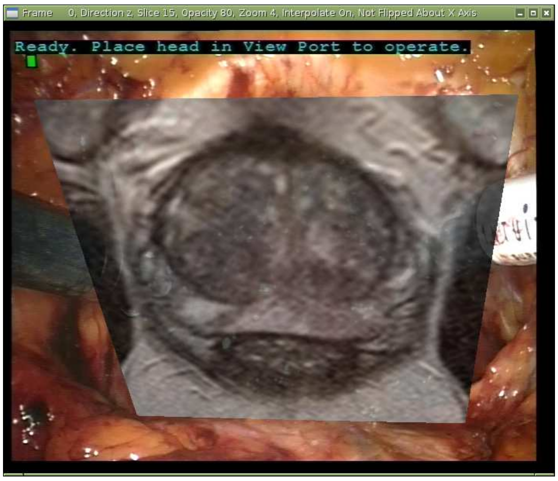

```
Frame     0, Direction z, Slice 15, Opacity 80, Zoom 4, Interpolate On, Not Flipped About X Axis
Ready. Place head in View Port to operate.
```

Figure 1.1: An example of a transverse slice of MRI overlaid on a typical endoscopic scene. The overlay software projects the MRI using a model of the endoscope so that the MRI appears to be coincident with the visible anatomy. The projection software allows the surgeon to change the opacity of the projection or move through the MRI data in the transverse, sagittal, or coronal planes.

1.4 System Description and Structure of the Thesis

The system uses a novel MRI to ultrasound registration algorithm to transform the preoperative MRI (X_M) into the tracker's coordinate system (X_O) in theatre, using equation 1.1.

$$X_O = T_{M \Rightarrow O} \times X_M \tag{1.1}$$

With the MRI placed in the tracking systems coordinate system, projection of the MRI can be achieved by equation 1.2.

$$X_{ES} = T_{O \Rightarrow ES} \times X_O \tag{1.2}$$

The aims of this thesis are twofold. Firstly it presents a method to determine the system wide transform, equation 1.3. Secondly it presents an analysis of the accuracy of the system based on analyses of the system components.

$$T_{M \Rightarrow ES} = T_{O \Rightarrow ES} \times T_{M \Rightarrow O} \tag{1.3}$$

The system consists of the following steps.

1. Patient is selected and consented for image guided surgery.

2. Preoperative MRI scans are collected.

3. Pelvic bone is segmented automatically from scans.

4. Endoscope cameras are calibrated prior to surgery.

5. Tracked ultrasound is used to find the pelvic bone in theatre.

6. Registration algorithm matches ultrasound data to segmented pelvic bone.

7. Endoscope is tracked during surgery.
8. Video stream from *daVinci*^{TM} is captured and fed to monitor where the preoperative model is overlaid.

For the purposes of development and analysis the system was broken into a set of subsystems. Each of these systems is described and analysed in a separate chapter of this thesis. Throughout the thesis we treat the image guidance system as a series of transformations from one coordinate system to another, with the complete transform $T_{M \Rightarrow ES}$ being made up of a series of transforms. The precise details of how these transforms are determined, and the error in their determination, forms the basis of this thesis and will be discussed in detail in the relevant chapters. In the following pages the transforms will be introduced and briefly explained. This forms both an introduction and a reference to which we will refer back to in later chapters.

In Chapter 4 we present a method for registering an ultrasound image of the patient's pelvic bone to the patient's MRI image. As the MRI image does not show the bone edge clearly we developed an algorithm to transform the MRI prior to registration. This algorithm is described in Chapter 3. The algorithm can be described as an inter-patient, inter-modality, non rigid registration. The algorithm takes a CT image from a different patient and warps it match the patient's MRI image. Figure 1.2 shows the transformations used.

The patient's pelvic bone can be imaged in theatre using a tracked ultrasound probe. The ultrasound image of the pelvic bone can be registered to the preoperative bone, using methods broadly similar to Penney et al. (2006). In Chapter 4 a new implementation of this method is described. The key difference to the work of Penney et al. (2006) is that the ultrasound is registered to a CT image from a different patient. The inter patient CT is first warped (using $T_{SM \Rightarrow M}$) to approximate the shape of the patient's pelvic bone. Figure 1.3 defines the transforms determined in Chapter 4. At the conclusion of Chapter 4 a transform $T_{M \Rightarrow O}$ (6 degrees of freedom in 3D) has been found that transforms the preoperative MRI image into the coordinate system of a tracking system (an Optotrak Certus). The next stage is to calibrate and track the endoscope to find $T_{O \Rightarrow ES}$, completing the image guidance system.

Chapter 5 details our methods to determine $T_{O \Rightarrow ES}$. Figure 1.4 details the transformations found in Chapter 5. The proposed method is not novel, using an existing camera calibration method to determine the projection parameters of the endoscope and a marker based tracking method to estimate the location of the endoscope. However, we do present a novel method to analyse the tracking error, which used more realistic marker error distributions, to yield a result significantly closer to observed values for this case. We show that the endoscope tracking error is, as a result, the largest error for the system and propose methods to reduce the error.

Development of the methods presented in Chapters 3 to 5 is ongoing and has yet to yield a practical image guidance system. In the mean time we have implemented a "bare bones" guidance system in theatre, in order to assess some of the other factors involved in an image guidance system. These factors include how the data is presented to the surgeon and how the surgeon interacts with the display, and assessing what anatomy is of importance to the surgeon during the operation. This image guidance

Chapter 3 Segmentation of Bone from MRI

$T_{SM \Rightarrow M} = T_{SM(nd9) \Rightarrow M}[nrr] \times T_{SM \Rightarrow M}[nd9]$

SM	Generic Shape Model
M	Patient MRI
SM(nd9)	Shape Model Scaled and Aligned to MRI
SM(nrr)	Shape Model Warped to MRI

$CS_{SM(nrr)} = CS_{SM} \times T_{SM \Rightarrow M}$

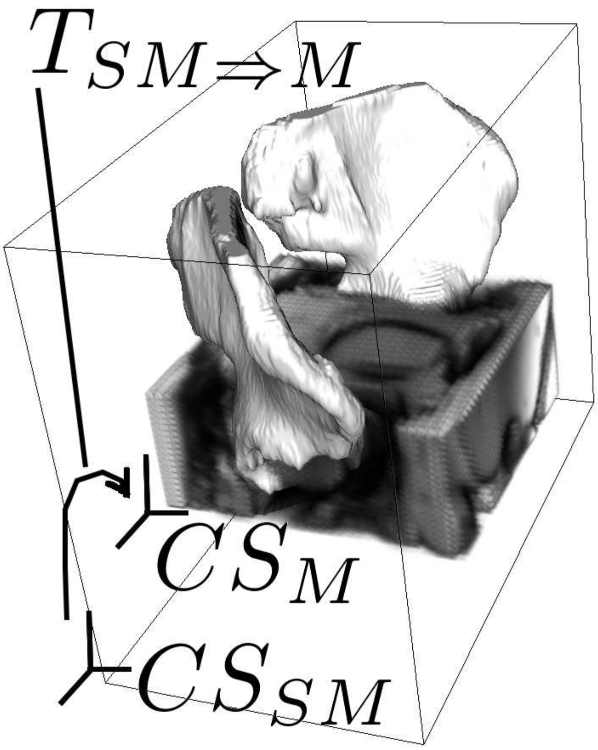

Figure 1.2: The transform $T_{SM \Rightarrow M}$ transforms voxels from the Coordinate System (CS) of a separate patient's CT image (effectively the mean shape of a shape model) to the coordinate system of the prospective RARP patient. A two stage process is used. The two images are first brought into alignment using a 9 degree of freedom transform ($T_{SM \Rightarrow M}[nd9]$, number of degrees of freedom = 9). This transform consists of rotation, translation, and scaling in three directions. With the two images thus aligned a non rigid transformation ($T_{SM(nd9) \Rightarrow M}[nrr]$, non rigid registration) warps the CT image to the MRI.

Chapter 4 Registration of Ultrasound to Fitted Bone Image	
$T_{UI\Rightarrow O} =$ $T_{UT\Rightarrow O} \times T_{UP\Rightarrow UT} \times T_{UI\Rightarrow UP}$ $T_{SM(nrr)\Rightarrow O} =$ $T_{UI\Rightarrow O} \times T_{SM(nrr)\Rightarrow UI}$ $T_{M\Rightarrow O} = T_{SM(nrr)\Rightarrow O}$ UI Ultrasound Image (2D) O Optical Tracker UT Ultrasound Tracking Collar UP Ultrasound Probe $CS_O = CS_M \times T_{M\Rightarrow O}$	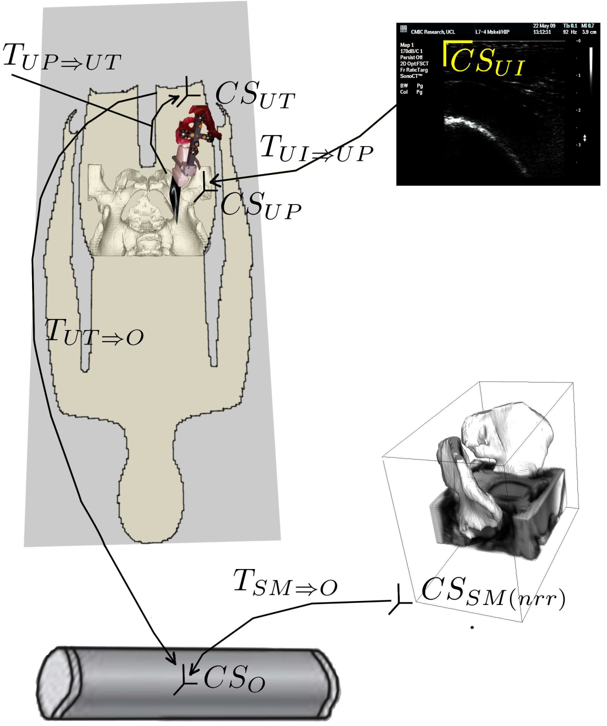 $T_{UP\Rightarrow UT}$ CS_{UT} $T_{UI\Rightarrow UP}$ CS_{UP} CS_{UI} $T_{UT\Rightarrow O}$ $T_{SM\Rightarrow O}$ $CS_{SM(nrr)}$ CS_O

Figure 1.3: As will be described in Chapter 4 a 3 dimensional ultrasound image of the patient's bone is first built from set of 2D slice images acquired using a tracked ultrasound probe. Pixels in the Coordinate System (CS) of the ultrasound slice image (points in 2D) are transformed to 3D points in the coordinate system of the ultrasound probe using $T_{UI\Rightarrow UP}$. These 3D points are transformed to 3D points in the coordinate system of a tracking collar attached to the ultrasound probe using $T_{UP\Rightarrow UT}$, a rigid body transform. Finally $T_{UT\Rightarrow O}$ transforms these points to points in the coordinate system of the Optotrak tracking system used in theatre. This ultrasound image is registered to the preoperative model found in Chapter 3 to give $T_{SM(nrr)\Rightarrow UI}$ and hence $T_{SM(nrr)\Rightarrow O}$. As the model is fitted to the MRI data $T_{M\Rightarrow O}$ is the same as $T_{SM(nrr)\Rightarrow O}$.

<table>
<tr><td colspan="2">Chapter 5 Endoscope Calibration and Tracking using an Optical Tracker</td></tr>
<tr><td>

$T_{O\Rightarrow ES} =$

$T_{EL\Rightarrow ES} \times T_{ET\Rightarrow EL} \times T_{O\Rightarrow ET}$

<table>
<tr><td>ES</td><td>Endoscope Screen (2D)</td></tr>
<tr><td>ET</td><td>Endoscope Tracker</td></tr>
<tr><td>EL</td><td>Endoscope Lens</td></tr>
</table>

$CS_{ES} = CS_{ES} \times T_{O\Rightarrow ES}$

</td><td>

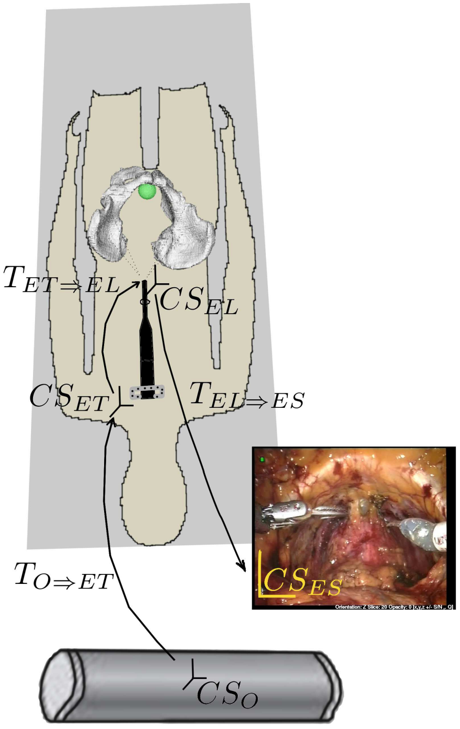

</td></tr>
</table>

Figure 1.4: In Chapter 5 the transformation ($T_{O\Rightarrow ES}$) between 3D points in the coordinate system (CS) of the tracking system (Optotrak) and 2D points on the endoscope screen is found. The "intrinsic" projection parameters of the endoscope define the transform ($T_{EL\Rightarrow ES}$) between 3D points relative to the endoscope lens and 2D points on the endoscope screen. The "extrinsic" parameters of the endoscope define the rigid (6 degrees of freedom) transformation ($T_{ET\Rightarrow EL}$) from the coordinate system of a tracking collar attached to the endoscope to the coordinate system of the endoscope lens. Both $T_{EL\Rightarrow ES}$ and $T_{ET\Rightarrow EL}$ are found using a camera calibration algorithm. The Optotrak tracking system estimates the position and pose of the tracking collar to give $T_{O\Rightarrow ET}$.

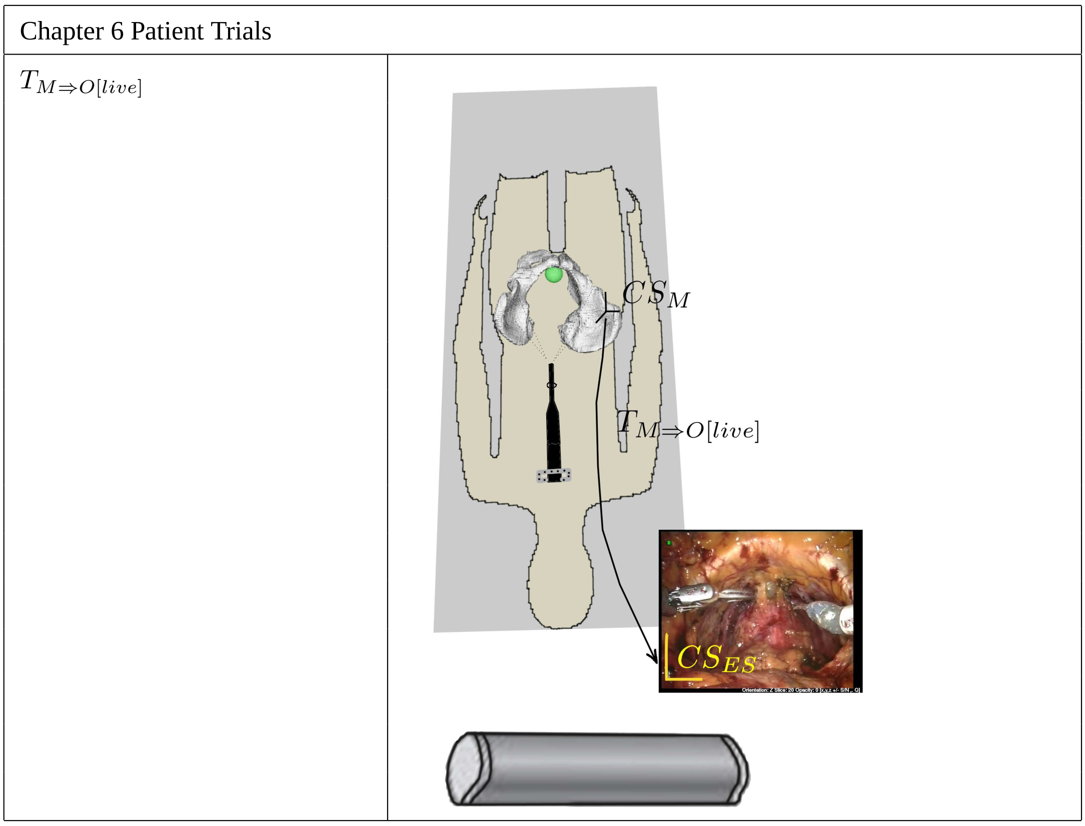

Figure 1.5: Chapter 6 presents a method we used to implement a simple image guidance system rapidly. In this system $T_{M\Rightarrow O[live]}$ is found directly using the inner surface of the pubic arch, which is visible both in the MRI image and in some of the endoscopic video frames. A manual alignment algorithm is used to find a revised estimate of $T_{M\Rightarrow O}$, referred to as $T_{M\Rightarrow O[live]}$. The manual alignment algorithm allows each degree of freedom to be incremented until the two images (MRI and endoscopic) of the pubic arch are overlaid.

system is described in Chapter 6. The system uses a direct manual registration between the MRI image and the endoscope image using surfaces visible in both images, see Figure 1.5 Chapter 6 also presents our experience in implementing the system in theatre on a total of 5 patients.

Preceding these chapters is a review of the existing literature, which has three aims. Firstly we introduce radical prostatectomy and specifically RARP. We then introduce image guided surgery. Image guided surgery can refer to a very great range of systems, so we define what we mean by image guided surgery for this application. We then define a generic endoscopic image guidance system consisting of a set of interconnected components. We show that the transformations shown in Figures 1.2 to 1.5 are core to any such system. Finally we look in depth at the state of the art for determining each of the transformations required for our system as defined in Figures 1.2 to 1.5.

Data Set	Number of Data Sets	Image Modalities	Chapter(s)
SM-X	21	CT	3 & 4
Patient-XX	5	MRI, ultrasound, video, & tracking	3 & 6
XMR-XX	2	MRI	3
Cadaver-XX	3	CT & ultrasound	4
Phantom	1	CT & ultrasound	4
Camera Calibration	1	Video & tracking	5

Table 1.1: A summary of the data used in this thesis.

1.5 Data

Each of the experiments in this thesis rely on data. A total of 32 data sets are used in this thesis. Full details of each data set are contained in Appendix A. The following is a brief introduction to the data sets and how and where they are used.

21 CT images of adult male pelvis are used in Chapter 3 to construct and test a statistical shape model (SSM) of the adult male pelvis. These images are designated SM-X, where X is a letter from A to U. The aim of the SSM is to perform bone segmentation from MRI images, so four data sets containing MRI images were used to test the algorithm. Two of these were taken from a previous study at Guy's Hospital, using co-registered CT and MRI images. These data sets are referred to as XMR-01 and XMR-02. Only the MRI data was used. MRI scans from two prospective prostatectomy patients were also used, referred to as Patient-01 and Patient-02.

For Chapter 4, it was necessary to have data sets containing CT and ultrasound images of the pelvic bone, together with the transformations between the two. For algorithm development and validation a custom plastic anatomy phantom was built. The phantom consists of a life size plastic pelvic bone, a spherical target representing the prostate, and eight fiducial markers. The phantom was imaged in both a CT scanner and with an ultrasound probe in a water bath. The gold standard transform between the CT and ultrasound images was established using the eight fiducial markers. The phantom data set is referred to by the prefix Phantom. The phantom is further described in Section 4.5.

As neither the CT nor ultrasound images of the phantom were representative of actual anatomical images the algorithms were validated using CT and ultrasound images from cadavers. Three cadaver data sets were used, denoted Cadaver-XX, where XX=01,02,03. The images were taken from a previous orthopaedic study. The gold standard transformation from ultrasound to CT images was determined using a set of bone implanted fiducial markers.

Chapter 5 uses tracking data collected during an in theatre calibration of the *daVinci*^{TM} endoscope, together with images of the calibration grid.

Finally, in Chapter 6 data from 3 more prostatectomy patients are introduced. These are denoted Patient-03, Patient-04, and Patient-05. These data sets contain MRI and ultrasound images, as well as endoscope tracking data. Table 1.1 summarises the data sets used in this thesis.

1.6 Measuring System Error and Summing Component Errors

As well as detailing the methods used by our surgical navigation system this report is largely concerned with the quantification and analysis of the systems' errors. It will be useful therefore to define the error measures used at the outset. We have used two measures of error in the report. The first is the registration error at the prostate. This is a measure of how far the preoperative image of the prostate is from the actual position of the patient's prostate. We use this is Chapters 3 and 4. The second measure is an on screen error visible to the surgeon. This is intended to give the surgeon an understanding of the system accuracy. This is used at the conclusion of each chapter to visualise the errors found in the chapter. We will now define how these errors are calculated.

1.6.1 Registration Error at Prostate

This is used when calculating the errors in the calculation of $T_{M \Rightarrow O}$. The registration error can be quantified by first defining a point in space, transforming it by the known correct $T_{M \Rightarrow O}(GldStd)$ [1] and the estimated $T_{M \Rightarrow O}(Est.)$. The registration error is simply the magnitude of the difference (in three dimensions) between the two transformed points. As we are interested primarily in the position of points near the prostate we have used such points in our error measure. Similar error measures are widely used in the literature and can be referred to as a target registration error (TRE). To allow direct comparison between the various data sets used in the work the same set of six points in space were used for this error measure. We refer to these as the "nominal" prostate surface points. The location of this was based on the position of the prostate in the plastic phantom used in much of the development work. Figure 1.6 shows the six points used. Doing this simplifies the calculation of errors and should be reasonably accurate as the position of the prostate does not vary significantly. The effects of this assumption are discussed further in Appendix D.

1.6.2 On Screen Error

From the surgeon's point of view the error measure of interest is the apparent error in the position of the prostate shown on the endoscope screen. Therefore each error source is ultimately converted to an on screen error. To do this in a consistent manner we first defined a point of interest in the preoperative MRI. We chose the apex of the prostate, as this is a landmark that can be seen through the endoscope and is clinically relevant. We then selected a frame of endoscopic video that shows the entire prostate and in our clinical experience was a good frame to overlay. The choice of this frame placed the apex of the prostate approximately 200 mm away from the lens, 20 mm above the horizontal centre line and 5 mm to the right of centre. The point is then projected onto the endoscope screen using intrinsic parameters found during a calibration of the endoscope, and detailed in Table 1.2. To visualise the errors, various errors are simulated in the projection process and the process repeated many times to give a distribution of on screen points. 1.7 gives an example of such an overlay.

---

[1] The gold standard transformation $T_{M \Rightarrow O}$ is the combination of $T_{SM \Rightarrow M}(GldStd)$ and $T_{M \Rightarrow O[nrr]}(GldStd)$. Model fitting errors are determined in Chapter 3 using simulation, so the gold standard transform is known beforehand. In Chapter 4 the ultrasound tracking and registration errors are determined experimentally using a plastic phantom and cadaver data. In both cases the gold standard transforms are determined using sets of fiducial markers.

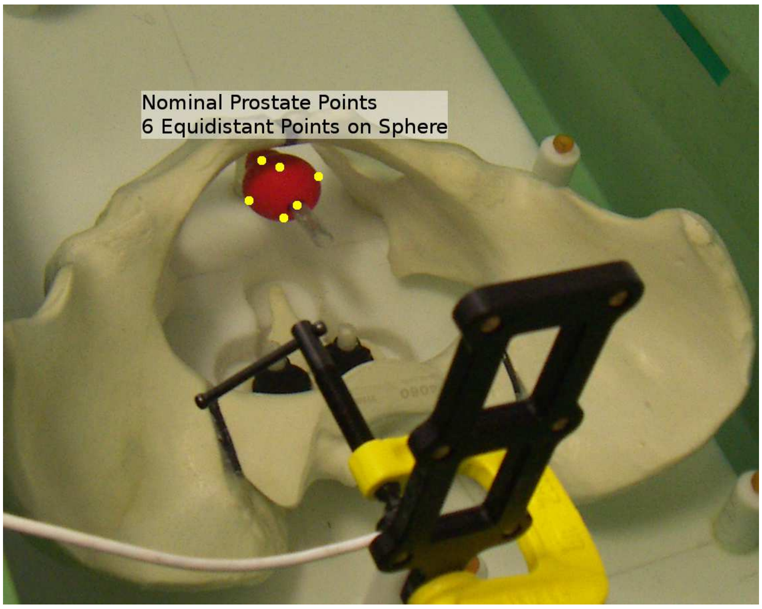

Figure 1.6: Definition of the six nominal prostate surface points used throughout to quantify the component errors.

Name	Horizontal	Vertical
Focal Length	852.02	923.14
Principal Point	415.38	288.12
Second Order Radial Distortion	-0.36	
Forth Order Radial Distortion	1.16	
Tangential Distortion	0.0085	0.0082
Sixth Order Radial Distortion	0.00	
Screen Dimensions (pixels)	720	576

Table 1.2: The Intrinsic parameters of the endoscope used for error visualisation. The endoscope is modelled as a pin hole camera with radial and tangential distortion as per Heikkila and Silven (1997).

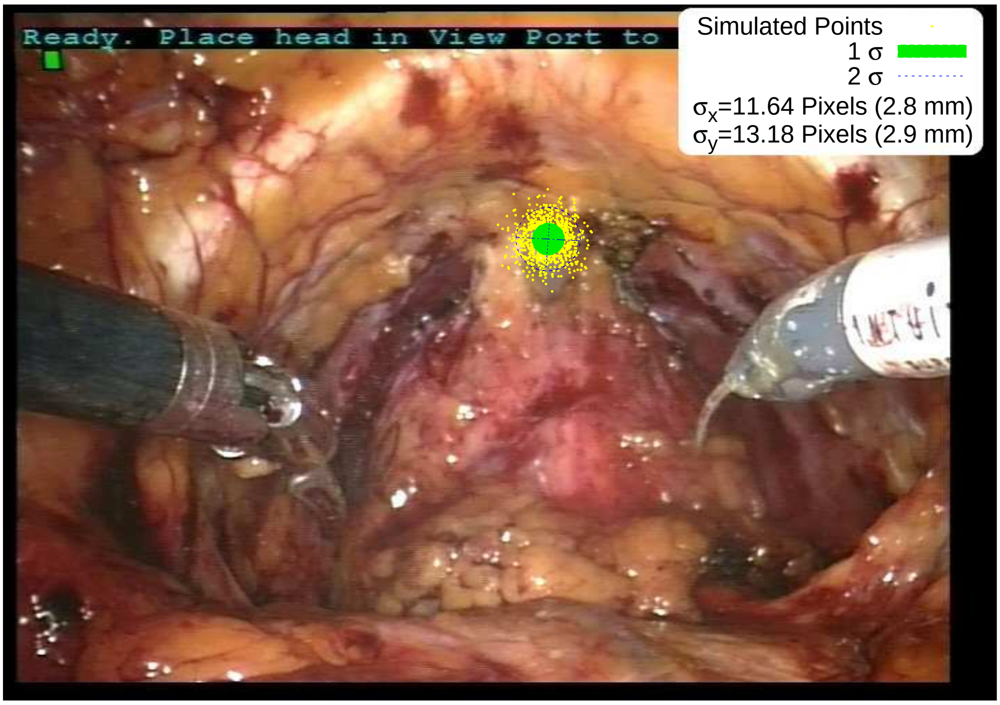
Ready. Place head in View Port to
Simulated Points
1 σ
2 σ
σx=11.64 Pixels (2.8 mm)
σy=13.18 Pixels (2.9 mm)

Figure 1.7: A landmark point (near the apex of the prostate) has been chosen and projected onto the endoscope screen under the influence of errors. The yellow points represent the projected points. The standard deviation of the point spread in the x and y directions is calculated and shown in the text box at the top right. Two ellipses are drawn on the image showing 1 and 2 standard deviations. The principal axes of the point distribution are also shown. To aid interpretation a value in mm has also been calculated. This is the pixel error back projected onto a plane 200 mm from the endoscope lens.

Whilst we recognise that the resulting error values expressed in pixels are very specific to our application and endoscope geometry making comparison with other work more difficult, they are none the less the most relevant to our application. To aid interpretation an error value in mm is also shown. This is the on screen error in pixels back projected onto a plane 200 mm distant from the endoscope lens (near the prostate apex). Measuring the error in different ways is of course possible and where appropriate we have done this also.

1.7 Ethics

Ethical approval for all studies involving prospective RARP patients was obtained from the Research Ethics Committee at Guy's Hospital (Reference Number 08/H0804/1). All prospective patients gave informed consent for the use of their MRI and ultrasound data, as well as the use of the tracking system during surgery.

Chapter 2

Literature Review

This chapter is structured as follows. We first give a brief introduction to prostate cancer and its treatment using RARP. Following this we introduce the key concepts of image guided surgery and explain where our system sits in this very broad field. Having defined an image guidance system as a connected group of subsystems, we then examine the state of the art for the subsystems.

2.1 Prostate Cancer and Radical Prostatectomy

In the UK prostate cancer is the most commonly diagnosed cancer in men, [Office of National Statistics (2008)]. It is also the second most common cause of cancer related death (after lung cancer) of men in both the UK and USA. Over 10,000 men die from Prostate cancer each in the United Kingdom, [Kirby et al. (2010)]. Around three quarters of men diagnosed will survive the disease beyond five years. Many more men will have prostate cancer but remain undiagnosed. Incidence will likely increase as screening improves [Kirby et al. (2010)]. Autopsy studies, [Breslow et al. (1977)] have found that up to 80% of men in their seventies have prostate cancer. Understanding the causes of and improving the treatment of prostate cancer is therefore an important challenge,

There are a wide range of treatments available for prostate cancer, these are listed in Table 2.1. Correct management of the symptoms is also of critical importance [Thompson et al. (2007)]. The choice of treatment will depend on the stage of cancer (whether it is confined to the prostate gland, or has spread beyond the prostate), the aggressiveness of the cancer, the life expectancy of the patient and the patient's wishes. Radical prostatectomy may be used when the cancer has not yet spread beyond the prostate gland, but is likely to do so if left untreated, and the patient is healthy enough to undergo surgery. The spread of the cancer is typically determined with an MRI scan. The aggressiveness of the cancer is assessed using a prostate biopsy and quantified using the Gleason score [Gleason (1977)]. Excision biopsies taken from nearby lymph nodes can also be use to assess the aggressiveness of the cancer. There are several approaches to radical prostatectomy. Radical prostatectomy removes the cancerous organ and the attached seminal vesicles. Biopsies can also be taken from the nearby lymph nodes allowing accurate assessment of the spread of the cancer. The nerve bundles that control potency run very close to the prostate. Depending on the location of the cancer within the prostate it is possible to attempt nerve sparing prostatectomy so that the patient may remain potent. Alternatively a non nerve sparing procedure

Established Techniques	
Active Surveillance	No treatment is undertaken, however the progress of the cancer is monitored.
Radical prostatectomy	The prostate is surgically removed.
External beam radiation therapy	Radiation is focused on the prostate from an external source.
Brachytherapy	Radioactive seeds are inserted into the prostate through the perineum. Cancerous cells are destroyed by radiation.
Experimental Techniques	
High frequency focused ultrasound	Cancerous cells are ablated using sound energy from a trans-rectal probe
Cryogenic ablation	Cooling fluid is passed through a needle to cool a target region. The cold temperatures destroy cancerous tissue.

Table 2.1: Some of the options for the treatment of localised prostate cancer. The option used will depend on the stage and aggressiveness of the cancer as well as the age and health of the patient.

Oncological outcomes (positive margin rate and biochemical PSA recurrence)
Time to and rate of urinary continence (Months, %)
Time to and rate of potency (if nerve sparing) (Months, %)
Unintended damage to surround anatomy (principally the rectum).
Length of hospital stay (Days)
Patient post operative pain
Time under general anaesthetic (minutes)
Cost of Surgery (£)

Table 2.2: Parameters that can be used to assess the success of a prostatectomy, allowing different methods to be compared. To be successful any alteration to the surgical method needs to demonstrably improve one of these.

can be performed. The success of a radical prostatectomy can be measured by a number of parameters. Table 2.2 lists parameters that can be used to measure the success or otherwise of a prostatectomy and thus compare alternative surgical approaches and assess the effectiveness of a proposed new method.

Traditionally, open surgery has been used, and this can either be done with a retro pubic approach (incision through the abdomen) or via an incision through the perineum. More recently, first reported in 1997 (Varkarakis et al. (2005)) laparoscopic approaches have been used. These reduce the patient's hospital stay and post operative pain. In the hands of a skilled laparoscopic surgeon laparoscopic surgery can have the same outcomes for positive margins, and time to continence/potency. The cost of surgery is higher due to more expensive consumables and the requirement for higher skilled surgeons. There is also a long history of robotics in urology (Challacombe et al. (2006)) as surgeons and scientists have strived to improve the available methods.

A more recent (since 2000, Takenaka et al. (2006)) alternative to laparoscopic surgery is robot assisted laparoscopic surgery. Here the standard laparoscopic tools are replaced with tools actuated remotely by the surgeon via the *daVinci*^{TM} “robot” [Guthart and Salisbury (2000)]. By removing the surgeon from direct control of the laparoscopic tools this allows the control interface to be designed to mimic direct control of the surgical tools, as opposed to the reversed control required for standard laparoscopic surgery. In 2008 there were 70,000 prostatectomies performed with a *daVinci*^{TM} robot worldwide, and there are now more than 1000 *daVinci*^{TM} systems worldwide [Tan et al. (2009)].

There is active debate over the benefits of the *daVinci*^{TM} system, however there is evidence [Tonet et al. (2006)] that use of the *daVinci*^{TM} allows surgeons without laparoscopic experience to perform to minimally invasive procedures, whilst the *daVinci*^{TM} provides limited benefit to an already skilled laparoscopic surgeon. Most trials show patient outcomes improving as the surgeon becomes more adept at using the *daVinci*^{TM} [Dasgupta and Kirby (2009)]. With the exception of time to continence outcomes are at least as good open surgery [Dasgupta and Kirby (2009)].

There is a need for randomised trials to assess the performance of the various surgical approaches (Dasgupta et al. (2006)). Robot assisted laparoscopic surgery is however a recently developed procedure and to date the procedure largely mimics standard laparoscopic surgery. It seems likely that as more surgeons use the system and the product itself evolves that the robot assisted procedure will diverge from the standard laparoscopic procedure, taking advantage of the increased degrees of freedom available at the end effectors, the ability to scale movements, the filtering of tremors, and the 3D endoscopic vision available on the *daVinci*^{TM} .

2.1.1 Description of robot assisted radical prostatectomy.

The following describes the robot assisted radical prostatectomy procedure currently in use at Guy’s Hospital. The patient is first prepared for surgery. A catheter is inserted into the bladder, and six ports are cut in the patient’s abdomen. A central port for the *daVinci*^{TM} endoscope. A port on either side for the two arms of the *daVinci*^{TM} and an additional three ports for laparoscopic tools controlled by an assistant. The patient can be insufflated with carbon dioxide through one of the three assistant ports. After port cutting the patient is placed head down in the Lithotomy position. This position is necessary to allow the *daVinci*^{TM} to be placed between the patients legs and gain access to the patients abdomen. At this stage the both the patient and the *daVinci*^{TM} are locked in position and do not move until the procedure is completed. The endoscope is inserted into the central port and the *daVinci*^{TM} tools into the side ports. Figures 2.1 to 2.9 give a pictorial step by step description of the prostatectomy performed on one of the studied patients. The use of medical terminology has been avoided to try and keep it accessible to the layman. As an indication of the time taken for the procedure the time elapsed since the start of port cutting is shown at the bottom right of each image.

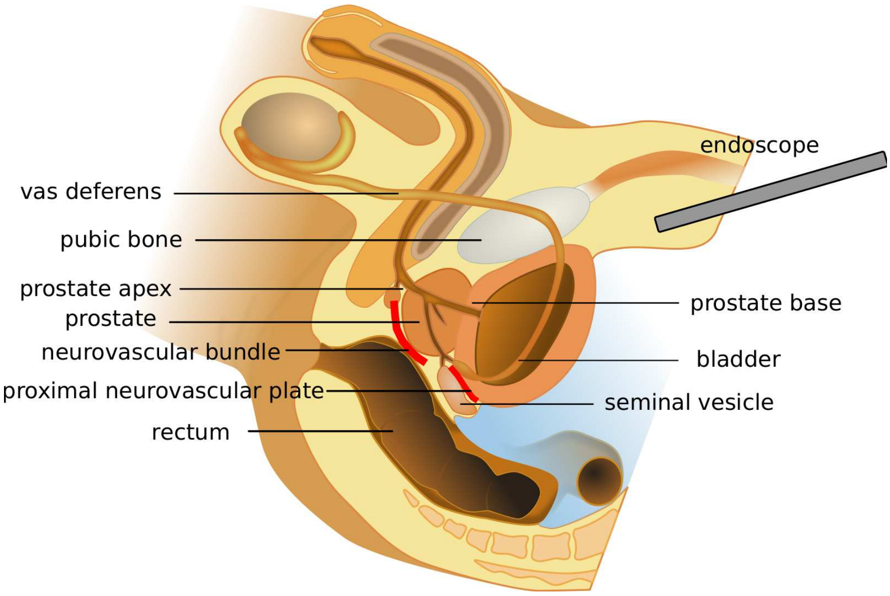
endoscope
vas deferens
pubic bone
prostate apex
prostate
neurovascular bundle
proximal neurovascular plate
rectum
prostate base
bladder
seminal vesicle

Figure 2.1: The prostate and surrounding anatomy as seen through the patient's mid sagittal plane [courtesy Tsaitgaist (2008)] with modifications. The nerve tissues, shown in red are not actually located on the mid plane but are slightly to the right and left of the mid plane. The endoscope and surgical tools are inserted through the abdominal wall into the retro-pubic cavity, to the right of the picture. The bladder is kept deflated during the first part of the procedure to aid access to the prostate.

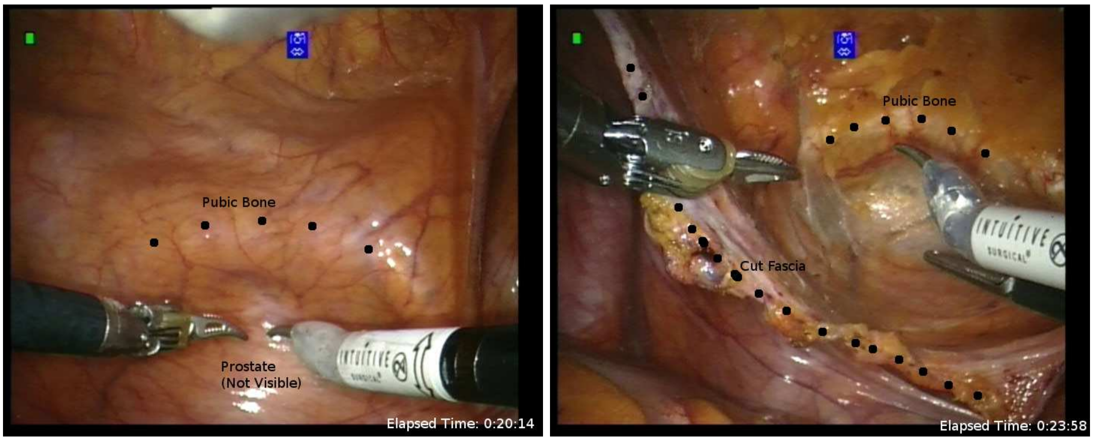
Pubic Bone
Prostate
(Not Visible)
Elapsed Time: 0:20:14
Pubic Bone
Cut Fascia
Elapsed Time: 0:23:58

Figure 2.2: On entry to the retropubic space the surgeon cannot see any relevant anatomy, though the position of the prostate can be estimated from the curvature of the pubic arch. The first stage is to remove a layer of fascia from the inner wall of the abdomen. This has been done on the right, allowing the pubic arch to be more clearly seen. The prostate, however, remains obscured.

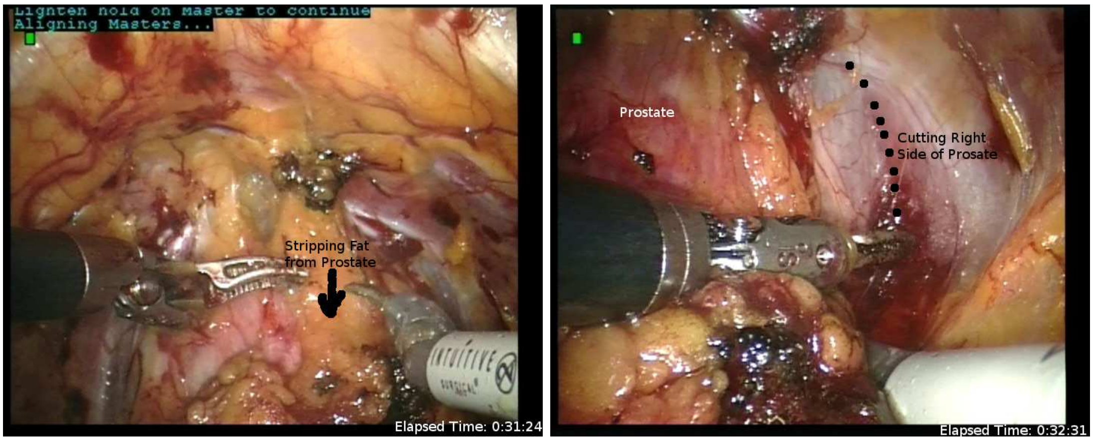
Lighten hold on Master to continue
Aligning Masters...
Stripping Fat
from Prostate
Elapsed Time: 0:31:24
Prostate
Cutting Right
Side of Prosate
Elapsed Time: 0:32:31

Figure 2.3: Tissue (membranes and fat) in front of the prostate is removed to reveal the prostate, on left. Incisions are made to along both sides of the prostate, right side incision shown on right. The extent of these incisions is limited to prevent damage to nerve fibres.

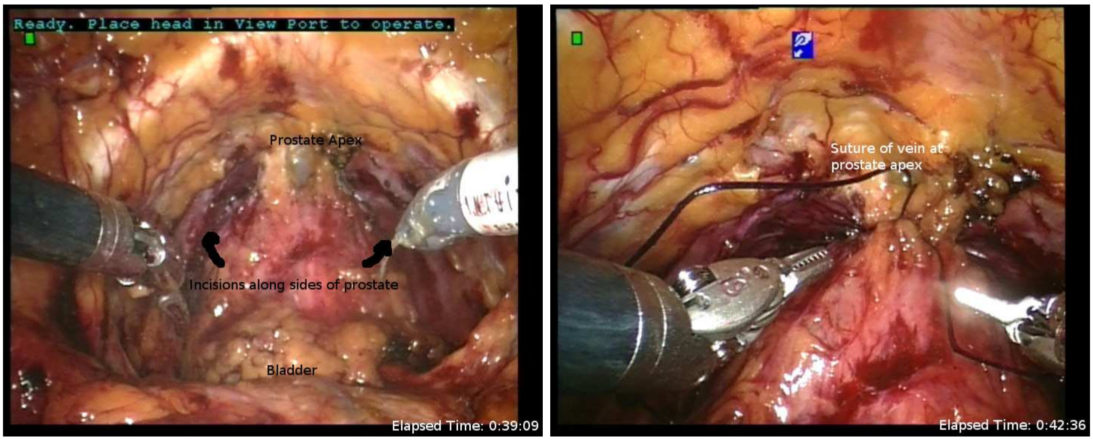
Ready. Place head in View Port to operate.
Prostate Apex
Incisions along sides of prostate
Bladder
Elapsed Time: 0:39:09
Suture of vein at
prostate apex
Elapsed Time: 0:42:36

Figure 2.4: On the left the prostate is now visible with a small incision on either side. The blood supply (the dorsal vein) at the apex of the prostate is then sutured (at right) to limit bleeding.

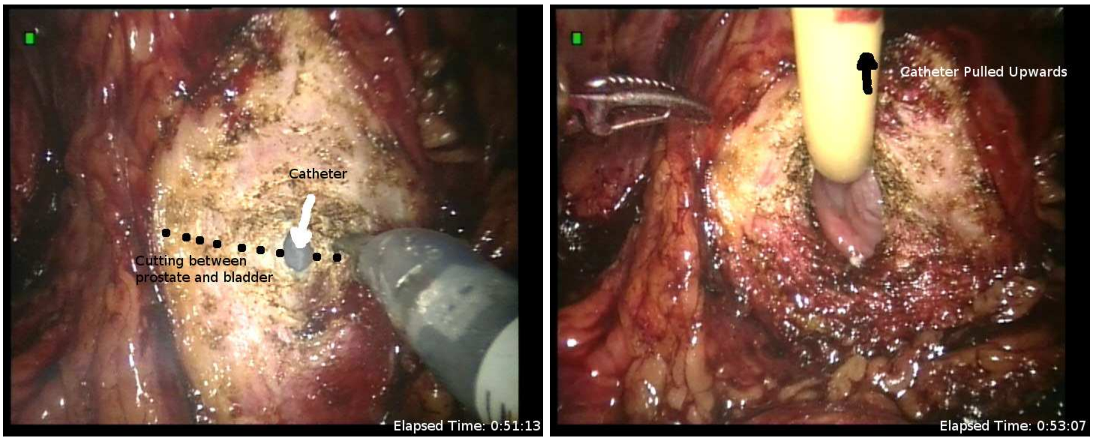
Catheter
Cutting between
prostate and bladder
Elapsed Time: 0:51:13
Catheter Pulled Upwards
Elapsed Time: 0:53:07

Figure 2.5: Dissection of the prostate from the neck of bladder is started. Knowledge of the size of the prostate's medial lobe is useful here as this can extend into the bladder. For good continence results it also important to minimise damage to the bladder neck. Once the dissection is sufficiently advanced the catheter can be pulled out of the urethra and pulled upwards, using an additional hook through the patient's abdomen, as shown on right. This pulls the prostate upwards allowing access to the far side.

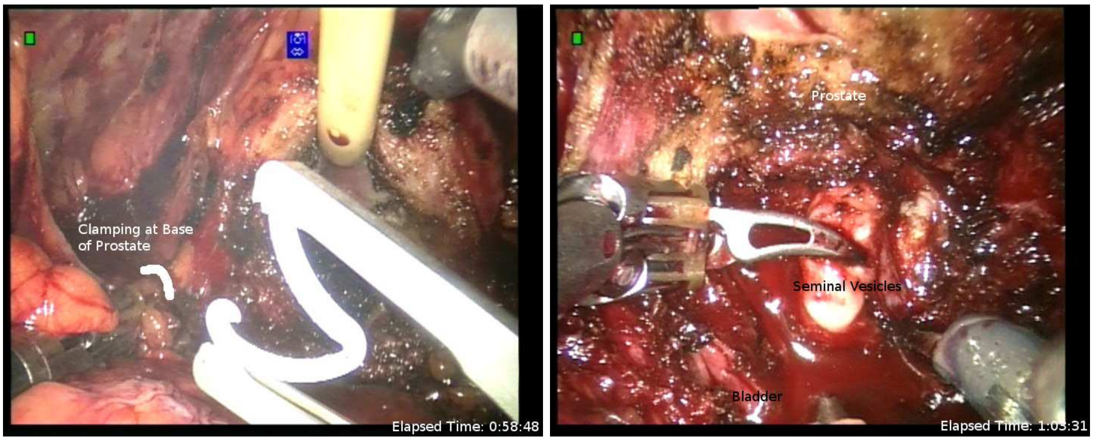
Clamping at Base
of Prostate
Elapsed Time: 0:58:48
Prostate
Seminal Vesicles
Bladder
Elapsed Time: 1:03:31

Figure 2.6: After dissection of the bladder neck is complete the base of the prostate can be freed. Clips rather than cauterisation are used to prevent blood loss due to the proximity to the neuro-vascular bundles. The left hand image shows a clip being applied to the left hand side of the prostate. After the base has been freed it is possible to reach under the prostate and pull the seminal vesicles and vas deferens forward.

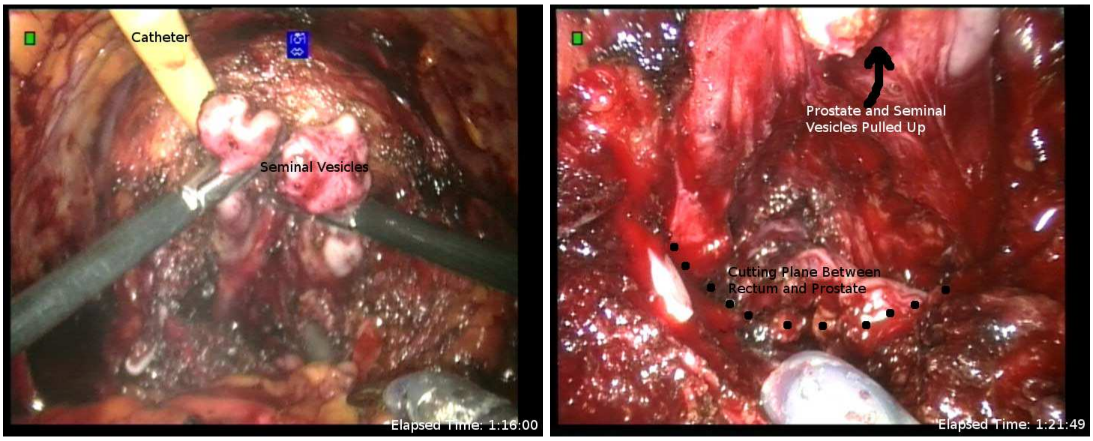
Catheter
Seminal Vesicles
Elapsed Time: 1:16:00
Prostate and Seminal Vesicles Pulled Up
Cutting Plane Between Rectum and Prostate
Elapsed Time: 1:21:49

Figure 2.7: With the seminal vesicles and prostate pulled upwards the rear face of the prostate can be dissected along the plane between the rectum and the prostatic capsule. The depth of this incision is dictated by the nerve sparing procedure in use. It is of critical importance not to cut into the rectum during this cut.

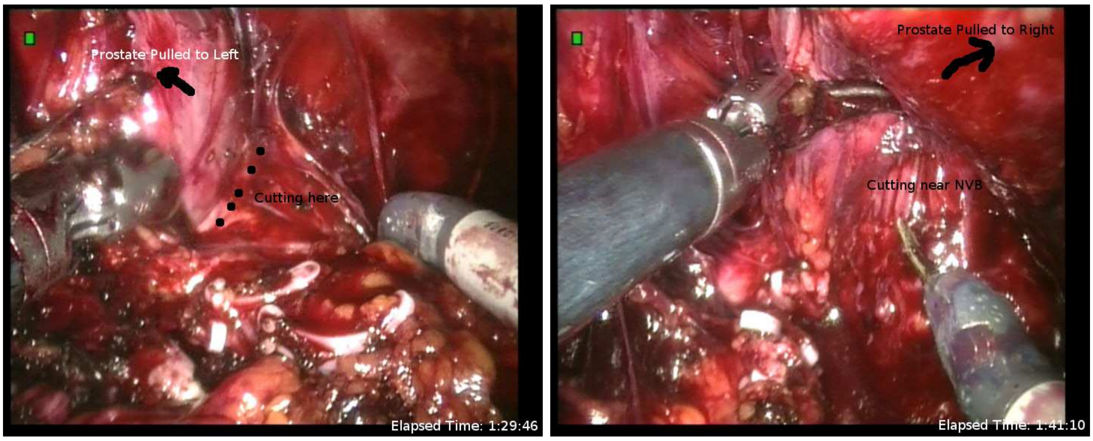
Prostate Pulled to Left
Cutting here
Elapsed Time: 1:29:46
Prostate Pulled to Right
Cutting near NVB
Elapsed Time: 1:41:10

Figure 2.8: The prostate is now relatively mobile and can be pulled to either side to allow the dissection along the sides of the prostate to be completed. Dissection here can either include or exclude the neuro vascular bundles depending on the nerve sparing plan in use. Dissection moves from the base to the apex of the prostate.

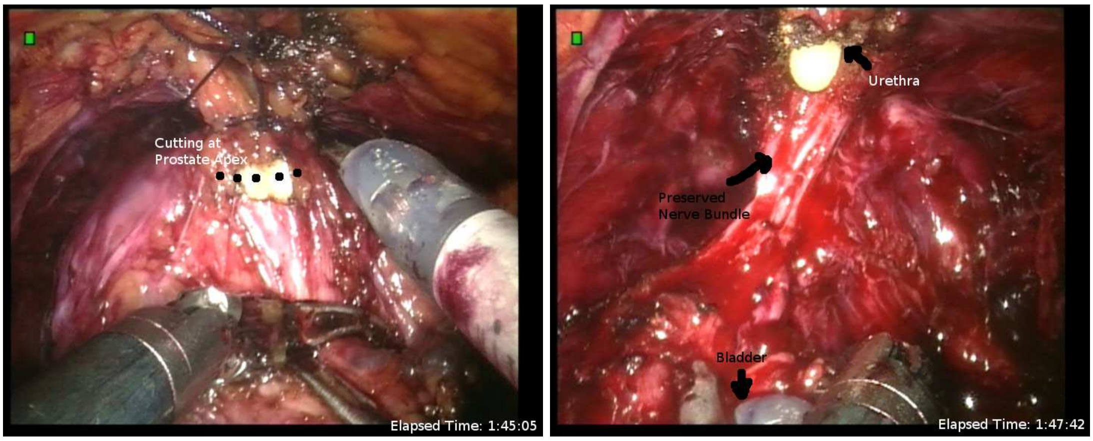

Figure 2.9: The prostate is left attached only through the apex, which is now severed, (on left). The prostate is put in a specimen bag and withdrawn from the operating field. The surgeon is now left with a gap between urethral sphincter and the bladder neck. This gap must be closed in order for the patient to regain continence. This is done by pulling the bladder neck to the urethral sphincter and suturing the two together.

2.1.2 Drawbacks and Opportunities of Robot Assisted Radical Prostatectomy

Whether or not use of the *daVinci*^{TM} robot improves patient out come is a subject of some debate in urology. Some studies show an improvement in patient outcomes [Mavrich (2006)] others are more cautious, while overall there seems little to pick between the two [Atug et al. (2006)]. It is our opinion that given the relative immaturity of the approach, improvements in patient outcomes are more likely to be gained using the RARP approach than remaining with an open approach. Most reports agree that laparoscopic procedures reduce the amount of time the patient must stay in hospital, thereby reducing hospital bed costs. Whether or not the system is cost effective therefore depends on an individual hospitals cost of a bed versus the additional cost of the equipment.

The main difficultly with a RARP procedure is that the surgeon is unable to feel the tissue, reducing the surgeon's ability to assess the spread of cancer. This reduces the surgeon's ability to make informed decisions about the size of the margin to leave around the prostate and whether or not to attempt a nerve sparing approach. There are two ways to adjust for this. One is to introduce a method of haptic feedback for the surgeon, which is an area of active research for many groups. The other is what this thesis is concerned with, which is introducing image guidance to the procedure. This takes advantage of the improving ability of preoperative imaging methods to detect the extent of cancer. If information on the location and extent of the cancer captured preoperatively can be sensibly overlaid onto the surgeons vision, then this can take the place of the tactile feedback available in an open procedure.

2.2 Image Guided Surgery

Before we begin an introduction to image guided surgery we must first define what we mean by the term for the purposes of this report. We start with a broad definition.

Image Guided Surgery 1 *An intervention during which the surgeon's view of the patient is augmented with images from a separate source.*

In a typical RARP the surgeon sees the patient through the endoscope using visible light. We propose adding information from preoperative MRI, making the procedure image guided. An alternative to using preoperative MRI is to supplement the visible anatomy with an intraoperative imaging method, such as ultrasound [Ukimura et al. (2006)]. Intraoperative ultrasound is also used by Leven et al. (2005), but here is laparoscopic and integrated with the surgeon's view through the endoscope. Most modern definitions of image guided surgery, for example that in Peters and Cleary (2008), imply that the additional imaging modalities must be registered to the main modality, so as to appear in the same coordinate system. Strictly speaking this is true, however it is wrong to assume that this must be done by some form of tracking system and/or an overlaid display. It is also possible to display the two modalities separately and allow the surgeon to implicitly align them mentally. Indeed this was the case for what is generally regarded as the first image guided surgery when a x-ray print was used to guide the surgeon during the removal of a needle from a patients hand in 1896 [Brailsford (1946)]. The surgeon mentally aligned the x-ray to the visible anatomy of the patients hand and used the x-ray to visualise the needle. Similarly, Ukimura et al. (2006) do not overlay the ultrasound onto the endoscope view, but rather relies on communication between the ultrasonographer and the surgeon to place the ultrasound data into the endoscope's coordinate system. The new generation of *daVinci*^{TM} systems , the *daVinci S*^{TM}[Tan et al. (2009)], allows image guided surgery with implicit registration by the surgeon by placing preoperative data on the machine's console next to the intraoperative endoscope view.

Such systems become difficult to implement when the anatomy visible in each modality differs. It is also not possible to quantify the accuracy of such systems independently of the surgeon using them. Image guidance using additional registration methods (stereo tactic frames or fiducial markers) is most frequently used in neurosurgery [Skrinjar et al. (2001),Warfield et al. (2000a)] where the surgical targets may be deep in the brain and only a very small section of brain surface may be visible to the surgeon. This has led to the development of equipment and algorithms for tracking surgical tools and registering multimodality images. The dominant registration method in use today is the point based fiducial registration method. For rigid anatomy it is relatively easy to implement and its accuracy is well studied [Wiles et al. (2008),Fitzpatrick et al. (1998)].

These methods lead to what is generally envisaged as the next stage of image guided surgery. Prototype image guided systems for digestive tract surgery [Soler et al. (2004)], heart surgery [Szpala et al. (2005)], RARP [Chen et al. (2008)], urology surgery [Ukimura and Gill (2009)] , and microscope assisted neurosurgery [Edwards et al. (2000)], all envisage overlaying rendered projections of segmented anatomy over the surgical scene. Figure 2.10 gives an example of this, presenting the same data as in Figure 1.1 but as rendered objects rather than MRI slices.

Hawkes et al. (2005) discuss the need to deform the overlaid renderings to account for intraoperative motion. In theory, by extracting surfaces from the endoscopic image [Visentini-Scarzanella et al. (2009)], registering these to the segmented surface and deforming the segmented surface to match

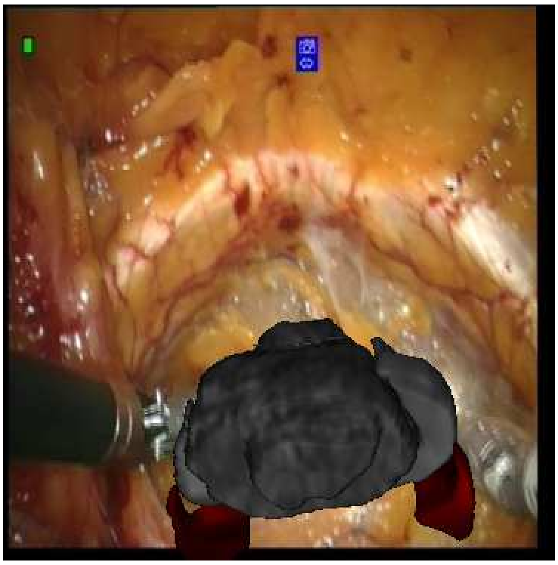

Figure 2.10: In contrast to the MRI overlay shown in figure 1.1 the prostate and neuro-vascular bundles have here been segmented from the MRI and are presented as 3D renderings.

[Skrinjar et al. (2001), Cash et al. (2005)] it should be possible to use such systems for image guidance in non rigid anatomy. Such systems are attractive as they present an intuitive representation of the anatomy to the surgeon and would also allow integration with planning and simulation tools. A segmented patient model aligned to the robots coordinate system would enable optimised port placement [Adhami and Coste-Maniére (2003)]. Such a model would also allow preoperative simulation of the procedure [Bro-Nielsen (1996)] using a patient specific model, alternatively the surgeon could interactively switch to a simulated patient mid operation in order to practice a difficult part of the procedure immediately prior to performing it on the patient. Constraints could also be placed on the operating tools to prevent cutting into non target anatomy [Davies et al. (1992), Ukimura and Gill (2009)]. We present a block diagram of such a system in Figure 2.11

Whilst we agree that such a guidance system represents the long term future of image guidance in urology surgery, we also realise that the implementation of such a system introduces risks and sources of error to the operating theatre that may not be strictly necessary.

There are two issues with presenting rendered segmentations of the anatomy to the surgeon. The first is that regardless of the method used to segment the anatomy there is an error inherent in doing this. This error needs to be properly understood. The second issue is that surgeons are experienced in the interpretation of medical images (MRI,CT and ultrasound) relevant to their speciality. So while rendered surfaces may present a more intuitive image on screen it is not certain that they will give the surgeon a better understanding of the anatomy than a more traditional three plane projection of a MRI volume.

The benefits of registering the auxiliary data to the patient also need to be properly demonstrated. As noted earlier, provided sufficient anatomy is visible the surgeon is capable of mentally aligning the multiple image types. Bringing automatic registration into the operating theatre introduces algorithms and or equipment that will have errors and may fail outright. Therefore it must be demonstrated that

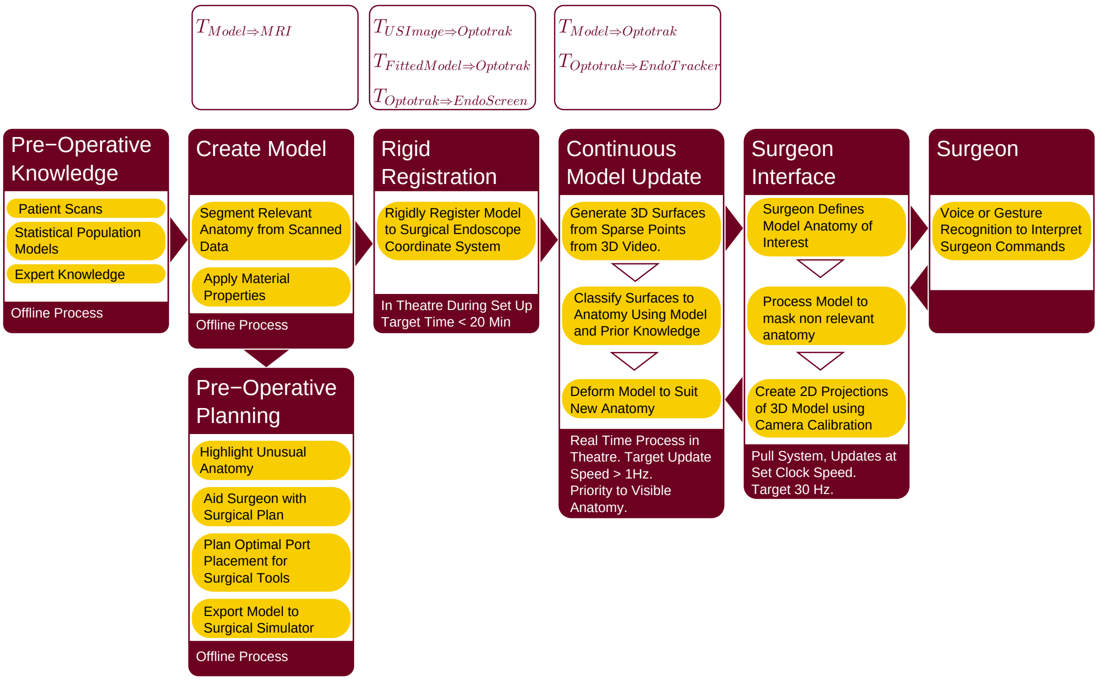
T_Model⇒MRI
T_USImage⇒Optotrak
T_FittedModel⇒Optotrak
T_Optotrak⇒EndoScreen
T_Model⇒Optotrak
T_Optotrak⇒EndoTracker
Pre−Operative Knowledge
Patient Scans
Statistical Population Models
Expert Knowledge
Offline Process
Create Model
Segment Relevant Anatomy from Scanned Data
Apply Material Properties
Offline Process
Rigid Registration
Rigidly Register Model to Surgical Endoscope Coordinate System
In Theatre During Set Up Target Time < 20 Min
Continuous Model Update
Generate 3D Surfaces from Sparse Points from 3D Video.
Classify Surfaces to Anatomy Using Model and Prior Knowledge
Deform Model to Suit New Anatomy
Real Time Process in Theatre. Target Update Speed > 1Hz. Priority to Visible Anatomy.
Surgeon Interface
Surgeon Defines Model Anatomy of Interest
Process Model to mask non relevant anatomy
Create 2D Projections of 3D Model using Camera Calibration
Pull System, Updates at Set Clock Speed. Target 30 Hz.
Surgeon
Voice or Gesture Recognition to Interpret Surgeon Commands
Pre−Operative Planning
Highlight Unusual Anatomy
Aid Surgeon with Surgical Plan
Plan Optimal Port Placement for Surgical Tools
Export Model to Surgical Simulator
Offline Process

Figure 2.11: Block diagram of proposed surgical image guidance system. System is built from seven discrete subsystems. These can be treated in isolation and built into an integrated system when complete.

these methods bring a benefit to the patient before their introduction.

What we propose is to implement a core system, focusing on a rigid body registration between preoperative MRI and the intraoperative endoscope image. It will then be possible to compare the clinical benefits (as measured by the parameters defined in table 2.2) with either a system that does not register the data (ie the *daVinci*^{TM} s) or a more sophisticated system that may warp the preoperative data to better match the visible intraoperative anatomy. Alternatively we can investigate the benefits of presenting the anatomy as rendered segmented bodies versus a MRI overlay. Furthermore a more complex image guidance system that uses endoscopic video data can be initialised by our system. Therefore it is important to understand the accuracy of our proposed system, and if the accuracy is to be improved it is important to understand the sources of error.

For completeness the remainder of this literature review is structured in line with the system areas defined in Figure 2.11. However the systems that deal with the core transforms, see Table (2.3), are dealt with in substantially greater depth.

2.3 Preoperative Knowledge

In order for an image guidance system to be useful it must be possible to image the anatomy of interest preoperatively. This section details what anatomy we intend to image and the methods available to do so. The first part details methods to image the surrounding anatomy of the prostate, while the second section details methods to image prostate lesions.

Transform	Description	Section
$T_{SM \Rightarrow M}$	Fitting a pelvic bone to the patient MRI / Segmentation of bone	2.4.1
$T_{SM(nrr) \Rightarrow O}$	Registration	2.6
$T_{SM(nrr) \Rightarrow UI}$	Image to Image Registration	2.6.2
$T_{UI \Rightarrow O}$	Ultrasound calibration and tracking	2.6.2 and 2.6.2
$T_{O \Rightarrow ES}$	Endoscope calibration and tracking	2.6.3 and 2.6.2
$T_{O \Rightarrow ET}$	Tracking the endoscope	2.6.2
$T_{M \Rightarrow O[nrr]}$	Adjust model based on endoscope image	2.7

Table 2.3: Summary of the core transformation dealt within in this thesis (from Figures 1.2 to 1.5) and where they are in the literature review.

2.3.1 Surrounding Anatomy

Chen et al. (2008) identify the prostate, urethra, rectum, seminal vescicles and surrounding nerve bundles as anatomy that would be useful to locate on a prostate image guidance system. This agrees with the description of nerve sparing RARP given by Tewari et al. (2008). Understanding the position of the neuro-vascular bundles is obviously important to avoid cutting through them. Similarly it is critical not to cut into the rectum. As shown in figure 2.6 the seminal vesicles and vas deference are located on the far side of the prostate so must be pulled through and out of the way. Here it would be useful to know their position and size. Locating the junction of the prostate and bladder prior to the dissection shown in Figure 2.5 would also be useful as would a measure of the size of the medial lobe of the prostate and its extension into the bladder. All this soft tissue anatomy is visible on a T2 weighted MRI. Typically the MRI used has an in plane resolution around 0.7 mm and a slice spacing between 3 and 4 mm. This places an upper limit on the accuracy of an image guidance system that uses data. Of course if it can be shown that higher resolution scans could improve the surgical guidance then this could be gathered at the expense of a longer scan time.

2.3.2 Prostate Pathology

Knowing where the cancerous lesions are within the prostatic capsule allows the surgeon to make an informed decision about the margin to leave around the prostate. The decision on whether to attempt a nerve sparing prostatectomy is generally made preoperatively, however as discussed in Section 2.1 it is possible to alter the plan during the procedure. In the case of open surgery this might be based on palpation of the prostate, while for RARP it may be result of a biopsy [Tewari et al. (2008)]. Provided the cancer can be well imaged prior to the operation image guidance could greatly assist this process.

Cancerous cells are visible on T2 weighted MRI to some extent. However they are not used diagnostically. The size and location of the tumor is generally determined using needle biopsy. The resolution of data from needle biopsies is limited by the number of biopsies taken and the ability to determine the position of each biopsy [Venugopal et al. (2005)].

There is great deal of research into techniques for the non invasive detection and characterisation of prostate lesions. Promising results have been shown by groups using fusion of MRI and MRI spectroscopy [Hricak (2005); Lee et al. (2005); Venugopal et al. (2005)]. Other groups have reported promising results using sonoelastography [Taylor et al. (2005)]. Braeckman et al. (2007) and Mohamed and Salama (2008) have developed systems using (TRUS).

The above is not intended to be an exhaustive review, but rather to highlight that prostate cancer detection is an area currently under intense research scrutiny. It is likely that in the short term future an accurate and reliable system for localising prostate cancer will be clinically available. The primary purpose of developing these systems is to assist in the selection of a treatment method, of which there are many. In the case where radical prostatectomy is chosen as the treatment method, a good knowledge of the precise location and extent of the cancer should enable the surgeon to make informed decisions about the surgical approach (nerve sparing or not) and to accurately judge the margins required.

2.4 Create Model ($T_{SM \Rightarrow M}$)

In addition to being able to image anatomy as in Section 2.3, if we wish to project renderings of segmented anatomy to the surgeon as in Figure 2.10 then the anatomy must first be segmented from the preoperative image. This section deals briefly with this, however as discussed in Section 2.2 presentation of rendered anatomy is not the focus of this work. As discussed in Chapter 1 our system uses the pelvic bone as a registration frame. Therefore what is of relevance here is the processing required to allow this registration in theatre, this will be discussed in depth. As we are more interested in segmenting the pelvic bone we will discuss this first.

2.4.1 Segmentation of Bone from MRI ($T_{SM \Rightarrow M}$)

The following is a brief explanation of our in theatre registration algorithm as an explanation of the requirements of our pelvic bone segmentation algorithm. A more complete description of the registration algorithm will follow in Section 2.6.2. The registration algorithm will use a b-mode ultrasound probe to image the patient's abdominal region and these images will be registered to the patients preoperative data using a image based registration method. We know from past work [Penney et al. (2006) and Barratt et al. (2006)] that it is possible to accurately register ultrasound images of the pelvic bone to CT images of the pelvic bone. Figure 2.12 shows example MRI, CT, and ultrasound images of the pelvic region. Figure 2.13 shows the data shown in Figure 2.12 re-sliced and interpolated to match the ultrasound slice. This more clearly shows the correspondence between the data sets, when a slice to volume registration algorithm is used. The ultrasound and CT both have strong intensity gradients at the bone edges. This indicates that computationally efficient similarity measures such as cross correlation can be used to drive the in theatre registration. The appearance of bone in MRI however is quite different. The cortical bone corresponding to the bright bone edges in the ultrasound appear as dark regions surrounding the brighter trabecular bone. Identifying the bone in the MRI therefore requires reference to both the individual voxel intensity and the intensity of surrounding voxels. It is possible that using a multimodal similarity measure such as normalised mutual information (NMI) might enable direct registration

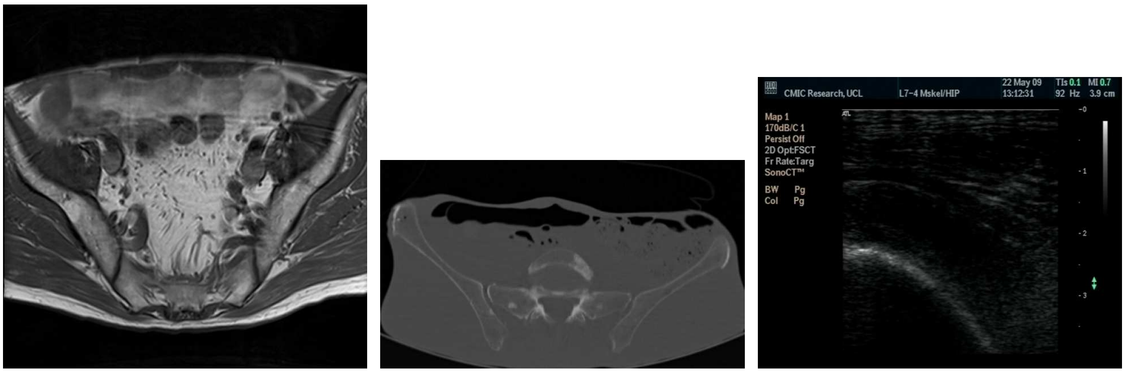

(a) MRI (b) CT (c) Ultrasound

Figure 2.12: Ultrasound (c) and CT (b) images of bone have broadly similar appearance, the bone edge appears as a bright line on the image. On a T2 weighted MRI (a) image of the same anatomy the bone edge does no appear as a bright line. The thick cortical bone appears as a dark line around the perimeter of the trabecular bone in the centre of the pelvic bone. Registering ultrasound to CT is a proven process. Registering MRI to ultrasound for the pelvic bone is untried.

of ultrasound to MRI, however no evidence of this having been done for pelvic bone was found in the literature. The aim of this section therefore is to investigate methods that could be used to process the preoperative data to enable a robust and rapid registration with ultrasound data. In short we require an algorithm to segment the bone from the MRI data.

We have discounted the use of manual segmentation by a radiologist for the following reason. Because of our unique application we are segmenting bone from a T2 weighted MRI image. This is not a task that is of clinical use, and so is not a task that many radiologists have done, nor is there a standard way of doing it. Therefore, in addition to being very time consuming (about 2 hours), the results of manual segmentation would be expected to vary significantly between radiologists. On the other hand an automatic segmentation will be repeatable.

Several algorithms exist in the literature that aim to segment bone from MRI data, either in 2D or 3D [Strzelecki (2004); Carballido-Gamio et al. (2004); Kapur et al. (1998); Lorigo et al. (1998)] using only the available image intensities. Rather than working directly on the image voxel intensities, these approaches process the image intensities to create some measure of the local texture around each pixel/voxel. This is based on the underlying assumption that unlike in CT data where each voxel of bony anatomy is brighter than a voxel of soft tissue, a bony voxel in MRI may have similar intensity to a non bony voxel, however they can be differentiated by looking at the intensity distribution of surrounding voxels. The consensus is that bone cannot be segmented using MRI intensity values alone due to overlapping intensities, intensity inhomogeneity over homogeneous anatomy due to scanning artefacts, and the lack of strong edges at the bone surface. Strzelecki (2004) and Carballido-Gamio et al. (2004) use texture information to segment bony structure in the ankle and lower spine respectively. Kapur et al. (1998) use region growing based on texture information followed by active contour smoothing to segment the femur and tibia. Lorigo et al. (1998) use an active contour directly to segment the femur and

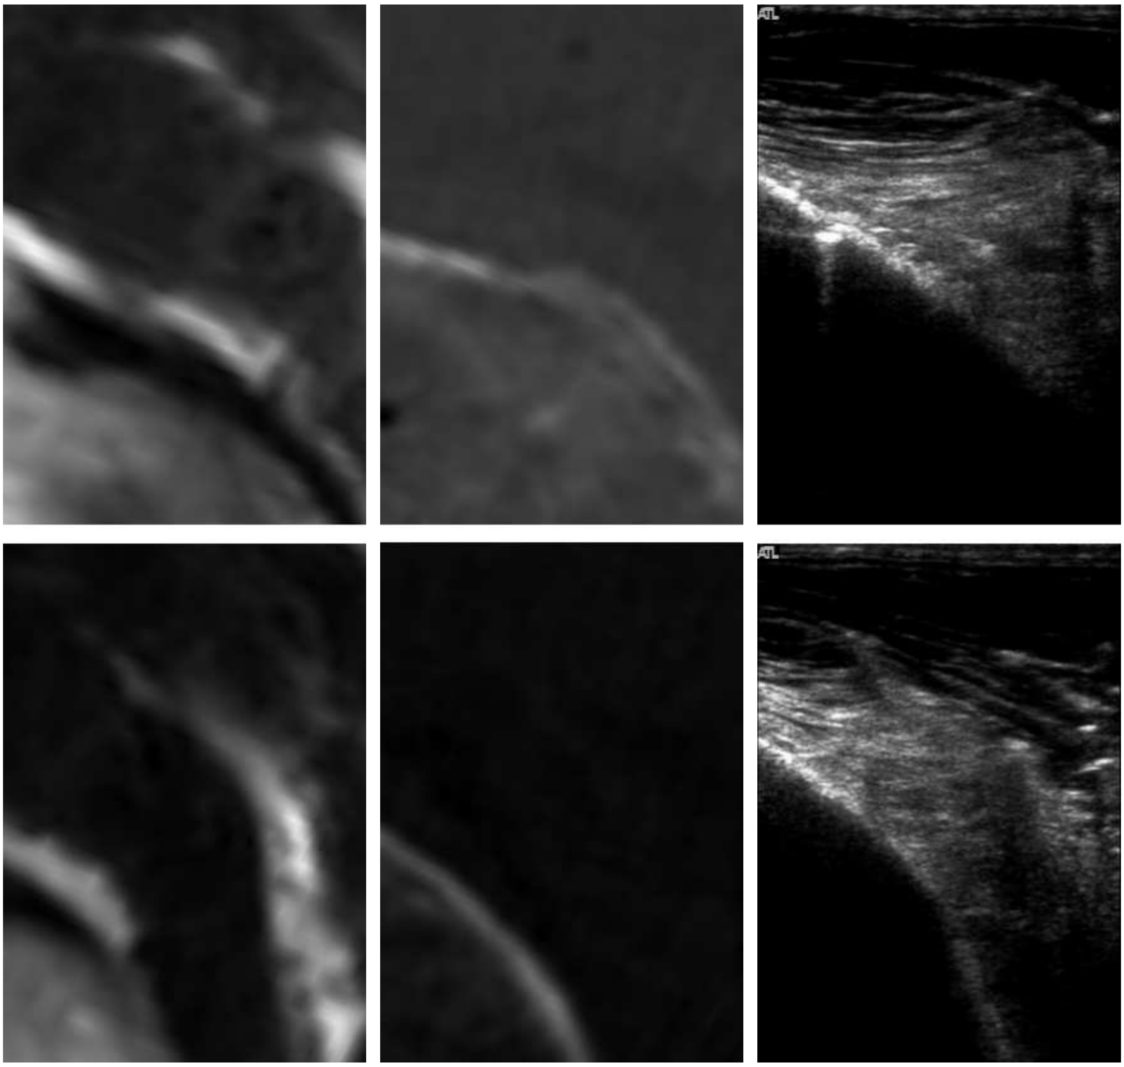

(a) MRI (b) CT (c) Ultrasound

Figure 2.13: This shows the same data sets as in Figure 2.12, but here the MRI and CT have been resliced and interpolated to match the ultrasound. This matches the registration procedure which matches ultrasound slices to volume images. It should be noted here that the ultrasound and CT come from the same patient, while the MRI is from a separate patient that has been registered to the CT, hence the match between these data sets is not perfect. They serve as a useful illustration nonetheless. In both cases the bone edge imaged by the ultrasound at the right is visible as a bright line in the CT. In the MRI the correspondence is not as straightforward, however the general appearance is similar, suggesting that registration using a multi-modal similarity measure has some chance of success.

tibia.

These methods use different approaches to drive the segmentation, however they all demonstrably fail when the chosen measures of image texture become similar for bone and surrounding tissue. In situations where the appearance of bone and soft tissue do not vary outside the segmentation thresholds these methods will work. Segmentation of the pelvic bone, however, is not such a situation. Due to the large size of the pelvic bone and the variety of tissue types immediately adjacent to it the appearance of the bone / soft tissue interface will vary significantly across the pelvis.

Failures in the appearance only algorithms are easily spotted using an understanding of what shape the bone should take. Building prior knowledge of the expected bone shape into the segmentation algorithm should therefore improve performance. This introduces the wide field of shape priors in segmentation. A shape prior can range from a constraint on the curvature of the segmentation to statistical population models for a given object or group of objects. Hoad and Martel (2002) introduce two shape priors when segmenting bone from MRI in the lower spine. They model the spinal column as an elliptical column and mask data anterior to the spinal column. Cootes and Taylor (2004) mention the use of an active appearance model of the knee based on MRI data to segment bone. Fripp et al. (2007) use a shape model to segment knee bone form MRI with some success. Their shape model is based on the active shape model (ASM) frame work introduced by Cootes et al. (1994). They use the minimum description length (Davies et al. (2002)) method to build the model. This works on the assumption that the model that describes the variation in the training population in the most compact way is the correct model. Fripp et al. (2007) also demonstrate that both concepts can be readily applied to multiple objects in the same image.

The shape of the pelvis has been described using a statistical shape model by Lamecker et al. (2004), who used their model to accurately segment bone from CT images. Yao and Taylor (2003) incorporate CT intensity information into their pelvic bone model. Chan et al. (2004) and Barratt et al. (2008) demonstrate the use of a pelvic CT statistical shape model for registration using ultrasound. None of these shape models have been used to segment bone from MRI. Lamecker et al. (2004) mention their intention to extend their model to MRI applications, however no further published work on this was found.

Two methods were found that attempt to automatically segment the pelvic bone from MRI. Boettger et al. (2008) presented their work on segmentation of pelvic bone from MRI for use in radiation therapy planning. However their method required specific MRI sequences and at present is not giving very good results for the pelvis. Schmid and Magnenat-Thalmann (2008) introduce a shape model based segmentation of T1 weighted MRI data for adult female pelves and femurs, driven by the need to do this for diagnosis of osteoarthritis. It is useful here to compare T1 and T2 weighted MRI. In general T1 weighting allows better delineation of bone. Figure 2.14 compares a T1 and T2 image of the same patient from our study. The T1 image on the left does show a better defined cortical bone, however it is still possible to see the bone on the T2 weighted image. In our study we are using T2 weighted scans as they show the soft tissues of surgical interest more clearly, so our results will not be

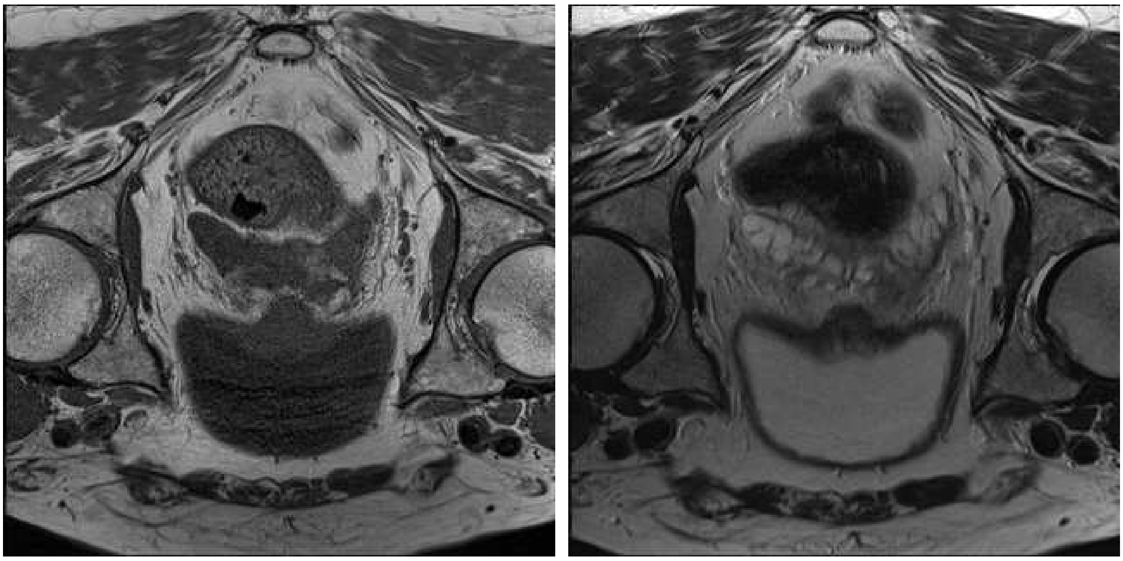

(a) T1 (b) T2

Figure 2.14: The same anatomy shown with a T1 (a) and T2 (b) weighted MRI image. The large circles at either side of these images are the heads of the femurs. The bladder and rectum are also prominent. The cortical bone shows as a dark line around the boundary of the trabecular bone. This line is more precisely delineated in the T1 image, though the difference is marginal.

directly comparable with Schmid and Magnenat-Thalmann (2008), though we should expect them to be similar. Schmid and Magnenat-Thalmann (2008) report a segmentation accuracy of 1.44 mm mean with a standard deviation of 1.1 mm. Assuming this is a Gaussian distribution this equates to an root mean square (RMS) error of 1.81 mm. This is calculated as the distance error between points on the segmented surface and points on a manually segmented gold standard surface. This figure appears excellent in comparison with other MRI bone segmentation algorithms, and in particular, referring to Table 2.5, seems to outperform several shape models based on higher resolution, higher contrast, CT pelves. There are several reasons to doubt the accuracy and generality of the published figure.

1. Results are published for 6 data sets out of a possible 29, without explanation of why these 6 where chosen.

2. The segmentations used to build the shape model and as gold standards are based on “expert” manual segmentations. No literature has been found examining the accuracy of manual segmentation of pelvic and femur bone from MRI, and it is not validated here. The results may vary significantly if the gold standards were prepared by a different expert. However this is a difficult problem to solve for a niche application such as pelvic bone segmentation.

3. The results presented appear to be averages across the pelvis and 2 femurs. It is possible that errors are less on the relatively simple shapes of the femurs and higher in areas of more complex shape such as the pelvis.

4. The method presented is multistage and multi-resolution with various tunable parameters at each level. There is an initial image force (based on intensities) versus a smoothing force. Then the position of the evolved surface within the model space is checked (using the Mahalanobis distance) and if it exceeds a certain threshold the mode coefficients are scaled. This is repeated at different model resolutions. Finally, the SSM is discarded and the segmentation is locally deformed assuming that errors from the true shape form a Markov Random Field. This stage also has a tunable smoothing and image force parameter. In total there are 5 declared tunable parameters. This suggests that while the method can be tuned to work in a given setting it may not perform reliably on a larger, more diverse data set.

Tables 2.4 and 2.5 summarize the relevant results from the literature. Table 2.4 covers the results that deal just with segmentation of MRI. The type of errors reported varies from study to study, as they are used for several different applications. As will be seen later, a system that matches the results reported by Fripp et al. (2007) would provide a good result in the context of the proposed image guidance system. Table 2.5 covers the results for pelvic shape models. The methods used to construct the models are as varied as are the results. The reported boundary errors are less then those reported by Carballido-Gamio et al. (2004), suggesting that a pelvic shape model approach to segmenting MRI will have the best chance of success.

Automatic segmentation of bone is essential for the proposed system. Based on the preceding literature review a system that combines measures of image texture with a statistical shape model would be the best approach. The use of statistical shape models is wide spread in medical imageing and there are many implementations. The following section serves as an introduction to statistical shape models and describes the statistical deformation model (SDM) method that we have chosen to implement.

2.4.2 Statistical Shape Modelling in Medical Image Segmentation

Statistical shape models can be used to parametrize arbitrary shapes. The underlying assumption is that the variation in shape seen across a population a set of N objects (for example the adult male pelves) can be described to a given precision by a set M parameters. If the shape of each member is defined by a vector of length n, then $M << n$. The number of description parameters will depend on the geometry of the shapes being described. To give a trivial case as an example, any number of spheres centred at the same point can be uniquely described a single parameter (the radius). In the case of medical images where complex shapes are being described the number of parameters required will be higher.

Cootes et al. (1994) were among the first to show how statistical shape models could be first constructed and then applied to the problem of image segmentation. Since then there have been innumerable publications on the use of shape models in medical imaging. In general they share the following methodology.

- A training set of images (of size N) is first captured or artificially generated.

- A set of n corresponding points between the training images is defined.

Paper	Method	Application	Result
Strzelecki (2004)	Appearance Only	Ankle Bone	No Quantitative Result
Carballido-Gamio et al. (2004)	Appearance Only	Spine	6.82 mm Boundary Error
Kapur et al. (1998)	Appearance plus regulisation	Knee	No Quantitive result for bone only.
Lorigo et al. (1998)	Appearance only	Knee	No Quantitive results
Hoad and Martel (2002)	Simple shape models	Spine	1.12 ± 0.15 TRE for registration using segmentation
Cootes and Taylor (2004)	Statistical Shape Model	Knee	No Results
Fripp et al. (2007)	Statistical Shape Model	Knee	Dice 0.94
Zhan et al. (2008)	Appearance based on combination of 2 MRI protocols	Pelvis	Poor
Schmid and Magnenat-Thalmann (2008)	Shape Model plus local deformation	Pelvis and femur	1.44 ± 1.1 mm = 1.81 mm RMS

Table 2.4: A summary of published results of bone segmentation from MRI. A minority the published papers presented quantitative results.

Paper	Method	Application	Result
Lamecker et al. (2004)	Statistical Shape Model	Pelvis	0.8 mm RMS Boundary erroor with source image in model , 2.4 mm with out.
Yao and Taylor (2003)	Statistical model	Pelvis	1.27 mm Average Boundary Error.
Chan et al. (2004)	Statistical Model	Pelvis	3.90 mm RMS Boundary Error
Schmid and Magnenat-Thalmann (2008)	Shape Model plus local deformation	Pelvis and femur	1.44 ± 1.1 mm = 1.81 mm RMS

Table 2.5: Results for published work using shape models for segmentation of the pelvic bone. These are all based on CT data except for Schmid and Magnenat-Thalmann (2008) which is MRI based.

- The shape statistics are calculated. By far the most common approach is principal component analysis (PCA), though other methods have been proposed.

- The last step is shape model fitting.

The goal of this thesis is not to examine the various methods of constructing and applying statistical shape models. The following text is an attempt to explain and justify the shape modelling method that was adopted.

Training Set Size

Statistical shape models are built on the shape variation observed across a training population. Therefore the shape model can only be as informative as the training set used to build the model. In cases where the shapes being modelled are simple, only a small training set is required. Returning to the example of a sphere, the population could be fully parametrized using a training set of two spheres of different radii. In the case of medical images where the shapes can be complex a larger number of training shapes will be required to properly describe the population variance. Mei et al. (2008) showed that in the case of 3D face images, 150 training images were sufficient, whereas for the left ventricle of the brain, greater than 319 training images would have been required.

In practical cases the number of training images is generally constrained by the amount of data available. There have been attempts (Cootes and Taylor (1995) and Koikkalainen et al. (2008)) to artificially enlarge the training set using finite element analysis. Unless these enlargement techniques have a grounding in a physical phenomena, for example Hu et al. (2008) where the shape model describes mechanical deformation due to a trans rectal probe, the enlargement amounts to little more than a relaxation of the shape constraints. As we have no physical basis for deforming the pelvic bone, artificial enlargement of the training set is not considered further.

Defining Correspondence

Each shape in the training population must now be defined relative to each of the other shapes. Typically an ordered set of points is defined on the boundary of each shape. Each shape has n points, with point n_i defining the same point in each shape. For relatively simple shapes in 2D, as in (Cootes et al. (1994)), it is practical to manually define a set of corresponding landmark points on the boundary of each shape in the training population. As the shape becomes more complex a larger number of landmarks points (n) is required along with a larger number (N) of training images. This makes manual definition of corresponding points impractical. Methods to automatically define corresponding are therefore an important area of research for the development of practical shape models for medical imaging. One promising method is the minimum description length approach presented by Davies et al. (2002). Using the minimum description length method automates the finding of corresponding points, however it remains necessary to segment the shapes of interest from the training population. A promising alternative is the statistical deformation model (SDM) (Rueckert et al. (2001)). Here the point to point correspondences are defined automatically without the need to segment the training images.

The SDM method uses the deformations required to map each training image to a target image to define the correspondences. A registration algorithm is used to find the transformation of each training image to a target image. Rueckert et al. (2001) use a b-spline registration algorithm (Rueckert et al. (1999)), though the SDM method should work with any registration algorithm. Once an average image has been found the registration process can be repeated, using the mean image as the target, if desired. The shape statistics are then calculated using the deformations to define the correspondences, rather than points, hence the term statistical deformation model.

Calculating Shape Statistics

Each shape in the training set (of N shapes) can now be represented by a vector X_i. In the following we will only deal with the 3D case. For a model defined by point correspondences as for Cootes et al. (1994) where n corresponding points are defined in 3D, each X_i will be $3n$ units long. Similarly for a model defined by deformation fields, each vector will be $3n$ units long where n now refers to the number of vectors in the deformation field. The mean shape or mean deformation field is first found, as;

$$\bar{X} = \frac{1}{N} \sum_{i=1}^{N} X_i \tag{2.1}$$

For each shape, or deformation field, the deviation from the mean is calculated as;

$$dX_i = X_i - \bar{X} \tag{2.2}$$

The covariance matrix for the training set can then be calculated as;

$$S = \frac{1}{N} \sum_{i=1}^{N} dX_i dX_i^T \tag{2.3}$$

The shape, or deformation field, variation seen in the training set can then be described by the eigenvectors of the covariance matrix S. The eigenvectors can be ranked according to their corresponding eigenvalues. Eigenvectors with large eigenvalues correspond to modes of variation that account for a significant part of the variation observed in the training set.

Shape Model Fitting

Any member of the training population can now be described by a weighted linear sum of the eigenvectors. The coefficients of the weighted linear sum form a vector B. For an arbitrary shape (the target), not in the training population, the shape model is “fitted” by finding the set of coefficients, B, that produce a shape that best matches the target, according to some measure of similarity. It is common to truncate the linear sum to only include a subset of eigenvectors, those with the largest eigenvalues. It is also possible to set limits on the allowable coefficient values.

The shape model fitting forms an optimisation problem for which there are a great many different solutions proposed in the literature. We elected to use NMI (Studholme et al. (1999)) as a similarity measure due to its ability to compare images of differing modalities. We chose to use a differential evolution optimiser (Price et al. (2005)) to fit the model. We have not attempted to benchmark these methods against the other methods in the literature.

Some authors allow the fitted shape model to deform further after shape model fitting to fit local edges, Schmid and Magnenat-Thalmann (2008) being one example. The effectiveness of this approach will vary greatly depending on the information content of the images being used. In general these approaches will have similar effects to the “training set enlargement” methods (Koikkalainen et al. (2008)) discussed previously.

2.4.3 Segmentation of Soft Tissue Targets from MRI

Segmentation of soft tissues presents similar challenges to segmentation of bone, along with the further difficultly that soft tissues can deform. The bladder and the rectum can change shape significantly depending on their state of fill. The prostate itself has been shown to translate rigidly and non rigidly deform depending on the patient position [Fei et al. (2003b)]. Hoogeman et al. (2004) show that prostate movement is about 2.4 mm if the general patient position is not radically changed and that prostate position is dependent on the filling of the rectum. For these reasons, a preoperative segmentation of these tissues from the MRI scan will only serve as a starting estimate of their intraoperative location.

MRI is generally acknowledged as a good modality for the delineation of soft tissue, studies by Rasch et al. (1999) indicate that the prostate is shown more accurately in MRI than by CT. This indicates that the soft tissues can be manually segmented from MRI data, the next stage is to investigate whether they can be segmented automatically. Freedman et al. (2005) have shown some success using a patient specific shape model to register the prostate and rectum over several scans (allowing the propagation of a manual segmentation), while Tsai et al. (2004) present a framework for the incorporation of several interacting shapes (prostate, rectum, and obturator muscles) into a single model. Both these studies only use intra patient models and would need extension to inter patient segmentation to be of use for the proposed guidance system. Using the bone shape segmentation from Section 2.4.1 to initialise a shape model incorporating the relationship between separate soft tissues [Tsai et al. (2004)] it should be possible in the future to build a system that can automatically segment relevant soft tissue from patient MRI data. This however is not considered further in this thesis. The detection of nerve fibres from MRI has not been investigated yet. Literature exists for imaging nerves, especially within the skull. It should also be possible to present likely positions of nerve bundles based on the location of other segmented anatomy.

2.5 Preoperative Planning

The existence of an accurate preoperative anatomical map will aid in the preoperative planning of the procedure. Depending on the resolution of the preoperative model it should be possible to draw the surgeons’ attention to unusual anatomy such as very large prostates that may alter the surgeons planned approach [Tewari et al. (2008)]. Similarly, enlarged medial lobes have been highlighted as causing difficulty when cutting the prostate from the bladder. Where these changes may affect the operation time suitable provisions can be made ahead of time. Similarly the model could be used in a surgical simulator to perform a trial run of the operation.

A popular research question in laparoscopic surgery is the optimal positioning of surgical tool ports. In the case of the *daVinci*^{TM} system three bulky arms are competing for space above the patient alongside the laparoscopic tools of two assistants. If ports are not placed correctly situations can arise where the laparoscopic tools cannot access the patient anatomy. For urology surgery there are no bones in the way to constrain the placement of ports however the port planning procedure is not trivial. Adhami and Coste-Maniére (2003) demonstrate the use of existing robot planning software for *daVinci*^{TM} based heart surgery, where the choice of port locations is constrained by the rib cage. Their method relies on an accurate model of the patient anatomy to enable simulation of the pivoting motion of the various laparoscopic arms.

2.6 Registration ($T_{M\Rightarrow ES}$)

This section deals with several of the technologies that are core to an endoscope based image guided surgery system. In essence our image guidance system takes a point in the preoperative MRI and projects it onto the endoscope screen using $T_{M\Rightarrow ES}$. This section introduces various methods to determine a transform from preoperative images to a video screen. This can be done in one of two ways. Firstly the transform can be determined directly using landmarks visible in both the video image and the preoperative image. Secondly the preoperative image can first be put into the coordinates of a tracking system, using a transform equivalent to $T_{M\Rightarrow O}$. The transformed points can then be projected onto the video screen using a transform equivalent to $T_{O\Rightarrow ES}$. The first approach requires that sufficient anatomy is visible on the video screen to perform the registration. As will be seen there are several applications where this is the case. In Chapter 6 we also use this approach based on the visible inner surface of the pubic arch, which becomes visible reasonably early in the procedure, see Figure 2.3. However in order for the system to be operational earlier than this we have adopted the second approach. The remainder of this chapter is divided into three sections. The first looks at direct methods to determine $T_{M\Rightarrow ES}$, the second at methods to determine $T_{M\Rightarrow O}$, and the third section looks at ways to determine $T_{O\Rightarrow ES}$. Although we refer to the Optotrak, MRI, and an endoscope in our transformation definitions it should be noted that we have not limited the literature review to systems that use this hardware.

2.6.1 Direct Registration ($T_{M\Rightarrow ES}$)

Provided the intrinsic parameters of the endoscope are known and the anatomy visible through the endoscope provides sufficient registration points it is possible to register a preoperative model to the intraoperative endoscope view directly. This has the significant advantage that it does not require any tracking of the endoscope, which as will be shown in Chapter 5 is a significant source of error. Mori et al. (2006) demonstrate this for a bronchoscopy (where the anatomy is well defined and largely static). Skrinjar et al. (2000) use a surface derived from a stereo image of the brain to register a preoperative brain model. If the endoscope's "intrinsic" projection characteristics are known it is possible to track the anatomy using either a set of visible landmarks, implanted fiducial markers, or by extracting a surface. Stefansic et al. (2000), and Ukimura and Gill (2009) both use implanted fiducial markers to track moving tissue. The drawback of this approach is that if registration to preoperative data is required the markers must also be visible in the preoperative data. This would necessarily be an invasive procedure. Several groups have demonstrated the tracking of surfaces using either mono or stereo endoscopes,[Hu et al. (2002),Visentini-Scarzanella et al. (2009),Cash et al. (2005)]. If these surfaces are also present on the

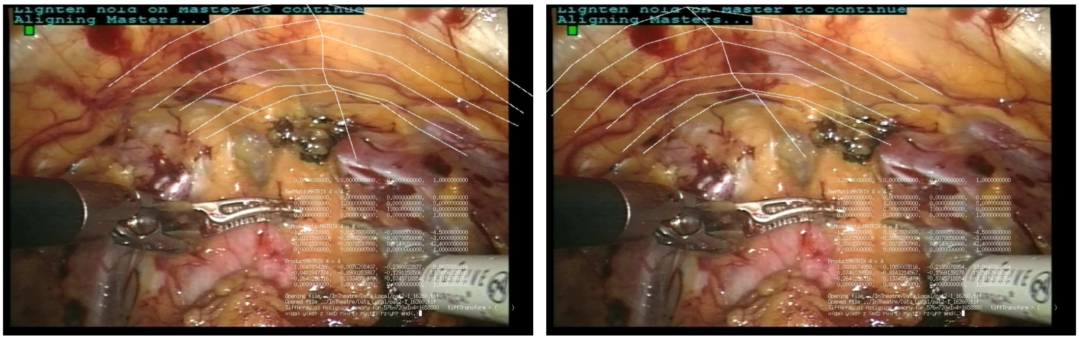

(a) Not Aligned (b) Aligned

Figure 2.15: Once the inner surface of the pubic arch becomes visible it is possible to use it to manually align the preoperative data to the visible anatomy. The set of lines overlaid on these images represent a set of points manually segmented on the pubic arch of the preoperative data. The centre line is shown and the most distant point matches the prostate apex. Provided that the starting estimate is reasonable manual alignment can be achieved in less than a minute.

preoperative data then direct registration should be possible.

In Chapter 6 we use a surface to surface registration procedure to perform a direct registration. Figure 2.15 shows the registration procedure using the visible inner surface of the pubic arch. This is not ideal as it prevents image guidance earlier on in the procedure, and is also difficult to estimate the accuracy of this procedure. Therefore we have favoured using a two stage tracking approach.

2.6.2 Registration $T_{M \Rightarrow O}$

There are many ways to perform the registration of the preoperative data to the intraoperative space, in most cases two steps are required. The patient (or a set of points attached to the patient) must be localised in the intraoperative space, then the preoperative data must be registered to this.

The most widespread method for registration is the use of fiducial markers. These are attached to the patient prior to the preoperative imaging and kept in place during the procedure. They define a set of points that can be seen on both the preoperative and intraoperative data and can thus be registered. The most commonly used registration method is the orthogonal Procustes formulation and singular value decomposition, [Fitzpatrick et al. (2000)]. The main advantage of this method is that an analytic solution exists, enabling rapid and robust registration. Another advantage is that the accuracy of the method is well studied [Fitzpatrick et al. (1998),Wiles et al. (2008)] allowing the likely error at the clinical target to be estimated. However it requires that fiducial markers can be placed in such a way that they do not move in relation to the clinical targets between imaging and surgery. In neurosurgery or orthopaedic surgery where this form of registration is widespread this is possible by screwing markers into the skull or bone, [Stefansic et al. (2000); Fried et al. (1997)]. In theory similar markers could attached to the patient's pelvis for urological surgery. This however makes the image guidance system quite invasive, limiting its acceptance by the patient and/or surgeon. Fiducial markers attached to the skin could also be used. However as the skin moves during insufflation these may not be very accurate. We therefore decided to avoid the use of fiducial markers, at least until any benefit they may or may not have can be properly determined.

The alternative to the use of fiducial markers is to use a second imaging modality that can be localised in theatre to image the patient's anatomy directly. The intraoperative images can then be registered to the preoperative images using an image to image registration method. There are several imaging modalities that are suitable for use during the intervention. These include radiography and cone beam CT [van de Kraats et al. (2005)], intraoperative MRI [Bharatha et al. (2001); Fei et al. (2003b); Warfield et al. (2000a); Fei et al. (2003a)], or ultrasound [Penney et al. (2006); Aylward et al. (2002); Porter et al. (2001); Leroy et al. (2004); Shao et al. (2006)]. Use of intraoperative CT or MRI requires large equipment that simply would not fit into our operating theatre alongside the *daVinci*^{TM} robot. Therefore they are not considered further here.

We elected to use tracked ultrasound to image the patient as the system is small in size and compatible with the *daVinci*^{TM} . It is also known that the pelvic bone can be imaged with b-mode ultrasound and registered to CT images of the bone [Penney et al. (2006)]. Hoogeman et al. (2004), Fei et al. (2003b) and Bharatha et al. (2001) all discuss the fact that the prostate will move relative to other tissues under various conditions. None deal with the specific conditions present on RARP and preoperative planning. However the point is that the position of the modelled prostate will not coincide with the observed prostate in surgery. Whilst we know that the prostate will move in relation to the pelvic bone, we expect that this movement will be small in comparison to our total system error. The apex of the prostate is closely coupled to the pelvic bone, so should not move excessively. The rest of the prostate is known to rotate about the apex based on the filling of the bladder and rectum. During prostatectomy these are both empty, so providing the MRI is also taken at the empty stage, the preoperative and intraoperative positions should be close. The geometry of the pelvic bone is nearly ideal for localising the prostate. When designing configurations of fiducial markers for registration the design goal it to have the clinical target at the centroid of the fiducial markers with large moments of inertia about the centroid [West and Maurer Jr (2004)]. With the prostate near its centre and iliac crests providing large moment arms the pelvic bone meets these criteria.

Using a tracked b-mode ultrasound probe creates a series of slices and their measured location in space. The pixels of each slice can be transformed into a set of points in three dimensional space relative to the tracking device or an arbitrary reference using $T_{UI \Rightarrow O}$. These points can be regarded as an image of the patient that can be registered to the preoperative data. As shown in Figure 1.3, $T_{UI \Rightarrow O}$ can be broken into three separate transforms. $T_{UI \Rightarrow UP}$ and $T_{UP \Rightarrow UT}$ define the ultrasound probe and tracker. They are analogous to the intrinsic and extrinsic parameters of a tracked endoscope and are similarly found using a calibration procedure. $T_{UT \Rightarrow O}$ is the tracking transform and is found using a tracking system. The next two sections discuss these transforms.

Paper	Method	Results
Prager et al. (1998)	Cross Wire Phantom	1.65 mm Reconstruction Error
Prager et al. (1998)	Three Wire Phantom	2.67 mm Reconstruction Error
Prager et al. (1998)	Single Wall Phantom	3.43 mm Reconstruction Error
Prager et al. (1998)	Cambridge Phantom	2.17 mm Reconstruction Error
Barratt et al. (2006)	Point Based Calibration	0.6 mm localisation error
Barratt et al. (2006)	Registration TRE for whole Pelvis with static Calibration	3.96 mm
Barratt et al. (2006)	Registration TRE for whole Pelvis with Optimised Calibration	2.17 mm
Barratt et al. (2006)	Percent Reduction in TRE	45%
Penney et al. (2006)	Registration TRE for whole Pelvis with static Calibration	2.13 mm
Penney et al. (2006)	Registration TRE for whole Pelvis with Optimised Calibration	1.63 mm
Penney et al. (2006)	Percent Reduction in TRE	23 %

Table 2.6: Three sets of published results for ultrasound calibration. The first four rows show that the type of calibration phantom will impact the accuracy of the calibration. The results shown are based on imaging the same point from a large range of angles. These results are not comparable with the 0.6 figure in row five as the methodology used to determine it was quite different. The large improvement shown in the registration TRE in rows 6 to 8 is based on results for the whole pelvic bone on 3 cadaver data sets. The last 3 rows show similar work using self optimising ultrasound calibration, using the same data set but different registration methods.

Ultrasound Calibration ($T_{UI \Rightarrow UP}$ and $T_{UP \Rightarrow UT}$)

In order to create the 3D volume and register it to the patient, it is necessary to calibrate the ultrasound probe with respect to the tracking infra red emitting diodes (IRED) attached to it. Numerous methods exist for performing this calibration [Prager et al. (1998)]. Many groups have come up with ways of performing calibrations using custom phantoms that are quicker/easier/more accurate etc. In general, however, the calibration accuracy is of the order of 1 to 2 mm, see Table 2.6

Barratt et al. (2006) note that calibration methods that use a phantom to calibrate the probe prior to use will suffer inaccuracy when applied to actual patients due to differences in the acoustic properties of the fluid baths used to perform the calibration and the actual patient anatomy. They demonstrate a method to overcome this by building the calibration parameters into the registration optimisation. Penney et al. (2006) implement a similar method by including a speed of sound scaling parameter into the registration optimisation. Table 2.6 shows both their results based on the improvement in TRE over the whole pelvic bone. These suggest that registration using the pelvic bone will incur an error in the order of 1 mm due to acoustic differences between the phantom and the patient. The discrepancy between the results of

Penney et al. (2006) and Barratt et al. (2006) indicates that the exact size of this error will depend very much on the registration method used.

Unfortunately the self calibration methods developed by Barratt et al. (2006) work under the assumption that the error due to calibration parameters forms the bulk of the registration error. The actual registration is performed using manually segmented points and surfaces from CT. It is implicitly assumed by Barratt et al. (2006) that the manual segmentation error is negligible in comparison to the errors due to the incorrect acoustic parameters. Here we are proposing to use a model constrained method for segmentation of pelvic bone from MRI, see Section 2.4.1. Based on the results shown in 2.4.1, we expect that the segmentation errors will not be negligible. Attempting to use the self calibration procedure proposed by Barratt et al. (2006) may lead to incorrect acoustic parameters being fitted to accommodate errors in the segmentation. This is not desirable. Penney et al. (2006) do not use a manual segmentation of either the ultrasound or CT data. Both image sets are processed to ”bone edge probability” images, effectively a probabilistic segmentation. They then optimise the acoustic parameters under the assumption that there is negligible shape error in the bone edge probability images. As discussed above, this is not the case for our segmentations.

Ultrasound Probe Tracking ($T_{UT \Rightarrow O}$)

Numerous methods for tracking ultrasound probes have been reported in the literature, including;

- Optical tracking systems, [Barratt et al. (2006); Lalonde et al. (2003); State et al. (1994)].
- Electromagnetic systems, [Aylward et al. (2002)].
- Mechanical or kinematic tracking or control, [Leven et al. (2005); Porter et al. (2001); Taylor et al. (2005); Shao et al. (2006)].

Electromagnetic systems can perform erratically when in the vicinity of metallic objects and electric motors. Kwartowitz et al. (2009) reported poor results when they attempted to use one in conjunction with a *daVinci*^{TM} . Mechanical systems are generally used when the motion of the probe is quite simple, for example TRUS systems where the probe moves in one dimension only. Construction of a mechanical system for scanning the whole pelvis would be quite difficult. Optical tracking has been used extensively for image guided surgery and are known to have good accuracy and reliability. [Penney et al. (2006); Barratt et al. (2006); State et al. (1994); King et al. (1999); Barnes et al. (2007)]. Therefore such a system will be used for tracking here. The accuracy of tracking using an optical tracking system will vary with the system used and the conditions it is used under. Dey et al. (2002) quote mean tip localising error of 1mm with a standard deviation of 0.3 mm. Housden et al. (2006) notes that the Polaris tracking system suffers from low accuracy for small movements. Most papers quote an error in terms of the tip or endoscope localisation error. As noted by Fitzpatrick et al. (1998) such an error depends not only on the accuracy of the tracking system in tracking a specific marker but also on the geometry of the markers with respect to the point of interest. For the Optotrak Certus system that will be used in this study (because it meets the requirements above and one is available) three published sources have been found that give an actual marker tracking error. These are shown in table 2.7. Barnes et al. (2007) gives

Paper	Method	Result
King et al. (1999)	Unknown	0.2 mm
Barnes et al. (2007)	Tracking of IRED mounted on micrometre stage	0.02 RMS
NDI (1992)	Unknown (NDI system specification)	0.2 mm RMS

Table 2.7: A selection of published figures for the tracking accuracy of the NDI Optotrak system.

a detailed account of how the error was determined. The figure shown in the table is the systems ability to localise a static IRED. They reported a higher figure when the tracked object was moving, due to the timing accuracy of the system. The range of motion and angles covered by the IRED they track is however very small in comparison to the motion we will require (they move the IRED over $1cm^2$. Whilst Wiles et al. (2004) does not provide a quantitative value for the tracking error they do demonstrate that the IRED tracking error is anisotropic (it is larger along the direction normal to the lenses of the tracking cameras) and dependent on the angle of the IRED to the lens' normal and the position of the IRED within the characterised volume of the Optotrak. This makes it likely that the errors we will experience will be greater than those reported by Barnes et al. (2007).

If the IRED tracking errors are known then it will be possible to estimate the tracking error at the centre of the ultrasound slice, either using an isotropic formulation Fitzpatrick et al. (1998) or an anisotropic formulation Wiles et al. (2008). Both papers also assume that the IRED tracking errors are normally distributed and independent along each axis. However if the errors are dependent on the position, angle and speed of the IRED then both these assumptions will be violated when IRED tracking errors are sampled over a range of positions. In order to properly estimate the tracking error for the ultrasound probe a better estimate of IRED tracking error than is available in the literature is required. The effect of the assumptions made by Fitzpatrick et al. (1998) and Wiles et al. (2008) can then be tested against numerical (Monte-Carlo) methods. This is dealt with in Chapter 5.

Using the above methods to determine $T_{UI \Rightarrow O}$ allows us to form a sparse ultrasound image of the patient in the coordinate system of the tracking system. Finding $T_{SM(nrr) \Rightarrow UI}$ will now allow us to put the preoperative MRI into the same coordinate system.

Throughout the above discussion the assumption has been made that the patient's pelvic bone will not move during the operation. This is based on observation of RARP procedures with our clinical partner. If this turns out to not be the case a strategy to account for patient movement will also need to be developed. This could be quite easily done by drawing marker points onto the patients skin and tracking their motion periodically throughout the procedure.

Ultrasound to Preoperative Image Registration($T_{SM(nrr) \Rightarrow UI}$)

Image to image registration is the process of finding a transformation of one image to another that maximises some measure of similarity between the two images. There are two broad classes of image to image registration methods in the literature, intensity based methods and feature based methods [McLaughlin et al. (2005)]. Lying between methods that can be classified as pure intensity or feature based there exists a range of methods that could be described as hybrids of the two. Intensity based methods have the advantage that they are generic, no knowledge is required of the actual image content. However, it is necessary for a relationship to exist between the intensity profiles of corresponding features in each image. They are most effective when registering two images of the same modality, though they will also often work on images from different modalities. Maes et al. (1997) demonstrate the use of mutual information to perform rigid registration between MRI and CT images.

Feature based methods first extract a set of features from the each of the two images. The images can then be aligned by minimising a distance measure between corresponding features. Assuming corresponding features can be extracted from each image, this method is independent of the imaging modalities used. The main drawback of feature based methods is that the registration accuracy is limited by the accuracy of with which the features can be found. For some image types this can be very low. Furthermore, the features used will be specific to the anatomy being registered, requiring prior knowledge of the type of features likely to be present. For example, Porter et al. (2001) extract tubular features (urethra and seminal vesicles) from both MRI and ultrasound. Bharatha et al. (2001) use the prostate centroid to rigidly register two MRI images.

For registering ultrasound slices to pelvic bones several authors have shown that the CT and ultrasound data can be registered directly using a simple transformation of the intensity values in both images, [Penney et al. (2006); Leroy et al. (2004)]. For cases where the preoperative image is MRI based the majority of papers do not perform registration directly between the MRI and ultrasound images. The MRI image is usually first segmented to create features that can be used for registration. We have effectively done this in Section 2.4.1. Aylward et al. (2002) register ultrasound images to a preoperative model showing segmented vascular anatomy, in this case they have segmented the anatomy from CT data, however, reference is made to performing the same procedure using MRI data. Porter et al. (2001) segment tubular anatomy (urethra and seminal vesicles) from a phantom MRI, then register these to the same anatomy segmented from ultrasound. Similarly, Lalonde et al. (2003) manually segment bone surfaces from MRI and ultrasound and register these. Maurer et al. (1999) do similar, registering A-mode ultrasound to bone surfaces manually segmented from MRI. Shao et al. (2006) manually segment the inner surface of the pelvis and register ultrasound data from a TRUS probe to this, they compare the use of three different similarity measures to perform this registration. In contrast to methods requiring segmentation of the MRI data, Penney et al. (2004) register ultrasound to MRI without requiring segmentation. Both ultrasound and MRI images are preprocessed to give a “vessel probability image” which for the case presented (liver) shows the probability of the voxel of interest containing a vessel. These probability images are then registered, with some success.

We now have a proposed method for putting the MRI into the coordinates of the tracking system ($T_{M \Rightarrow O}$). This is made up of three discrete transforms. The method we have chosen is very similar to that of Penney et al. (2006) so it is reasonable to assume that the error in placing the MRI into the tracker

Paper	Method	Result
Penney et al. (2006)	Ultrasound to CT Cadaver Pelvis	1.63 mm Ave. RMS TRE
Barratt et al. (2006)	Ultrasound to CT Cadaver Pelvis	2.17 mm Ave. RMS TRE
Aylward et al. (2002)	Ultrasound to CT Phantom	2.3 Mean TRE
Lalonde et al. (2003)	Manually segmented Ultrasound to MRI on patient pelvic bones	2.6 mm Translation $1.5°$ Rotation
Maurer et al. (1999)	Ultrasound to CT Phantom Skull	1.00 mm Mean TRE
Shao et al. (2006)	Ultrasound to Manually Segmented MRI Pelvis (inner pubis) Patient Data	1.87 mm Mean Translation $2.55°$ degrees mean rotation error (compared with self registration)
Penney et al. (2004)	Ultrasound to MRI Patient Liver (moving)	3.66 RMS TRE

Table 2.8: Published results for registrations based on matching ultrasound data to preoperative images. With the exception of the first two rows none of the results are directly comparable as the size of the objects used for TRE calculation differs or different error measures are used. However it appears that for pelvic registration using ultrasound, a TRE for the bone surface of less than 2 mm should be achievable.

coordinate system will be similar, around 2 mm. The next task is to project the MRI onto the screen of the endoscope using $T_{O \Rightarrow ES}$.

2.6.3 Endoscope Tracking and Calibration ($T_{O \Rightarrow ES}$)

Transforming points from the tracking system's coordinate system into points on the endoscope screen involves two transformations found by endoscope calibration $T_{ET \Rightarrow EL}$ and $T_{EL \Rightarrow ES}$ and and endoscope tracking transformation $T_{O \Rightarrow ET}$. The next two sections deal first with the calibration procedure, then with the tracking. Although the *daVinci*^{TM} endoscope is stereoscopic we have decided to treat only one channel at present. In theory it should be possible to repeat the calibration for the second channel to achieve a stereoscopic projection.

Endoscope Calibration $T_{ET \Rightarrow EL}$ and $T_{EL \Rightarrow ES}$

Endoscope calibration is a well studied procedure. Calibration can be achieved by viewing a calibration object of known geometry, typically a grid of some description. The endoscope is typically modelled as a pinhole with four intrinsic parameters (focal length in two directions and a principal point) together with some radial and tangential distortions. The *daVinci*^{TM} endoscope is corrected so does not suffer from the large radial distortions seen in some endoscopes, however we have kept them in our model to allow for generalisation to other endoscope types. Paul et al. (2005), Dey et al. (2002) and Shahidi et al. (2002) each describe calibration procedures that could be utilised for this study. The method described by Mourgues and Coste-Maniére (2002) applies the method developed by Zhang (1999) to the specific case of the *daVinci*^{TM} stereoscopic endoscope, extending the method to the stereo case and achieving an error due to the intrinsic parameters of around half a pixel. In Chapter 5 we have used a fully automatic formulation of Zhang (1999) method, freely distributed by Wengert et al. (2006). King et al. (1999) show that camera calibration accounts for only a small amount of error in an image guidance system with RMS errors around 0.3 mm. This is small compared to the likely errors in bone segmentation, ultrasound based registration and tissue motion. For the RARP two endoscope lens are used, one is straight and the other angles the lens at 30°. Therefore two calibrations are required. It is also noted that the surgeon occasionally adjusts the endoscope focal length during the procedure. Clearly this would invalidate our determination of $T_{EL \Rightarrow ES}$. However, the adjustment made is small, so we believe that the error induced will be small in comparison to other system errors. If this is shown not to be the case then Stoyanov et al. (2005) propose a method for automatic recalibration of the endoscope during the procedure to allow for changes in the endoscope focal length by the surgeon.

Endoscope tracking ($T_{O \Rightarrow ET}$)

Two methods present themselves for tracking the endoscope. One is to use the same tracking system as for the ultrasound tracking. This would involve attaching a set of IREDs to the *daVinci*^{TM} endoscope. This won't be discussed further here as tracking using the Optotrak has already been discussed in Section 2.6.2. We could find no literature discussing tracking the *daVinci*^{TM} endoscope in this way, though tracking endoscopes in this way is commonplace [Shahidi et al. (2002)].

The more commonly used approach is to use the *daVinci*^{TM} in-built kinematic data [Leven et al. (2005)]. This has the advantage that no IREDs are required to be attached to the endoscope. However it has two drawbacks. Firstly, as the MRI data has been placed in the coordinate system of the Optotrak, use of the *daVinci*^{TM} kinematic data would require an additional transform from the *daVinci*^{TM} coordinate system to the Optotrak coordinate system, introducing an additional source of error. The second drawback is that based on the published literature the *daVinci*^{TM} kinematics have limited accuracy. Mourgues and Coste-Maniére (2002) had limited success with this, with the error due to inaccurate tracking of the endoscope being around 4 pixels on screen. Kwartowitz et al. (2009) present the best results that we have found to date using the *daVinci*^{TM} kinematics (for tracking the instruments rather than the endoscope). They realised that the bulk of the errors in the kinematic system result from the passive joints that are adjusted manually during equipment set up. By combining optical tracking of the passive joints with kinematic tracking of the active joints they achieved errors at the tool tip of 1.39 mm. It should be possible to achieve this error or less using just optical tracking, so our intent is try and track the endoscope using attached IRED. If this can't be done for technical reasons we may investigate the use of the *daVinci*^{TM} kinematic data.

2.7 Continuous Model Update

As mentioned in the previous section the soft tissues in the model are unlikely to be aligned with the soft tissues of the patient due to non rigid motion in between the preoperative MRI and the operation. The system we are proposing does not account for this motion, the surgeon is required to mentally compensate for this motion. However there are methods in the literature that could be used to update a patient specific model to account for soft tissue deformation in near real time. We review some of these methods here to see what directions our system could take in the future.

As discussed in Section 2.6.1 there are several methods for extracting surfaces from endoscopic video, for example Mountney et al. (2006) and Hu et al. (2007). Another rich source of data may come from simultaneously tracking the *daVinci*^{TM} end effectors, either using the endoscope vision or the *daVinci*^{TM} kinematics. Kuhnapfel et al. (1999) demonstrate how the motion of end effectors can be used to predict tissue motion in a surgical trainer. Lin et al. (2005) demonstrate the use of the daVinci kinematics to track the end effector motions and interpret the surgical action. Matching the two would provide a prediction of the anatomical motion that can be compared with the anatomical motion observed through the endoscope. A useful outcome of this process would be the existence of a surgical simulation model in parallel to the actual procedure. Should the surgeon wish to test their next move prior to executing it they could temporarily switch to operating only on the simulation.

The next step is to classify the visible surfaces based on the preoperative model and deform the model. The preoperative model could be continuously updated to reflect the observed position of relevant anatomy during the procedure. Similar ideas have been presented by Warfield et al. (2000b). The updating of the model will depend on the observed anatomy and on the material properties assigned to the model components. These properties will vary throughout the model and throughout the operation. Parts of the model will be very stiff and rigidly positioned (pelvic bone). The prostate itself will require both a degree of elasticity and its allowed movements will change as it is excised. Initially it will be fairly tightly connected to the inner surface of the pelvic bone, but as the operation progresses it will become free to move. Material properties may also include expected visual appearance of organs based on prior knowledge to aid in their classification.

An iterative process will be required to match the observed anatomy with the deforming model, probably borrowing fast finite element modelling (FEM) solvers and graphics processing unit (GPU) implementations from the field of surgical simulation [Bro-Nielsen (1996)].

2.8 Surgeon Interface / Surgeon

Presentation of the preoperative information and allowing the surgeon to control what is displayed is crucial to the acceptance of a surgical guidance system. A large body of literature exists looking at different ways to present information to the surgeon. Some interesting papers are [Paul et al. (2005); Dey et al. (2002); State et al. (1994)]. We have intentionally kept our interface and display as simple as possible. More sophisticated interfaces could then be benchmarked against this. From studies for surgical simulation we know that visual updating must be achieved at at least 30 Hz to be perceived as smooth. It is possible that in the near future haptic feedback will become available on the *daVinci* [Guthart and Salisbury (2000)], if this was the case then the possibility of using haptic signals to guide the surgeon should also be explored. In this case the refresh rate required is much higher, around 500 Hz [Goksel et al. (2005)].

Communication with the system could be achieved either via an auxiliary control pad, voice control, or using gestures on the existing end effectors. It will be necessary to selectively mask model anatomy, depending on the stage of the operation. This could be determined prior to the procedure and/or interactively defined by the surgeon during the procedure.

2.9 Summary

From the preceding discussion it is clear that there are many opportunities for useful research in the field of image guided surgery. Many groups have presented prototype systems that present nice overlays of segmented anatomy. However these systems are not in use in practice and the clinical benefit of the various components cannot be demonstrated. Our aim is to develop a basic guidance system that can be used on a significant number of patients with minimal requirement for additional labour. Coupled to developing the system we will develop a robust understanding of the system errors and the sources of these errors.

The next three chapters deal with three separate error sources. Chapter 3 deals with errors due to preprocessing the MRI, (finding $T_{SM \Rightarrow M}$). Chapter 4 deals with errors in registering the processed MRI to the in theatre tracking system (finding $T_{UI \Rightarrow O}$). Chapter 5 deals with errors in projecting the registered anatomy onto the endoscope screen (finding $T_{O \Rightarrow ES}$). The combination of these three transforms provides a minimalist guidance system, that if properly understood can form a building block for ever more ambitious approaches. This is followed by a chapter detailing our experience of using the system in practice. This focuses on the surgeons experience of the system and should be used to inform our next steps.

Chapter 3

Segmentation of Bone from MRI

3.1 Introduction

As discussed in the preceding chapters we are planning to register the preoperative MRI in theatre using ultrasound images of the bone. In order to do this we plan to pre-process the MRI image to better delineate the bone. In T2 weighted MRI images cortical bone is visible as a dark band around the trabecular bone tissue. However the thickness of the band varies and there are other tissues that have similar intensity. Segmentation of bone therefore requires analysis of not just individual voxel intensities but also the surrounding voxels. In this chapter we detail our algorithm for warping a CT image from a second patient to fit the MRI image of the prospective prostatectomy patient. The algorithm described is an inter-patient, cross-modality, non-rigid registration for the pelvis. No published accounts of similar algorithms have been found. The registration algorithm is constrained using a statistical shape model of the adult male pelvis. The warped CT can either be converted to a bone segmentation using an intensity threshold or used directly for registration to ultrasound images. In addition to describing the algorithm we have attempted to quantify the registration error at the prostate due to using the algorithm as opposed to taking a CT scan of the prospective patient. Although the algorithm has been developed for image guided surgery it could also be used for bone segmentation for osteoarthritis (Schmid and Magnenat-Thalmann (2008)) or in radiotherapy planning (Boettger et al. (2008)).

By using a shape model we ensure that the fitted CT image represents a realistic pelvis, however we expect there will be an error due to the difference between the shapes allowed by the shape model and the actual shape of the patient's pelvis. We quantify this error in three ways. Firstly as a point to closest point on surface distance. This is useful for comparison with other results. Secondly we use the segmented bone shape in simulated registration experiments to estimate the error as a registration error at the prostate surface. Finally we convert registration error at the prostate to an on screen error on the surgeon's console.

3.2 Segmentation Algorithm

Maes et al. (1997) demonstrate that by using a suitable image similarity measure to drive the optimisation it is possible to register MRI images to CT images with a rigid (6 degree of freedom (DOF)) transform when the underlying anatomy is the same, that is rigid intra-subject registration. In our application we do not have CT data of the patient to register to the MRI. The algorithm we propose will register CT from another patient to MRI image of the prospective RARP patient. To achieve a accurate match the registration will be non rigid, allowing the CT data to warp to best match the MRI. Thus we have a non rigid multi-modality inter patient registration. As it is multi-modal we will use NMI [Studholme et al. (1999)] as the similarity measure to drive the registration. As discussed in Section 2.4.1 the non rigid and inter patient nature of the segmentation means that strong regularisation is required to keep the fitted CT realistic. Therefore we use shape constraints derived from the shape differences observed in a training set of 21 adult male CT pelves, ie a statistical shape model.

3.2.1 Shape Model Construction

Data

CT scans of adult male pelves were available from a previous study in orthopaedic surgery. The patient demographics for this study are similar to what would be expected for prostate surgery. Hence we have used this data to build a shape model. The details of the 21 data sets used are given in Appendix A. Each data set has been given a unique alphabetic identifier, from SM-A to SM-U.

Affine Registration

The 21 data sets were manually inspected. A data set that appeared to lie near the average shape was chosen (SM-K[1]) and the remaining data sets were aligned to it. 9 degrees of freedom were used for the alignment (6 rigid degrees of freedom and scaling in three directions). The alignment was done using a using a gradient descent optimiser using NMI as the image similarity measure. As this registration was intra modal (CT to CT) it was not strictly necessary to use NMI as the similarity measure. However, as it is used later to perform inter modal registrations it was used here as well.

The volumes were tri-linearly interpolated to provide a common voxel size across all data sets. Common voxel size is required due to the choice of registration algorithm used to define the correspondences for the SDM. The deformation fields used to define the correspondences for the SDM must be defined in the same coordinate system. The registration algorithm used, see below, defines the deformation field per voxel. Therefore all the CT images must be the same size and have the same voxel dimensions.

Defining Correspondence

The SDM uses correspondence between the individual voxels of the training set images to define the shape variation. The training set images were first transformed to a common coordinate system, using the affine alignment described above. Voxel to voxel correspondences of the transformed images were found using a voxel-wise non rigid registration algorithm. One of the training data sets (SM-U) was chosen as the target image, and the remaining data sets were registered to it using a diffeomorphic non-rigid registration algorithm [Crum et al. (2005)]. At this stage the accuracy of both the affine and non rigid registrations were checked visually.

Data set SM-U was chosen as the target image because the models built with it gave the best segmentation results when used in the leave one out testing that follows in this chapter. We do not know why SM-U makes a better target image, though we assume that it is closer to the model average than the others and so it is easier to register the remaining shapes to it. A more correct approach would to be to create an average shape, then repeat the registrations using the average as the target image. We did not do this here for the following reason. We are going to use leave one out testing to test the performance of the model, that is, for each data set we are going to build a model with the remaining 20 data sets and fit the model to the left out data set. As every model would have a slightly different average shape the non rigid registrations used to build the model would also be different, requiring us to manually check 420 non rigid registrations. By using the same target shape throughout only 21 non rigid registrations require checking, allowing a practical validation. Furthermore, in the context of the system wide errors, the accuracy gained by using the average shape for registration is very likely to be negligible.

---

[1] See Appendix A for details of the data used.

The non rigid registration algorithm [Crum et al. (2005)] outputs a 3D deformation vector for every voxel that defines how to warp the source to the target image. We concatenate the individual voxel vectors into a single vector of dimension $3 \times N_{voxels}$. We have 21 such vectors, one of which is a zero vector generated by registering the target image to itself. These 21 vectors form the basis of our SDM. As we are only interested in the registration of the pelvic bone we apply a mask to these deformation vectors so that any voxels that do not lie within 6mm of the pelvic bone are removed. The mask was based on a manual segmentation of data set SM-U. The mask leaves over 1.7 million voxels active, so each of the 21 vectors has over 5 million dimensions.

Principal Component Analysis

We now construct a statistical deformation model by finding the eigenvectors and corresponding eigenvalues of the covariance matrix of these masked deformation vectors. This would typically be done using singular value decomposition. However the high dimension of the vectors makes it impractical to perform singular value decomposition (SVD) on the resulting very large covariance matrix. Fortunately an alternative method for finding the eigenvectors is described by Cootes in Chapter 7 of Baldock and Graham (2000). An average deformation vector $\bar{X}$ is first calculated. The average deformation vector is subtracted from each of the 21 deformation vectors (X_i). The 21 normalised deformation vectors then form the columns of a 5170869 by 21 matrix D as in equation 3.1.

$$D = ((X_1 - \bar{X}), ..., (X_{21} - \bar{X})) \tag{3.1}$$

Now define a matrix T, as in equation 3.2.

$$T = \frac{1}{21} D^T D \tag{3.2}$$

T is a 21 by 21 matrix, whose eigenvectors and eigenvalues can be found readily by SVD. Given e_i, the eigenvectors of T, the 21 eigenvectors of the full covariance matrix are given by De_i and the eigenvalues are the same as those of T.

The eigenvectors of the covariance matrix are ranked in descending order of their eigenvalues, and these form a statistical model of the deformations observed in the training set. The 21 data sets used

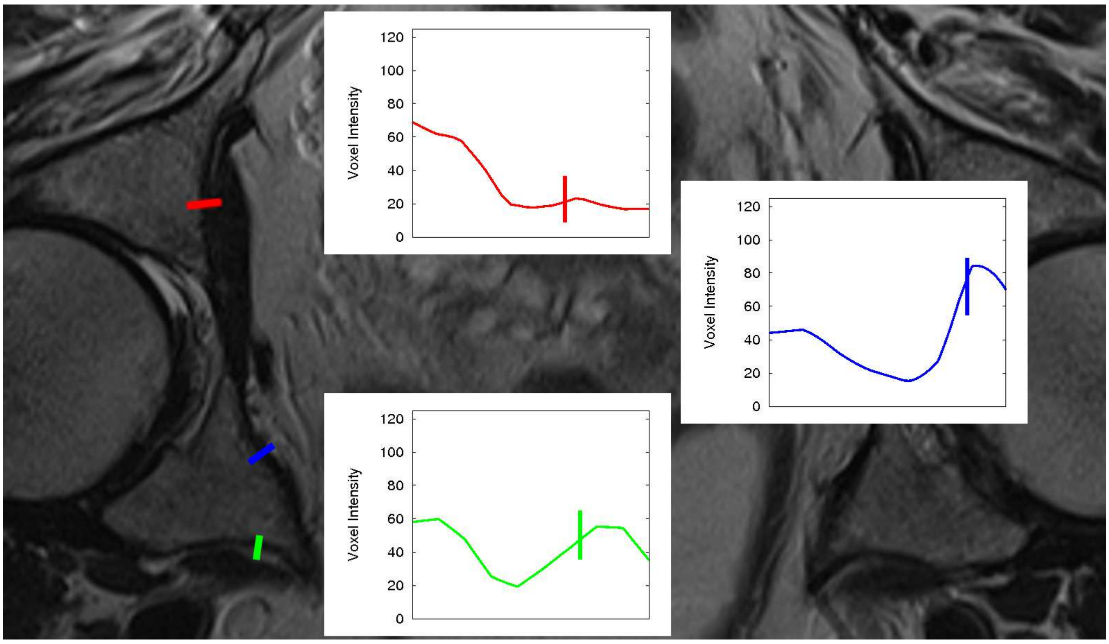

Figure 3.1: Shown here is a single slice of T2 weighted MRI. Three example profiles have been made across the bone edge. The vertical lines across each profile indicate our estimate of where the bone edge lies. Even within this single slice the intensity profiles across the edges are variable. Across the full pelvic bone and between different patients we expect even greater variance. Hence our decision to use shape constrained registration.

will not cover the full range of shapes possible for adult male pelves, so using this shape model for segmentation will result in a segmentation error that will manifest as a registration error at the prostate in the complete guidance system.

There are many strategies in the literature for modifying shape models to achieve higher segmentation accuracies in spite of incomplete training data, Koikkalainen et al. (2008) being a good recent example. However they all amount to a lessening of the shape regularisation provided by the model to allow freer deformation. Another common approach is to use the shape model as an initialisation to a free non rigid registration. This will work well in situations where the model is necessary to find the right edges across the whole image, but boundaries are well defined locally. When images are noisy or, as in our case, the images to be matched are from different imaging modalities this reduction in regularisation can lead to a poorer registration as the free registration moves to the wrong edges. Figure 3.1 gives an example of the intensity profiles across the bone edge in three positions on a single slice of T2 weighted MRI. Even at this local level the intensity profile at the bone edge varies significantly across the image. We believe therefore that trying to optimise the segmentation by relaxing the shape constraints will have limited success.

A further danger of adding a relaxation stage occurs when, as in our case, a small number of validation data sets is available. At present we are testing our algorithm on only four MRI data sets. With enough variables in the algorithm it would be possible to optimise the algorithm to achieve better results than we report here. With only a small validation data set however this could be due to an over optimisation of the segmentation parameters, rather than an improved algorithm. The results may not be repeatable on a larger data set. Therefore rather than investigate ways to fine tune the model performance we have instead tried to keep the model simple and to quantify the size of the registration error due to using the model. If this error is significant we could seek to improve the algorithm when we have a large and more varied validation population.

3.2.2 Model Fitting

With the model defined it can now be fitted to an individual data set, referred to here as the $Target$. This is done as follows. The $Target$ data set is first transformed to the shape model mean shape using $T_{SM \Rightarrow M}[nd9]^{-1}$ from Figure 1.2, to give $Target_{SM}$. The affine transform $T_{SM \Rightarrow M}[nd9]^{-1}$ translates, rotates, and scales the $Target$ data set so that $Target_{SM}$ is aligned with the model average image. The optimal alignment is found by optimising NMI between $Target_{SM}$ and the shape model mean shape using a gradient descent algorithm. We also also re-sample $Target$ so that $Target_{SM}$ has the same voxel dimensions as the shape model mean shape. This is not strictly necessary but was done to simplify the algorithm. The shape model mean shape was compared to $Target_{SM}$ visually to check that the registration process was successful.

We now fit the model mean shape to $Target_{SM}$ as follows. The shape model is fitted by varying a set of model shape parameters b such that the image similarity between $Target_{SM}$ and a model derived image S is optimised. Once again, NMI was used as the similarity measure. S is found by transforming the model target image SM-U using $T_{SM(nd9) \Rightarrow M}[nrr]$. $T_{SM(nd9) \Rightarrow M}[nrr]$ is found using equation 3.3.

$$T_{SM(nd9) \Rightarrow M}[nrr] = T_{ModelTarget \Rightarrow ModelMean} + b\Phi \tag{3.3}$$

$T_{ModelTarget \Rightarrow ModelMean}$ is a row vector of 5170869 dimensions. It is a concatenation of 1723623 3 dimensional vectors that define the voxel to voxel deformations to transform SM-U to the mean of the images used to make the shape model. The dimension of the row vector b is the number of modes of the shape model we choose to use, N. We test the effect of the number of modes used in the Section 3.3. Φ is a N by 5170869 matrix of the N eigenvectors of the shape model. Fitting the shape model is done by changing the N scalar values that make up b so as optimise the NMI between S and $Target_{SM}$.

Optimising the shape parameters b is done using a differential evolution algorithm, [Price et al. (2005)]. This provides a robust optimisation that can be readily constrained. It also does not require the computation of derivatives which could be very time consuming in this case.

Differential Evolution Algorithm

Differential evolution works by evolving a population of possible solutions by generating new solutions based on weighted differences between existing solutions. The new solution is compared with an existing solution, and if it is an improvement it replaces the existing solution in the population. The size of the population (N_p) remains constant. Over time the population of solutions converges. Once the diversity of the population falls below a preset limit the algorithm can be halted, alternatively a limit can be placed on the number of generations.

Parameter	Description
Value to be optimised	Image similarity (NMI)
Convergence Limit	When the standard deviation of the solution population falls below this limit the algorithm is halted.
Maximum Generations	The algorithm is halted after this many populations regardless of convergence.
Population Size (N_p)	The number of candidate solutions in a generation
Starting Population Bounds	The values of the solution parameters in the first generation are constrained to fall between these limits.
Constraint Bounds	It is also possible, but not essential, to constrain all subsequent solutions to fall between preset limits.
Crossover Probability (cr)	A floating point number between 0 and 1, controlling the degree of parameter mixing between candidate solutions
Weight (F)	A floating point number between 0 and 1, controlling the generation of new candidate solutions.
Constraint Method	Different methods were tested for constraining the solution.
Vector Generation Method (Geneticist)	Different methods for creating new candidate solutions. See Table 3.2 for details.
Crossover Implementation	Binomial or Exponential

Table 3.1: The parameters used by a differential evolution solver.

Differential evolution has been shown to be capable of robustly optimising a variety of functions and is fairly robust to the choice of optimiser parameters used [Price et al. (2005)]. Table 3.1 lists the various parameters used by a differential evolution solver.

To start the solution process it is first necessary to generate a starting population of N_p solutions. Each candidate solution is a potential value of b from equation 3.3. We will refer to candidate solutions for generation G as $x_{i,G}$ where i is an integer from 1 to N_p. Each candidate solution $x_{i.G}$ is a row vector of length N, the number of shape modes being used. Price et al. (2005) suggest generating the starting population $x_{i,0}$ using random sampling from the range of allowable parameter values using a uniform probability distribution. For our application we want to constrain the solution based on what was observed in the shape model training data. The eigenvalues of the covariance matrix determined in Section 3.2.1 represent the observed variance along each principal direction of the training data. If the training data is representative of the greater population and we assume a normal distribution then 99.7% of all adult male pelves will fall within three standard deviations of the mean. Therefore we constrain the shape parameters that make up b to fall within the range $\pm 3\sigma$ where σ is the square root of the eigenvalue corresponding to each of the eigenvectors in Φ. We have used the same constraints to limit the random generation of candidate solutions for the first generation.

At each generation a new set of trial solutions, $\underline{v}_{i,G+1}$ are generated and compared with the existing solutions, $x_{i,G}$. We have used a freely available C++ implementation of the differential evolution algorithm[2]. This implementation allows five schemes for generating the trial solutions. The five schemes are detailed in Table 3.2. Two global parameters are used throughout the optimisation, the first is the "weight" (F) whose use is shown in the equations in Table 3.2. This is a floating point number between 0 and 1. A second parameter, the "Crossover Probability" (cr), allows the mixing of the individual shape parameters between the trial solution and the existing solution to which it will be compared. This again is a floating point number between 0 and 1. A crossover probability of 0.5 will result in each trial solution having on average 6 shape parameters from the generated trial solution and 6 parameters from the existing solution. Price et al. (2005) recommend a crossover probability of 0.9 (so 90% of the trial solution shape parameters will come from the generated solution on average) and a weight of 0.8 as suitable for most optimisations.

Two methods of applying crossover are implemented in the algorithm we used. The first method, referred to as exponential crossover randomly chooses a point along the shape vector $v_{i,G+1}$, then a random number between 0 and 1 is sampled from a uniform distribution. If this is less than cr the parameter is taken from the trial solution. A new random number between 0 and 1 is generated and if this is less than cr then the next shape parameter along the vector is also taken from the trial solution. This process is repeated until either cr is exceeded or all N parameters of the vector $v_{i,G+1}$ have been taken from the trial solution. An alternative implementation is binomial crossover where a uniform random number between 0 and 1 is generated for every shape parameter and those parameters where this is less than cr are swapped with the trial solution.

In addition to this, the standard differential evolution algorithm was modified for this application. In its standard guise the algorithm will not constrain the solution. Whilst the starting population is constrained to lie within 3 standard deviations of the model mean, subsequent generations can lie beyond these boundaries. For fitting a shape model it is desirable to constrain the solution. This is easily done for the differential evolution algorithm and Price et al. (2005) discusses several methods for doing this. Three methods were implemented here. These are shown in Table 3.3

Once the differential evolution algorithm is halted, either when the population converges or when a maximum number of generations is reached, the best solution is selected. The best solution is the one that gives the highest value of NMI between S (the image derived from the shape model) and $Target_{SM}$ (the patient image). This solution gives b in equation 3.3, hence $T_{SM(nd9)\Rightarrow M}[nrr]$ is determined. $T_{SM\Rightarrow M}[nd9]^{-1}$ is inverted to give $T_{SM\Rightarrow M}[nd9]$. $T_{SM\Rightarrow M}$ is shorthand for applying $T_{SM(nd9)\Rightarrow M}[nrr]$ followed by $T_{SM\Rightarrow M}[nd9]$. Applying $T_{SM\Rightarrow M}$ to the shape model target image, SM-U, gives the shape model approximation (S) of the patient image ($Target_{SM}$). S is a warped CT image of the patients pelvic bone that can be used for in theatre registration.

The remainder of this chapter is divided into sections, each reporting an experiment performed to validate the algorithm. We first test the algorithm's sensitivity to the various parameters used by the

---

[2] DESolver.h version 1.0 written by Lester E. Godwin, PushCorp, Inc.,Dallas, Texas, 972-840-0208 x102, godwin@pushcorp.com

ID	Geneticist	Cross Over	Method
0	Best1	Expon.	$\underline{v}_{i,G+1} = \underline{x}_{best,G} + F \times (\underline{x}_{r_1,G} - \underline{x}_{r_2,G})$
1	Rand1	Expon.	$\underline{v}_{i,G+1} = \underline{x}_{r_1,G} + F \times (\underline{x}_{r_2,G} - \underline{x}_{r_3,G})$
2	RandToBest1	Expon.	$\underline{v}_{i,G+1} = \underline{x}_{i,G} + F \times (\underline{x}_{best,G} - \underline{x}_{i,G}) + F \times (\underline{x}_{r_1,G} - \underline{x}_{r_2,G})$
3	Best2	Expon.	$\underline{v}_{i,G+1} = \underline{x}_{best,G} + F \times (\underline{x}_{r_1,G} + \underline{x}_{r_2,G} - \underline{x}_{r_3,G} - \underline{x}_{r_4,G})$
4	Rand2	Expon.	$\underline{v}_{i,G+1} = \underline{x}_{r_1,G} + F \times (\underline{x}_{r_2,G} + \underline{x}_{r_3,G} - \underline{x}_{r_4,G} - \underline{x}_{r_5,G})$
5	Best1	Binom.	$\underline{v}_{i,G+1} = \underline{x}_{best,G} + F \times (\underline{x}_{r_1,G} - \underline{x}_{r_2,G})$
6	Rand1	Binom.	$\underline{v}_{i,G+1} = \underline{x}_{r_1,G} + F \times (\underline{x}_{r_2,G} - \underline{x}_{r_3,G})$
7	RandToBest1	Binom.	$\underline{v}_{i,G+1} = \underline{x}_{i,G} + F \times (\underline{x}_{best,G} - \underline{x}_{i,G}) + F \times (\underline{x}_{r_1,G} - \underline{x}_{r_2,G})$
8	Best2	Binom.	$\underline{v}_{i,G+1} = \underline{x}_{best,G} + F \times (\underline{x}_{r_1,G} + \underline{x}_{r_2,G} - \underline{x}_{r_3,G} - \underline{x}_{r_4,G})$
9	Rand2	Binom.	$\underline{v}_{i,G+1} = \underline{x}_{r_1,G} + F \times (\underline{x}_{r_2,G} + \underline{x}_{r_3,G} - \underline{x}_{r_4,G} - \underline{x}_{r_5,G})$
			$\underline{x}_{best,G}$ is the best solution in the current generation, and $\underline{x}_{r_{1,2,3,4,5},G}$ are randomly chosen members of the current generation.

Table 3.2: Table of methods used for generating the trial solution. Five different trial vector generation algorithm are implemented by the standard differential evolution algorithm, together with 2 implementations of crossover. This gives a total of 10 possible methods for generating the trial solution vectors.

differential evolution solver and the number of shape modes used. We then test the algorithm on MRI data and attempt to measure the error as an on screen projection error.

3.3 Effect of Number of Shape Model Modes on Error

3.3.1 Aim

When fitting a shape model to data the number of modes used is an important parameter. In shape models where there are a large number of training sets in relation to the dimension it is likely that the higher order modes (those corresponding to the smallest eigenvalues) contain mostly noise rather than genuine anatomical information, therefore there is a good argument for restricting the optimisation to the lower modes only (those corresponding to the largest eigenvalues). For a given shape model it is possible to plot the contribution of each mode to the shape space in terms of cumulative variance. In many examples in the literature a cut off value of cumulative variance is then chosen, and only modes below this are used. Often values of 90% or 95% are used.

More thorough methods have been developed to determine the number of modes to retain in order for the shape model to best capture the training data. Mei et al. (2008) use a method that analyses the convergence of models made with subsets of the training data to determine the number of modes required and the number of samples required to build the model. They test their method on shape models of up to 1581 dimensions (527 point correspondences in 3D). In an anatomical (face data) model

ID	Description
0	If any shape parameter in the trial solution vector is beyond the constraint boundaries, the entire trial solution vector is discarded. The trial solution vector is replaced with a solution vector generated with the same algorithm used to make the first generation, ie. uniform random sampling of the allowable solution space.
1	If a shape parameter in the trial solution vector is beyond the constraint boundaries, this parameter only is clipped to the constraint boundaries.
2	If any shape parameter in the trial solution vector is beyond the constraint boundaries, the entire trial solution vector is discarded. The trial solution vector is replaced with a new trial solution generated by repeating the procedure from Table 3.2 with different randomly selected $\underline{x}_r$.

Table 3.3: Three strategies were trialled to constrain the solution. The first strategy will prevent the population converging if the global maxima is beyond the constraints. Using it may reduce fine scale searching for the global maxima. However it may be beneficial if the algorithm veers towards a local maxima that lies beyond the constraints. The second strategy should not effect convergence in the same way, however it is pointed out in Price et al. (2005) that it will reduce the diversity of the difference vector population reducing the effectiveness of the optimiser. The third strategy is probably most compatible with the unconstrained differential evolution algorithm, however it may fall into very long loops while it searches for a new random trial solution that fits within the constraints.

with 1500 dimensions they show that 70 modes should be used and approximately 150 samples are necessary to build the model, that is 0.05 modes should be retained per dimension and 0.1 samples are required per dimension. Whilst these figures cannot be directly applied to our case as the numbers will depend on the anatomy being modelled, we can safely say we have insufficient training samples. In all likelihood we are at least one order of magnitude away from having sufficient samples (we have 21) to build a complete model of the male pelvic bone. This is partly an artefact of the way we have constructed our shape model. We have a very high dimensional shape model because we have used voxel to voxel correspondence, rather than the more common method of segmenting surfaces and defining point correspondences using a sparse point set on these surfaces. Alternatively we could use a b-spline registrations algorithm [Rueckert et al. (1999)] and use the node points to define the correspondences [Rueckert et al. (2001)]. Chan et al. (2004) do this for the pelvic bone, resulting in a 1185 dimensional model. However it is clear that we still won't have enough samples. Therefore we realise that our model will give only an approximation of the variation in the adult male pelvis. None the less the model as it stands can still be used for our application, and we will show that the results it gives compare well with other published segmentation methods. The aim of this experiment is therefore not to determine how many modes we need, rather to determine if changing the number of modes used within the range we have available (0 to 20) has a significant effect on our results.

3.3.2 Data

Two tests were performed in this experiment. The first tested the model performance on CT data using the 21 adult male CT data sets SM-A to SM-U. This was followed by validation on MRI data, Patient-01, Patient-02, XMR-01, and XMR-02. Details of all data can be found in Appendix A.

3.3.3 Method

Two tests were done. Firstly the CT data used to build the model was utilised. Each CT volume was approximated by a model built from the remaining 20 data sets, as described in Section 3.2.1. The model based approximations were repeated using different numbers of retained shape modes. 4, 8, 12, 16 and 19 modes were retained, 19 being the maximum number of modes possible for a shape model built from 20 data sets.

Approximating a Data Set with a Shape Model

Each data set (SM-X) from the shape model training data can be approximated using a shape model built from the remaining 20 data sets as per equation 3.4.

$$\text{SM-X} \approx T_{SMApprox} \times \text{SM-Target} \tag{3.4}$$

Rather than using the numerical optimisation detailed in Section 3.2.2 to estimate $T_{SMApprox}$ we can use a direct analytic method, because we know T_{NRR}. T_{NRR} is the transform found when building the full shape model in Section 3.2.1. We can solve equation 3.5 for b using an analytic least squares method. This prevents any errors induced by a failed optimisation affecting our results here.

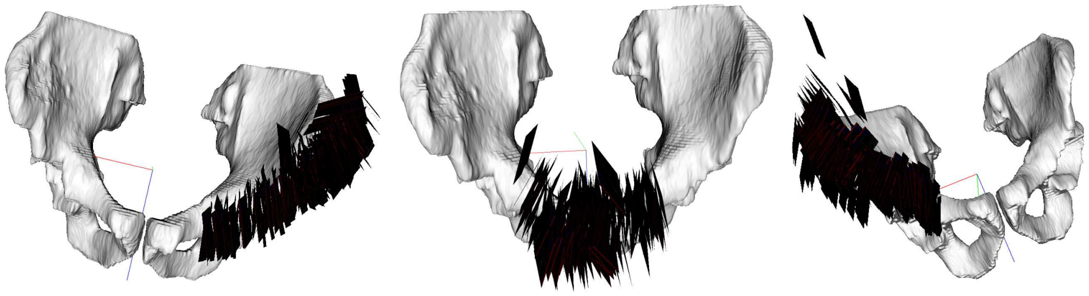

Figure 3.2: The slice distribution used for the registration algorithm. From left to right, 228 possible left iliac slices, 175 pubis slices, and 247 Right iliac slices. Note that left and right refer to the patient's left and right.

$$T_{NRR} \approx T_{ModelTarget \Rightarrow ModelMean} + b\Phi \tag{3.5}$$

The value of b found by solving equation 3.5 can then be used in equation 3.6 to give $T_{SMApprox}$.

$$T_{SMApprox} = T_{ModelTarget \Rightarrow ModelMean} + b\Phi \tag{3.6}$$

Manual segmentations of the shape model target image can then be propagated using both $T_{SMApprox}$ and T_{NRR}. We will refer to the two resulting segmentations of SM-X as $S_{SMApprox}$ and S_{NRR} respectively. Comparing the two segmentations produced gives a measure of the shape approximation caused by using the shape model. By changing the length of the vector b we can see how this changes as we change the number of shape modes used.

We are interested in two measures of the error here. The first is the registration error at the prostate first discussed in Section 1.6.1. As this cannot be calculated directly from the two segmentations but relies on a separate registration algorithm we also present a more direct measure of boundary error.

Prostate Registration Error due to Model Segmentation Error

The following method was used to estimate the registration error at the prostate. A set of simulated "ideal" ultrasound slices are first generated by re-slicing (using tri-linear interpolation) S_{NRR}. S_{NRR} contains the entire pelvic bone surface. However, only limited sections of the bone can be imaged intraoperatively using b-mode ultrasound. Therefore before re-slicing S_{NRR} we applied a mask to remove surfaces not visible to the ultrasound probe. Re-slicing was performed using physically realistic ultrasound planes taken from the phantom registration experiment described in Chapter 4. These slice locations are shown in Figure 3.2.

These simulated ultrasound slices were registered to the model based $S_{SMApprox}$ segmentation using a gradient descent registration algorithm. The algorithm is described in full in Chapter 4. In brief it first converts the slice images to list of points in 3D space (using $T_{UI \Rightarrow O}$). Points with an intensity below a threshold (in this case 1) are discarded. The result is a list of coordinates in 3D space and corresponding intensities. These points are transformed using an estimate of $T^{-1}_{SM(nrr) \Rightarrow UI}$ and the voxels of $S_{SMApprox}$ sampled using nearest neighbour interpolation to find a matching model intensity value. The model intensity and ultrasound intensity are multiplied to give an intensity product. This is done for all points in the ultrasound image and the intensity products summed. The value of this sum is maximised by varying the estimate of $T^{-1}_{SM(nrr)\Rightarrow UI}$. This is done using a gradient descent approach, with multiple step sizes.

The registration algorithm is sensitive to the starting position of the registration and the precise distribution of slices used. Therefore the registration was repeated many times from different starting positions and with different randomly chosen ultrasound slices. The starting positions were randomly and uniformly spread over a six dimensional (one dimension for each of the six degrees of freedom) hypersphere of radius 10mm. (Rotations were scaled as $1^\circ \approx 4\mathrm{mm}$)[3] . The radius of this hypersphere was chosen ensure that the registration algorithm reliably converges to the correct solution. The intent here is to test the effect on registration accuracy, not the capture range of the registration algorithm.

The registration error was then measured as the TRE at six points on the surface of a “nominal” prostate (see Figure 1.6) that was the same for all 21 data sets. A nominal prostate was used so the errors for each data set were directly comparable. In reality the position of the prostate would be different for each data set. The effect of this simplification is discussed in Section D.

Average Segmentation Boundary Error

$S_{SMApprox}$ and S_{NRR} are binary segmentation images. Points on the bone surface have an intensity of 1 and all other points have an intensity of 0. Average boundary error was found by finding the nearest surface voxel (with intensity 1) in S_{NRR} for every surface voxel in $S_{SMApprox}$. The mean value is reported.

Validation on MRI Data Sets

To confirm the trends observed using the CT data, model based segmentation was repeated on four MRI data sets. Here a shape model built from all 21 CT training data sets was fitted to each MRI data set using the optimisation algorithm detailed in Section 3.2.2. Segmentation errors were calculated by comparing the model based segmentations with manual segmentations of the MRI data. Only average boundary errors were calculated. The process was repeated using 12 and 20 modes only, as the results from the CT data showed minimal improvement when more than 12 modes were used.

3.3.4 Results

Fitting to CT Data

Figure 3.3 gives the results of the evaluation made using the CT data sets. A steady improvement in TRE can be seen as the number of modes used is increased. This is in line with expectations as the model is more able to fit the target shape. Using a more direct measurement of the model fitting accuracy, average boundary error, also illustrates this point, however here the improvement is very small.

Figure 3.3 shows that the gain between using 12 modes and 19 modes is relatively minor, about 0.2mm, and there is no gain between 16 and 19 modes. We attribute the slight increase in TRE at the

---

[3] This was chosen so that over the pelvic region a step change of rotation gave the same average displacement as a step change in translations, see Chapter 4 for more details.

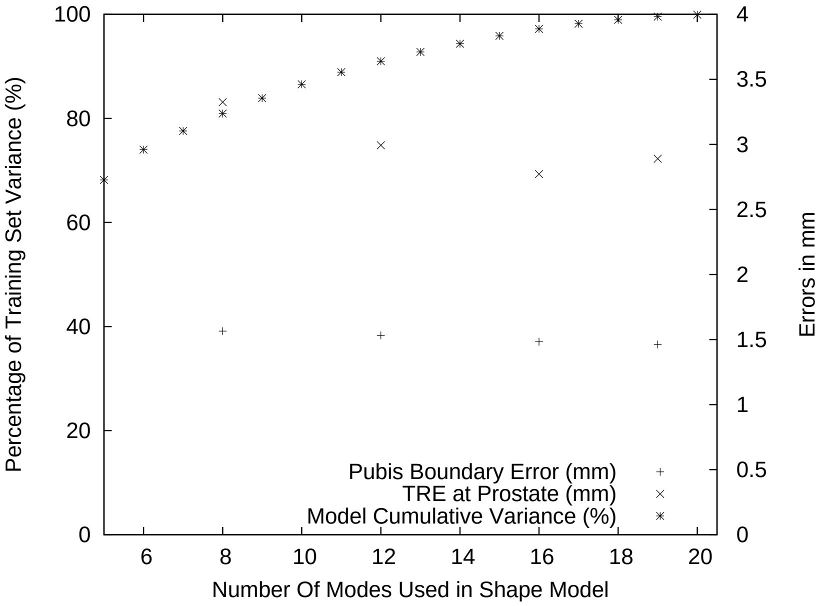

Figure 3.3: Plots of the model cumulative variance, the TRE at the prostate and the the boundary error, based on the CT data. It can be seen that the TRE steadily reduces in a similar trend to the model variance. The average boundary error also reduces, though the gains are small in absolute terms. The slight increase in the TRE at prostate from 16 to 19 modes, together with the slight decrease in boundary error suggests that the gains in segmentation accuracy are masked by the accuracy of the registration method used to assess the TRE at prostate.

Data Set	12 Mode		20 Mode	
	NMI	Ave BE	NMI	Ave BE
Patient-01	-1.017590	1.45 mm	-1.018702	1.61 mm
Patient-02	-1.014382	3.08 mm	-1.016171	3.17 mm
XMR-01	-1.022439	2.20 mm	-1.022332	2.12 mm
XMR-02	-1.016674	1.70 mm	-1.016611	1.59 mm

Table 3.4: Performance of the differential evolution model fitting algorithm on four MRI data sets, in terms of the optimisation similarity metric (NMI) and the average boundary error. Increasing the number of modes used does improve the similarity metric for Patient-01 and Patient-02. For XMR-01 and XMR-02 there is no improvement. In Section 3.6 we will show that repeated runs of the optimisation algorithm with the same data give NMI values with a standard deviation of 6×10^{-5}. This indicates the slight degradation in NMI for XMR-01 and XMR-02 can be attributed to the repeatability of the algorithm. The changes in boundary error are less predictable, reflecting the fact that boundary error and NMI are not well correlated, as will be discussed in Section 3.7.

prostate from 16 to 19 modes to the accuracy of the registration method we are using to estimate the registration error. That is, even though we get a more accurate segmentation of the bone, as measured by boundary error, using 19 modes, the limited accuracy of the registration method prevents this being realised in TRE at the prostate.

Fitting MRI Data

Table 3.4 gives the results of model based segmentations of the four MRI data sets, when using 12 and 20 modes. We show the results as an average boundary error and also show the optimised value of NMI reached. Whilst allowing more modes to be used appears to allow a slightly better NMI to be reached in two cases, this did not correlate with a reduced boundary error. In the other two cases the optimiser was unable to achieve a higher NMI despite the additional modes being available.

Table 3.4, together with the CT results above suggests that even though using more modes in the shape model can deliver a more accurate bone segmentation, the accuracy of both our optimisation algorithm and our registration algorithm prevents this improved segmentation potential being converted to a better TRE at the prostate above 12 modes.

3.4 Extrapolation of the Iliac Shape Using a Shape Model

3.4.1 Aim

It will be shown in Chapter 4 that in order to register the prostate to the patient accurately in theatre it is necessary to register using data from both the iliac crests and the central pubic region. This conforms to our expectations based on our understanding of fiducial marker configuration from Section 2.6.2. Figure 3.4 shows an example of the ultrasound slices used to register the model to the patient during the in theatre trials discussed in Chapter 6. One advantage of using the CT based shape model to perform the

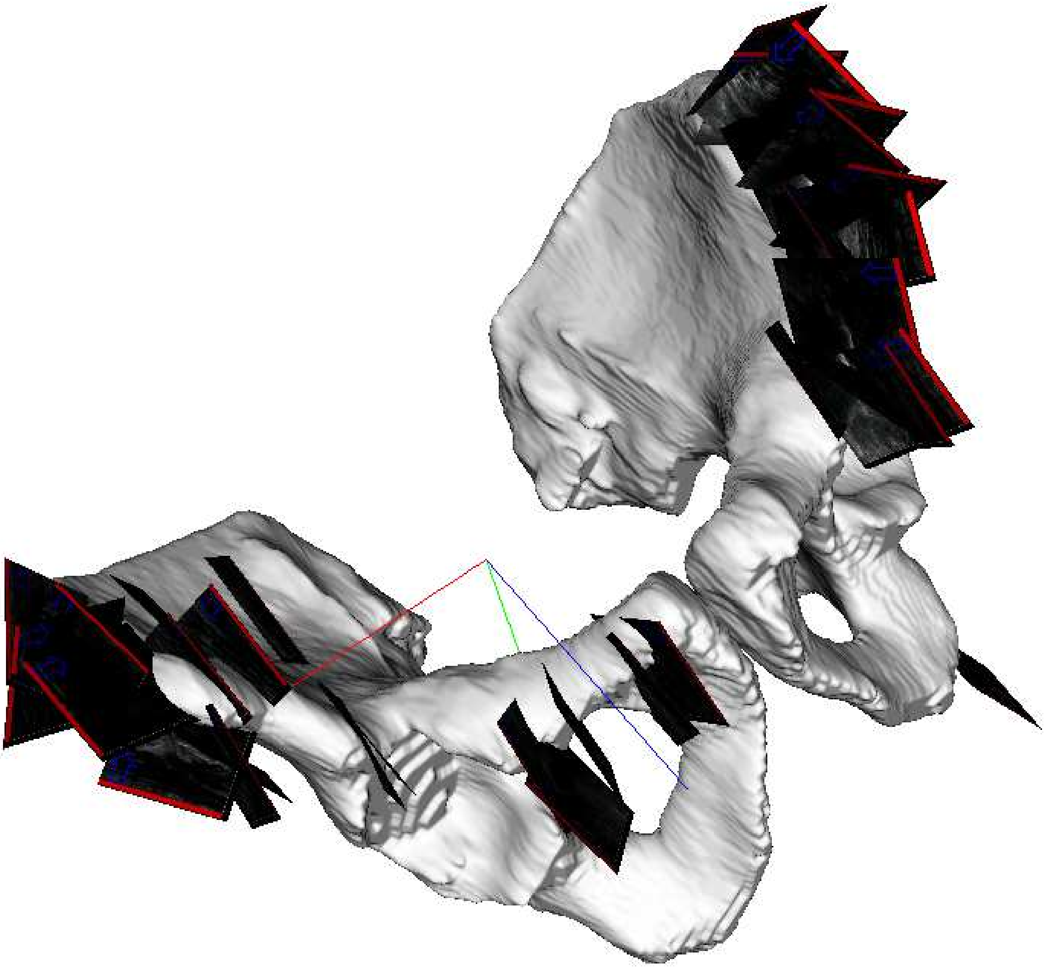

Figure 3.4: An example set of ultrasound slices positioned on a pelvic bone segmentation from the CT based shape model. Slices on both the iliac crests and pubis are necessary to accurately locate the prostate.

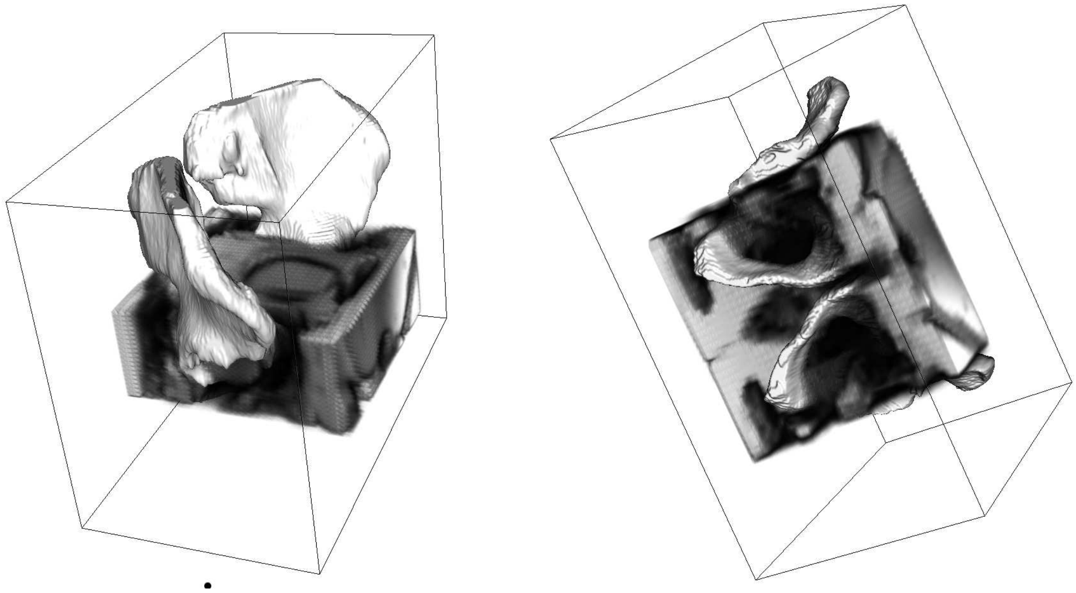

Figure 3.5: Full segmentation of pelvic bone with typical prostate MRI overlaid. This volume spanned by this MRI volume was used as a mask when fitting the shape model to the CT data.

segmentation is that in cases where the complete bone surface is not visible on the MRI an estimated bone surface can be extrapolated by fitting the model to the known data. The MRI volume shown in Figure 3.5 is typical of the preoperative prostate MRI collected at Guy's Hospital during this project. As can be seen it does not cover the entire pelvic bone. One solution to this problem would be to alter the MRI protocol to collect a more complete scan, however this cannot be done retrospectively. An alternative is to fit the model to the available data, then build a full pelvis using the data present in the model. This section aims to determine whether this is a valid method for this application.

3.4.2 Data

This experiment used the 21 adult male CT data sets used to build the shape model. The MRI data from Patient-01 was used to mask the CT data when fitting the shape models, so that only those regions visible in the MRI were used for shape fitting.

3.4.3 Method

The same methodology as described in Section 3.3.3 was used to determine the model segmentation error and the resulting TRE at the prostate. This process was done twice. Firstly with SM-X masked to cover a region of 6mm around the entire pelvic bone, and secondly with SM-X masked to cover a region of 6mm around the pelvic bone only in the volume defined by the available MRI data, see Figure 3.5.

As will be discussed in Chapter 4, the registration algorithm is sensitive to the spatial distribution of ultrasound slices used, yielding best results when there are comparatively more pubis slices. As it would be expected that the pubis region should be more accurately approximated when the iliac regions are masked out, the registration test was repeated with a uniform distribution of slices and one with more pubis slices. In each case 100 repeat registrations of each data set were used to assess TRE.

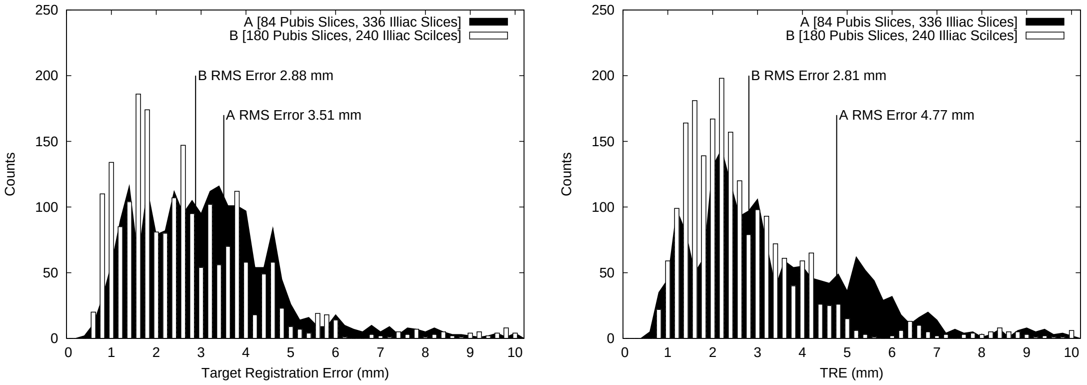

Figure 3.6: Histograms of the target registration errors at the prostate surface using 2100 repeats of 4 different registration methods. On the left are registrations using the shape model fitted to the full pelvic bone, on the right are registrations using the shape model fitted only to data inside a mask representing a typical prostate MRI volume. In each case two different distributions of simulated ultrasound slices were used, one placing more weight on data from the pubic region (Method B) than the other(Method A). When using less pubic slices the masked segmentation performs substantially worse than the unmasked (4.77 mm RMS vs 3.51 mm) segmentation. When more pubic slices are introduced there is a very small but statistically significant improvement (2.81 mm RMS vs 2.88 mm RMS) in the registration errors. Outliers have been removed from data shown. Outliers were defined as results where the TRE was greater than 3 standard deviations away from the mean value. For the unmasked data method A had 83 (4.0%) outliers and method B 68(3.2%). For the masked data method A gave 55 (2.6%) outliers and method B 60 (2.9%).

3.4.4 Results

TRE at Prostate Surface

Figure 3.6 presents histograms showing the TRE at the prostate for both the masked and unmasked model based segmentations. In each case two distributions of ultrasound slices were used. Registrations with a TRE greater than 3 standard deviations from the mean TRE for each method were rejected as outliers. Comparing the corresponding results for the masked and unmasked model based segmentations it can be seen that when using a larger number of pubic slices there is very little difference between using the masked and unmasked data. As the results are not normally distributed a Kolmogorov-Smirnov test was used to test the significance of the difference between the masked and unmasked results. This gave a p-value of 7.286e-09, suggesting that the improvement seen is statistically significant, however the difference is clinically insignificant. (2.81 mm vs 2.88 mm).

Boundary Errors

Table 3.5 gives the boundary errors for the two methods of model based segmentation. For boundaries in the iliac region the average boundary error is greater when using the masked model, while for boundaries in the pubic region the average boundary error is less. The significance of these differences was tested using Kolmogorov-Smirnov tests. It was found that the difference for the pubic region was statistically

Method	Iliac BE (Average mm)	Iliac BE (Std Dev. mm)	Pubis BE (Average mm)	Pubis BE (Std Dev. mm)
Full Pelvis	1.26	0.31	1.28	0.34
Masked Pelvis	1.82	0.50	1.25	0.33

Table 3.5: The average boundary errors over two regions of the pelvis for both the masked and unmasked model fitting. Results are the average boundary error for each of the 21 segmentations averaged over the 21 segmentations and the corresponding standard deviation. In line with the results shown in Figure 3.6 the masked model has a substantially greater boundary error on the iliac region and a slightly smaller boundary error in the pubis region.

insignificant with a p-value of 0.987, while the differences for the iliac region were significant with a p-value of 0.0054. The boundary error results indicate that when the model is fitted using a mask that excludes the iliac regions, the boundary error in the iliac regions is greater. The boundary error in the pubic region is either unaffected or very slightly less, due to the model being able to more accurately fit the more limited data.

More usefully the TRE results show that for the application of image guided prostatectomy, a segmentation based on a limited sample of the pelvis can be used without loss of registration accuracy. One drawback should be noted however. The process we have just described determined $T_{SM(nd9)\Rightarrow M}[nrr]$. The full registration $T_{SM\Rightarrow M}$ also requires $T_{SM\Rightarrow M}[nd9]$. We used an estimation of $T_{SM\Rightarrow M}[nd9]$ based on the unmasked volumes for the preceding results. We found that determining $T_{SM\Rightarrow M}[nd9]$ using the masked anatomy gave poorer, though still usable results, ie. the masked model does lead to a lead to a loss of registration accuracy. We are looking at ways to overcome this by improving the process to find $T_{SM\Rightarrow M}[nd9]$.

3.5 Results as TRE at the Prostate

3.5.1 Aim

As discussed in Section 1.6.1 it is our aim to present the results of this section as a TRE at the prostate. The process for calculating this is provided in Section 3.3.3. This method registers simulated ultrasound slices taken from a “gold standard” segmentation of the actual data to the model derived segmentation. “Gold standard” segmentations of the iliac crests are available for the CT data. For the MRI data sets the iliac crests are not imaged, so no gold standard segmentation is available for registration. At present therefore we can only assess the MRI results as boundary errors in the region were MRI data is present. The goal of this experiment is to determine if the error as prostate TRE can be predicted from the boundary surface error. The resulting correlation may enable estimation of the TRE at the prostate for the MRI data, as will be done in Section 3.7.

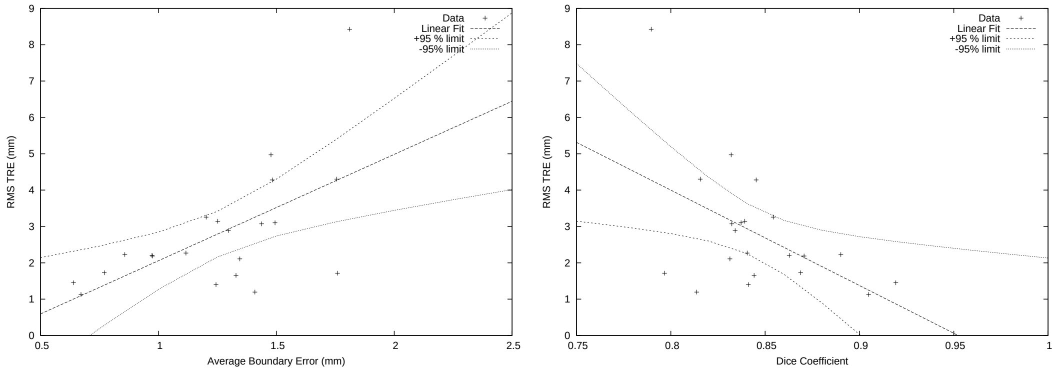

Figure 3.7: TRE at prostate boundary versus Dice coefficient and RMS boundary error. For each of the 21 data sets used in Section 3.4 three measures of error, the RMS TRE reported in Section 3.4, Dice's coefficient and the boundary surface error. By plotting these it is possible to see whether they are correlated and thus whether it is possible to predict TRE from either the Dice coefficient or from the average boundary error. These plots show the data, a fitted linear regression and the $\pm 95\%$ confidence limits. The linear regression between boundary error and TRE has a Pearson's correlation coefficient of 0.57. The linear regression between Dice's coefficient and TRE has a correlation coefficient of 0.47. These linear regressions will be used in Section 3.7 to predict the TRE at the prostate for a given average boundary error.

3.5.2 Method

We calculated boundary surface error in the central pubis region for the results from Section 3.4. We also calculated Dices coefficient [4] for the same data. Linear regressions were then fitted between boundary error and prostate TRE and Dice's coefficient and TRE.

3.5.3 Results

Figure 3.7 plots the three error measures. It is possible to estimate the RMS TRE at the prostate for a given boundary error or Dices coefficient. Both measures are correlated with prostate TRE with correlation coefficients of approximately 0.5 in both cases. It is clear from Figure 3.7 that such predictions will not be highly accurate, nonetheless they will provide an indication of likely TRE.

3.6 Differential Evolution Optimiser

3.6.1 Aim

As discussed in Section 3.2.2 the differential evolution optimiser is set up with a number of parameters. This experiment aims to see what effect, if any, varying these parameters has on the segmentation accuracy. From this we define the optimum parameters for our application. Furthermore we would like to quantify the effect that failure to find the global optima may have on TRE at the prostate.

---

[4] Dice's coefficient [Dice (1945)] measures the degree of overlap between the shape model derived segmentation and a "gold standard" segmentation.

Variable	Range
Weight F	0.10 0.28 0.46 0.64 0.82 1.00
Crossover cr	0.10 0.28 0.46 0.64 0.82 1.00
Geneticist	0 1 2 3 4 5 6 7 8 9
Constraint	0 1 2
Total Runs	$6 \times 6 \times 10 \times 3 = 1080$

Table 3.6: Parameters tested for evaluation of differential evolution optimiser. Values for weight and crossover were uniformly sampled from the range 0.1 to 1.0. Setting either of these to 0 will create an unchanging population. Please refer to Tables 3.1, 3.2 and 3.3 for explanations of these parameters.

3.6.2 Data

MRI data from XMR-02 was used to test the effect of optimiser parameters. This data set appears representative of the other three MRI data sets, so the results should be similar. As the results indicated that the differential evolution algorithm was reasonably insensitive to the parameter values, it was decided not to repeat the experiment on the remaining data sets.

3.6.3 Method

To evaluate the various selection and constraint strategies and the effect of varying weight and crossover probability model fitting was repeated for different combinations of parameters, as shown in Table 3.6. In all cases 12 shape modes and a population N_p of 60 were used.

The results of all optimisations were evaluated by comparing the optimum value of NMI achieved. At present their is no time pressure on this part of the procedure, so parameter combinations that achieved slow convergence and hence longer run times were not penalised.

Once the optimum parameters were identified, the repeatability of the algorithm was tested. By default the starting population, whilst randomly generated, uses the same seed so the same starting population was used for all runs. To test whether the results of the optimisation were dependent on the starting population the optimisation was repeated 20 times with different random starting populations using only the optimum parameters.

3.6.4 Results

Figure 3.8 plots NMI achieved for different values of the four parameters.

A predicted TRE is given on the right hand axis of each graph. This is calculated from the NMI using a pair of linear regressions. The first converts boundary error to prostate TRE and was developed in Section 3.5, see Figure 3.7. The second converts NMI to boundary error and was found by calculating the boundary error for this data set using the method from Section 3.3.3, then fitting a linear correlation to NMI. This will not provide a very accurate estimate of the TRE but it is included to help visualise the likely system wide effects of using a suboptimal combination of solver parameters.

Based on the results in Figure 3.8 the parameters values shown in Table 3.7 were used for further

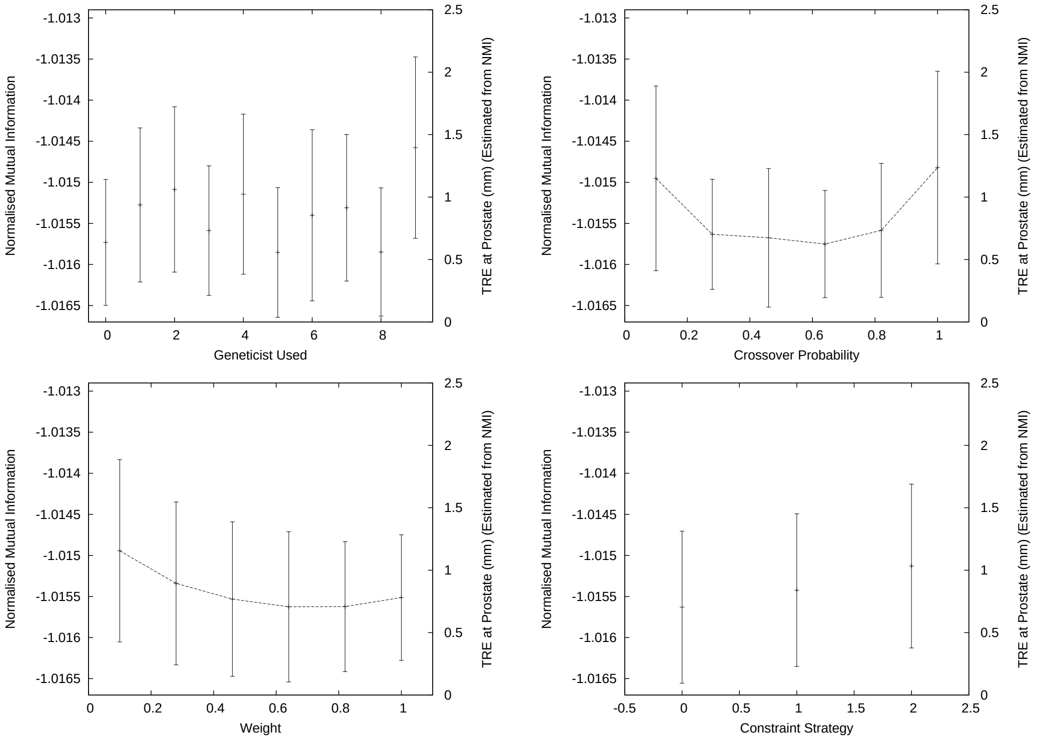

Figure 3.8: The results of varying the four parameters of the optimisation algorithm. The error bars represent ± 1 standard deviation of all the results at that level. These are generally large as for each parameter the remaining parameters continue to play a significant role. Based on these charts, Geneticist ID 5 was selected as the best solution generation method, and constraint method 0 (replacement with random solutions) was used. Both the continuous variables, (weight and crossover) show minima in the range 0.6 to 0.8. More points could be tested in this range, giving an unambiguous minima, however as the curves appear fairly flat in this range this would probably not yield a significant improvement in results. Therefore the weight value was rounded to 0.8 and the crossover used was 0.64. The predicted TREs are given on the right hand axes to help interpretation. Conversion form NMI to boundary error was done by fitting a linear regression between the two based on the data here. Then boundary error was converted to prostate TRE using the regression from Section 3.5. The resulting formula is $TRE = 661.99 + 651.07 \times NMI$. It is not expected that this will be a very accurate prediction of TRE however it allows the systems sensitivity to these parameters to be visualised.

Parameter	Value
Similarity Measure	NMI
Convergence Limit (Population Standard Deviation)	0.001
Maximum Generations	100
Population Size	60
Starting Population Bounds ± Standard Deviations	3
Constraint Bounds ± Standard Deviations	3
Crossover Probability	0.64
Weight	0.80
Constraint Method	0 - Replace with random solution
Vector Generation Method (Geneticist)	Best1
Crossover Implementation	Binomial

Table 3.7: The optimal set of parameters for running the differential evolution algorithm in this application.

optimisations. It appears however that apart from extreme values of the crossover probability or a low weight, the choice of parameters does not have a large influence on our system results.

The choice of constraint strategy prevents the algorithm's convergence criteria from being met as each time a solution falls outside the constraint boundaries a new random solution is created. The new solution will almost certainly be far from the existing population. Therefore the algorithm generally runs to 100 generations before being halted. On current equipment this takes around 12 hours, which is sufficient for the application when the MRI is taken well in advance of the surgery. If run time was added as a performance metric to be balanced against similarity measure achieved then the results of this experiment would be substantially different.

The algorithm was now rerun 20 times with these parameters unchanged, but with different randomly generated starting populations. Figure 3.9 plots the NMI achieved on each run together with the Mahalanobis distance from the solution with the best NMI. Plotting the Mahalanobis distance enables a visualisation of the 12 dimensional solution space. It is clear that the algorithm does not find a true global minima on each run, however the difference in terms of predicted TRE is negligible (TRE standard deviation of 0.04 mm, NMI standard deviation is 0.000066).

3.7 Fitting the CT shape model to MRI Data

3.7.1 Aim

The aim of this experiment is to determine how accurately the CT statistical shape model can be fitted to MRI data. This requires optimisation of the shape vector to best fit the available MRI data.

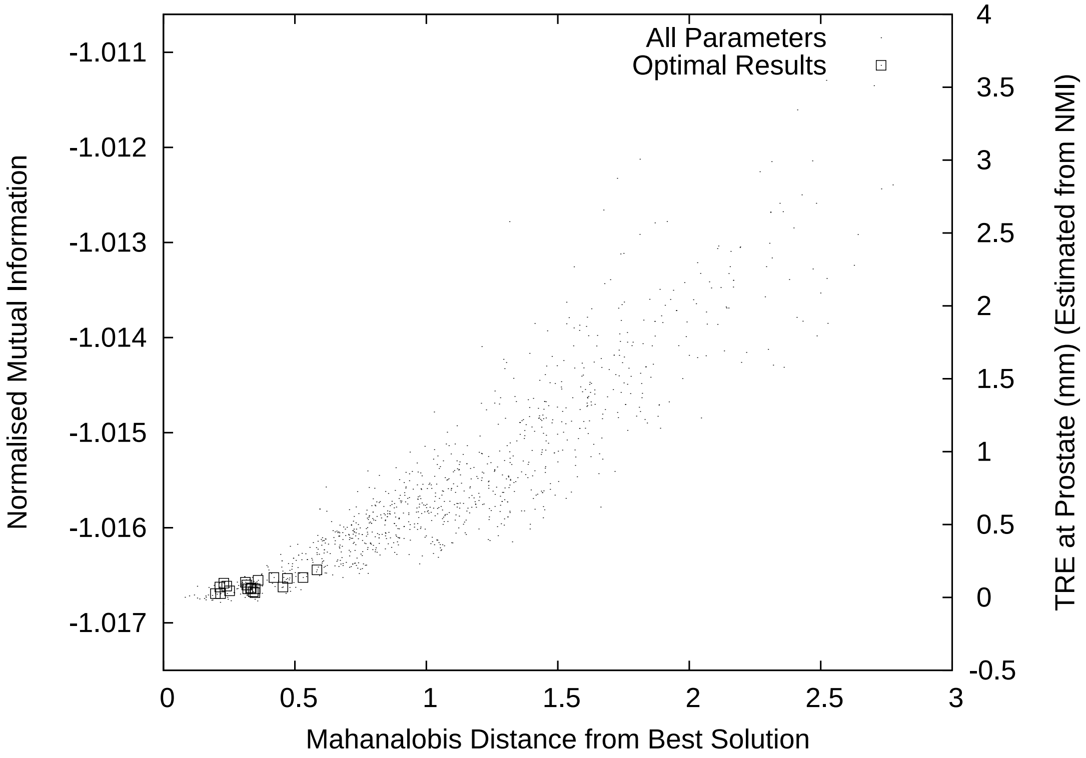

Figure 3.9: By defining the solution with the optimal NMI mutual information as the origin of a 12 dimensional parameter space, it is possible to plot each other solution based on its Mahalanobis distance from the origin. This together with the solution's NMI gives an idea of the shape of the solution space. The fact that there are many solutions with high NMI near the optimal solution indicates that there are many local minima near the global minimum. When using the best optimiser parameters the algorithm appears stable, consistently getting near the best solution. The effect on predicted TRE is negligible, the standard deviations of the repeat solutions being 0.04 mm TRE.

Data Set	Registered To	Attempts
Patient-01	SM-Average	1
Patient-02	SM-T	2
XMR-01	SM-K	1
XMR-02	SM-K	1

Table 3.8: The MRI data sets were aligned to the CT shape model using a gradient descent optimiser working on normalised mutual information. The quality of these affine alignments was checked visually. In cases where the alignment appeared poor the alignment was repeated, either from a different starting position or using a different data set from the shape model training set. The table shows which data set each MRI set was successfully registered to and how many target data sets were used.

3.7.2 Data

Four MRI data sets (Patient-01, Patient-02, XMR-01, XMR-02) were used to validate the model, details of these are provided in Appendix A.

3.7.3 Method

For each data set a model based segmentation was compared to a gold standard segmentation. As the data sets do not contain the full pelvic bone it is not possible to perform a simulated ultrasound to bone registration as in other experiments, slicing the gold standard segmentations in the iliac regions gives empty slices. Therefore the segmentation accuracy was measured using boundary errors, as described in Section 3.3.3. The TRE at the prostate can be estimated using the linear regression from Section 3.5.

Gold Standard Segmentations

The four MRI data sets were manually segmented, by the author, by fitting a spline to points on the bone soft tissue boundary for each slice. This was not done by a radiologist as there is no accepted protocol to do this. For data set XMR-01 and data set XMR-02 matching CT data was available, so this was was used as a reference to inform the manual segmentation process.

Model Based Segmentation

Each data set was first aligned to the shape model with a 9 DOF affine transform ($T_{SM \Rightarrow M}[nd9]^{-1}$). This registration was optimised using a gradient descent optimiser using NMI as a similarity measure between the two images.

As all the images in the shape model have already been placed into affine alignment, the same affine transform can be used to put the MRI into alignment with any of the shape model data sets. This is useful if the first attempt at the affine alignment fails. After the affine alignment the aligned data set is compared with the average shape of the shape model by a manual visual inspection. If the alignment is poor the registration can be repeated to a different shape in the model training set. Table 3.8 shows which data sets from the shape model each MRI set was aligned to.

Once the MRI image had been aligned to the mean shape of the model, the model was fitted to the

Data Set	Patient-01	Patient-02	XMR-01	XMR-02
Average Boundary Error (mm)	1.45	3.08	2.20	1.70

Table 3.9: Average boundary errors for shape model based segmentation versus manual segmentation of MRI data

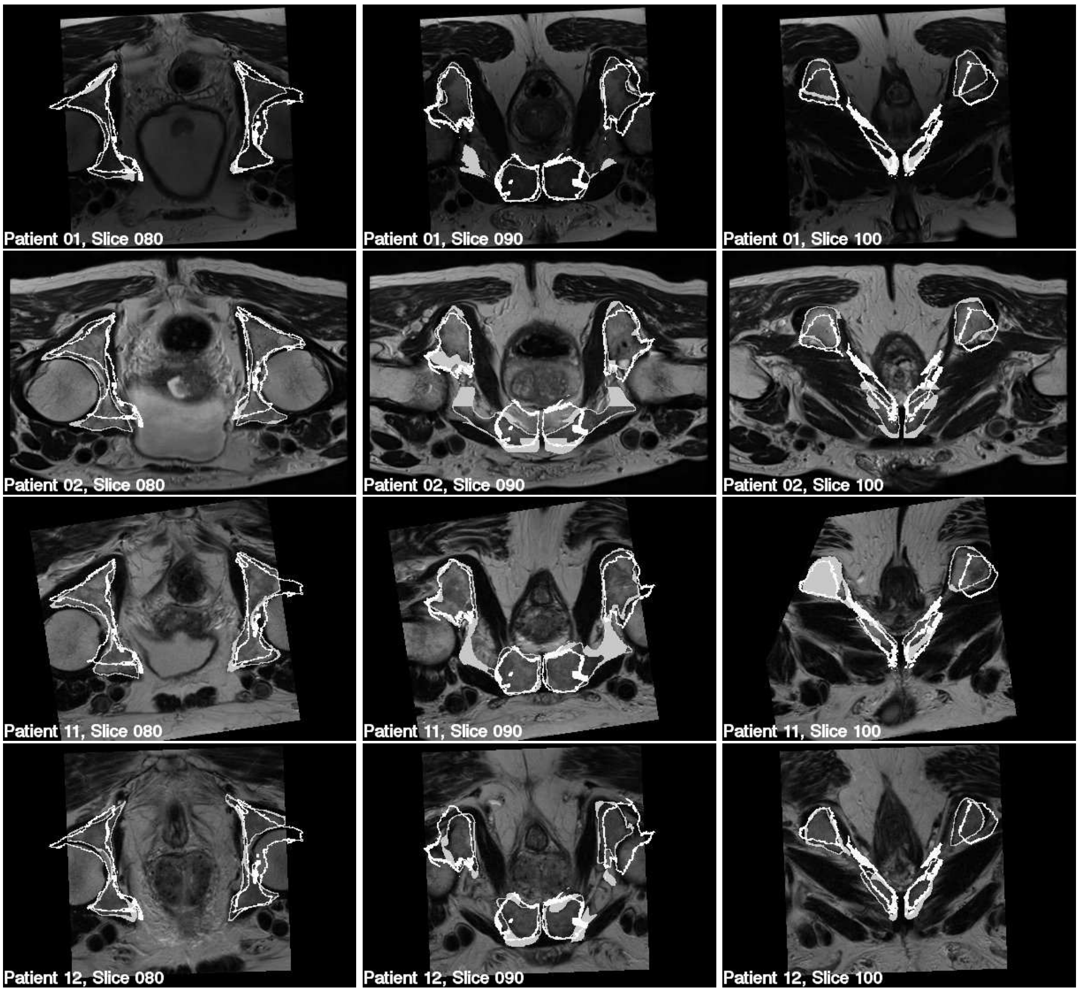
Patient 01, Slice 080
Patient 01, Slice 090
Patient 01, Slice 100
Patient 02, Slice 080
Patient 02, Slice 090
Patient 02, Slice 100
Patient 11, Slice 080
Patient 11, Slice 090
Patient 11, Slice 100
Patient 12, Slice 080
Patient 12, Slice 090
Patient 12, Slice 100

Figure 3.10: Example slices from each of the four MRI data sets. The manual segmentations are shown in grey, while the shape model propagated segmentations are shown in white. Patient-02 is notably worse than the others. It is also interesting that all data sets fail in the same region to the top right of slice 100, suggesting that this region is not well described by the shape model.

MRI using the differential evolution optimiser as per Section 3.2.2

3.7.4 Results

Each of the four data sets was segmented using the shape model. Error was measured in terms of boundary error across the pubic region. Table 3.9 presents the boundary error results for the four data sets. Figure 3.10 presents pictorial results of example slices for the segmentation.

The results presented indicate that the segmentation method works well. The boundary errors are

Data Set	Predicted TRE	Lower 95 Percent Confidence TRE	Upper 95 Percent Confidence TRE
Patient-01	3.38	2.65	4.11
Patient-02	8.14	4.64	11.63
XMR-01	5.56	3.67	7.45
XMR-02	4.11	3.06	5.16
Average	5.29	3.50	7.08

Table 3.10: Predicted TRE at the prostate surface for the four data sets based on fitting a linear regression for TRE vs pubic boundary error for results obtained using full pelvic CT data. Upper and lower limits of the predicted values are also given for a 95% confidence interval

slightly larger than those observed for the CT data in Section 3.4. This may be chance or more likely there is a small optimisation error added to the shape modelling error.

The boundary error results can be converted to prostate TREs using the linear regression from Section 3.5. The results of this are shown in Table 3.10.

3.8 Projection of Prostate TRE to on Screen Error

3.8.1 Aim

Our goal is to present the error due to segmentation as an on screen error. Here we project the prostate TRE onto the screen.

3.8.2 Method

The projection method described in Section 1.6.2 is used to project a point at the apex of the prostate onto the endoscope screen. The apex point is perturbed in two ways, firstly by an isotropic Gaussian distribution with an RMS of 5.29 mm as per Table 3.10. This assumes that the registration errors are isotropic and normally distributed. As the figure of 5.29 mm is based on an extrapolation of our results we cannot test the effect of this assumption directly. However we can test the effect of this assumption using our results from Section 3.4. Here we have both an RMS error figure and the estimates of $T_{SM \Rightarrow M}$ used to generate this. By comparing a projection of the prostate apex perturbed by a normally distributed error with those of the prostate apex transformed by the various estimates of $T_{SM \Rightarrow M}$ we can test whether the assumption of a normally distributed registration gives a valid projection error. We used the results of Method B for the masked anatomy for this (an RMS error of 2.81mm).

3.8.3 Results

Figure 3.11 shows the result of projecting an isotropic Gaussian with RMS error of 5.29 mm onto the screen. Figure 3.12 compares the isotropic normally distributed error assumption with our actual results. Figure 3.12 indicates that our assumption of an isotropic normally distributed registration may cause an underestimation of the on screen error. It is possible however that this is due to the limited size of our

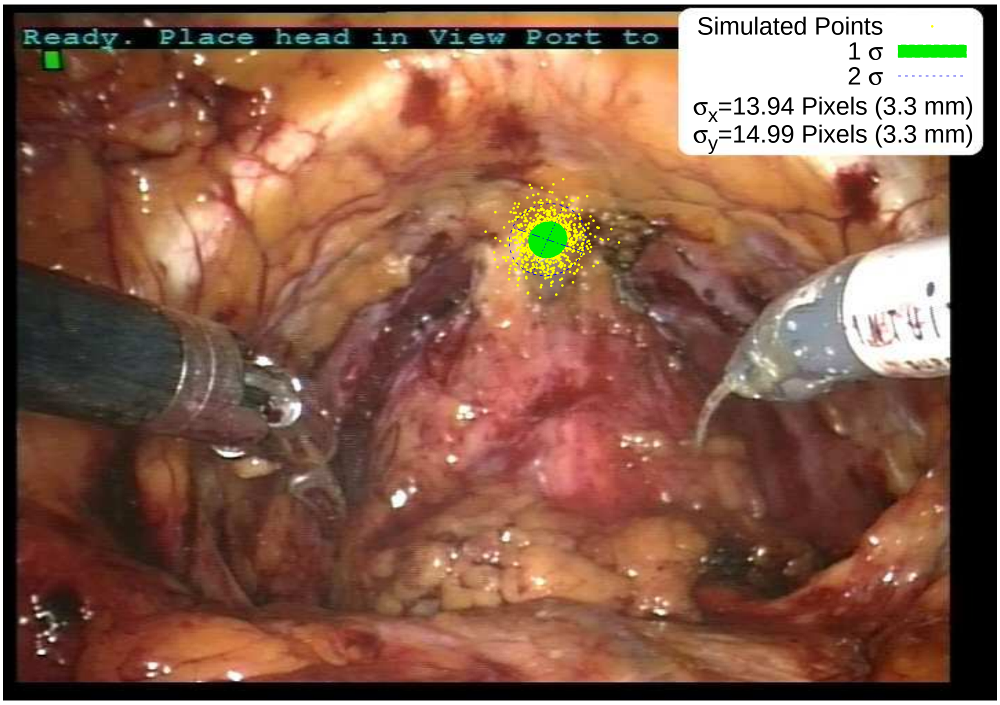
Ready. Place head in View Port to
Simulated Points
1 σ
2 σ
σ_x=13.94 Pixels (3.3 mm)
σ_y=14.99 Pixels (3.3 mm)

Figure 3.11: A landmark point (near the apex of the prostate) has been chosen and perturbed by a random Gaussian error in 3D equivalent to the bone segmentation error (5.29mm) 1000 times. The resulting projections give an indication of the on screen error due to this error source.

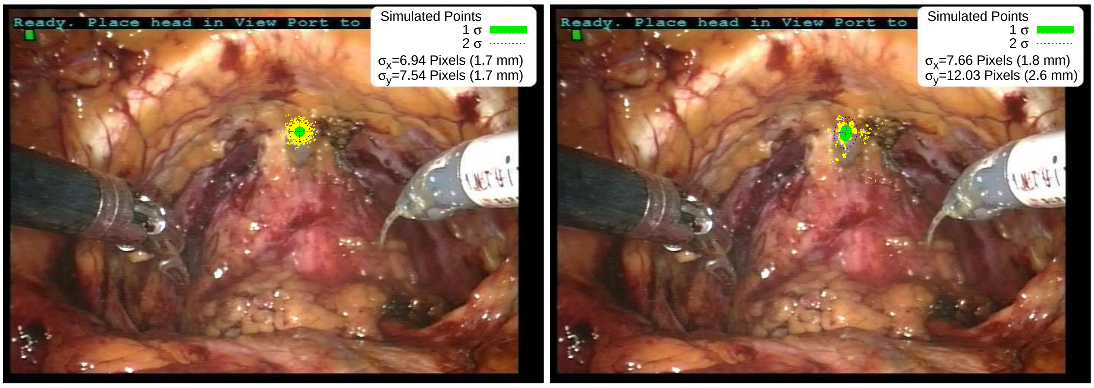
Ready. Place head in View Port to
Simulated Points
1 σ
2 σ
σ_x=6.94 Pixels (1.7 mm)
σ_y=7.54 Pixels (1.7 mm)
Ready. Place head in View Port to
Simulated Points
1 σ
2 σ
σ_x=7.66 Pixels (1.8 mm)
σ_y=12.03 Pixels (2.6 mm)

Figure 3.12: On the left a landmark point (near the apex of the prostate) has been chosen and perturbed by a random Gaussian error in 3D equivalent to the bone segmentation error (2.8mm) from Section 3.4 2100 times. On the right the 2100 estimates of $T_{SM \Rightarrow M}$ used to estimate the segmentation error have been used to perturb the point directly. It is apparent that the two methods yield substantially different results.

data set. What is shown on the right of Figure 3.12 is not a projection of a single distribution of 2100 samples. Rather it is a superposition of 21 distributions, each of 100 samples, representing the 21 data sets used to build the shape model. If the anisotropy is caused largely by the model fitting error, then the results in Figure 3.12 are based on an effective data size of 21. It is possible that this has caused the distribution to appear more anisotropic than it actually is. The addition of more data may reduce the anisotropy. We cannot, however, rule out the possibility that the segmentation and registration methods cause an anisotropic TRE at the prostate. In the absence of more data this must be left as future work. For the purposes of this thesis we will assume that the error due to the segmentation of the of MRI images with a shape model can be modelled as an isotropic normally distributed error with standard deviation 5.3 mm.

3.9 Discussion

3.9.1 Comparison of Pelvic Shape Model with Results from the Literature

Using the linear regressions between Dice's coefficient and prostate TRE and boundary error and prostate TRE from Section 3.5 it is possible to compare our results with those from the literature. When doing this it should be remembered that the correlations are not strong (with correlation coefficients approximately 0.5), so the comparisons must be treated with caution. Taking the Dice data in Figure 3.7 and extrapolating the result for a Dice coefficient of 0.94 as achieved on knee MRI by Fripp et al. (2007) gives an RMS TRE very near zero. In reality there is a lower limit of around 0.6 mm to the TRE that can be achieved due to the registration algorithm, which is discussed in detail in the next chapter. This suggests that (for the knee at least) model based segmentation of MRI is accurate enough to render manual segmentation or a separate CT scan unnecessary. However there are two significant differences between Fripp et al. (2007) and the work here, firstly the knee (more specifically the distal end of the femur) is a less complex shape than the pelvis, and secondly the MRI acquisition protocol used was better suited to bone delineation. Schmid and Magnenat-Thalmann (2008) quote a boundary error of 1.44 mm for a shape model based pelvic bone segmentation from MRI, which would yield a TRE of 3.4 mm, slightly better than our results. Redrawing Table 2.5 with equivalent results for this study yields Table 3.11. The results presented here are all boundary surface errors, calculated by searching through the shape model derived boundary surface segmentation for boundary points, then locating the nearest boundary point in the corresponding gold standard segmentation. The published results are similar to those found using the shape model.

3.9.2 Correlation of Optimisation Metric with Boundary Error

For the optimiser to work well three things must occur. Firstly the pelvic shapes seen in the MRI data must be somewhere in the shape space of the statistical shape model, this was tested in Section 3.4. Secondly the differential evolution optimiser must effectively optimise the similarity metric (NMI), this was tested is Section 3.6. Lastly, the similarity measure must correlate well with the boundary error measure used for evaluation.

To see how well NMI correlates with boundary error for our MRI data the optimisation process was

Paper	Result from Literature	Our Result
Lamecker et al. (2004)	2.4 mm RMS Boundary Error without source in model	2.43 mm
Yao and Taylor (2003)	1.27 mm Average Boundary Error.	1.64 mm
Chan et al. (2004)	3.90 mm RMS Boundary Error	2.43 mm
Schmid and Magnenat-Thalmann (2008)	1.81 mm RMS Boundary Error	2.43 mm

Table 3.11: Comparison of shape model results with published studies using shape models on pelvic CT data, except for Schmid and Magnenat-Thalmann (2008), based on MRI and including a local deformation step. Given the very different methodologies used to construct and fit the shape models, the broad agreement between them is encouraging.

Data Set	Pearson Correlation Coefficient	Boundary Error
Patient-02	-0.292	3.08
XMR-01	0.574	2.20
XMR-02	0.790	1.70
Patient-01	0.844	1.45

Table 3.12: Data sets arranged in order of correlation between the image similarity measure and the average boundary error. The degree of positive correlation has a dominant effect on the final error measure.

monitored to see the progress in similarity measure and boundary error. The differential evolution algorithm works by progressively breeding a population of 60 candidate solutions to improve the similarity measure. By also evaluating the boundary error metric for each of the candidate solutions, it is possible to plot how both metrics change during optimisation. Figure 3.13 presents these plots for each MRI data set. From Figure 3.13 it can be seen that in each case the optimisation algorithm steadily improves the NMI . In three cases the boundary error metric improves alongside the image similarity measure. For Patient-02 however the boundary error increases. The correlation between NMI and boundary error varies between data sets. The strength and direction of the correlation between the two metrics was measured using Pearson's correlation coefficient. Table 3.12 indicates that the sign and magnitude of the correlation between NMI and boundary error is a good predictor of the boundary error achieved.

Understanding the reasons for this variation in correlation between data sets could yield a more effective segmentation algorithm, or an improved similarity measure for MRI to CT registration. The quality of the manual segmentation being used as a gold standard will have a significant effect on our results. If the manual segmentation is wrong the correlation would appear worse as the optimiser would

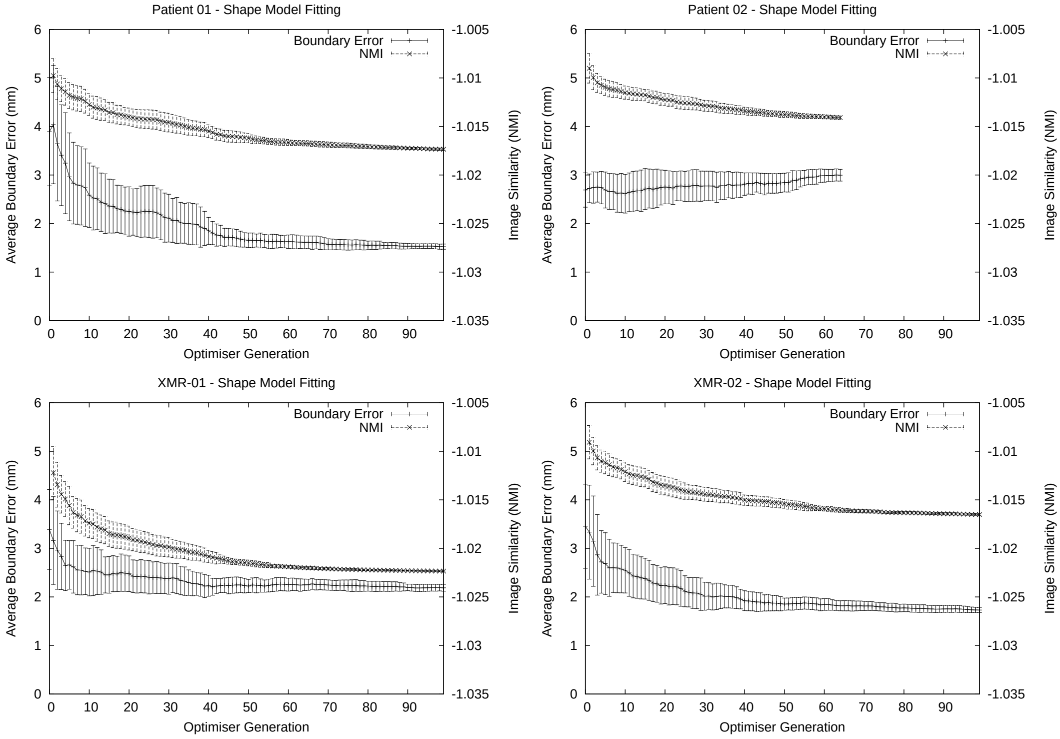

Figure 3.13: Plots showing the image similarity measure being optimised (normalised mutual information) and the evaluation metric (average boundary error in the pubic region) over the course of optimisation for the four data sets used. In all cases there is a correlation between the image similarity measure and the boundary error, this is negative for Patient-02. Table 3.12 gives the correlation coefficients for these correlations.

Error Source	RMS TRE (mm)
Optimisation	4.48
Shape Not in Model	2.81
Registration Algorithm	0.23

Table 3.13: Summary of RMS TRE at prostate surface due to the three errors in using a shape model based segmentation algorithm.

move the segmentation to a more correct value. It would be worth repeating these manual segmentations and comparing them between multiple users before looking for other causes of the variable correlation coefficients.

3.9.3 Result Summary

Table 3.13 summarises the key results from this chapter. The average predicted TRE shown in Table 3.10 is made up of three components. The registration algorithm error, the optimisation error, and the error induced by the shape model. The registration algorithm error is the performance of the registration algorithm when matching identical images. This was determined by repeating the experiment from Section 3.4, but registering the simulated ultrasound to the gold standard segmentation, rather than a model derived segmentation. The RMS TRE for this was 0.23mm. This registration error can be subtracted from the TRE squared error to give an error due to segmentation alone of $\sqrt{5.29^2 - 0.23^2} = 5.29mmRMS$.

The shape model error can be taken from the Section 3.4 and is 2.81 mm. Again this can be subtracted from the total error to yield the error due to shape model optimisation. $\sqrt{5.29^2 - 2.81^2} = 4.48mmRMS$

3.10 Conclusion

We have presented a novel cross-modality, inter-patient, non rigid registration algorithm. The method uses a shape model to constrain the allowable deformations. We have shown that using this algorithm it is possible to fit pelvic CT from a separate patient to pelvic MRI. Doing this as opposed to a manual segmentation of the MRI or taking a separate CT scan induces a registration error observable at the prostate surface. This error has been quantified here at 5.29 mm RMS. This error would appear to the surgeon as shown in Figure 3.14

Whether the error of 5.29 mm is small enough for for clinical use is unknown. A system of this accuracy would certainly be useful for showing the surgeon the general location of specific tumours. Based on this the surgeon could decide on the size of the margin around each side of the prostate. Similarly it may help the surgeon locate important anatomy such as the neuro-vascular bundles. The surgeon could not rely on the system when making incisions however, and would instead have to base their decision on the visible anatomy. If it were shown that this error was significant in the overall system it may be possible to reduce it by improving the shape model. In the first instance this could be done by adding more data sets to the training set.

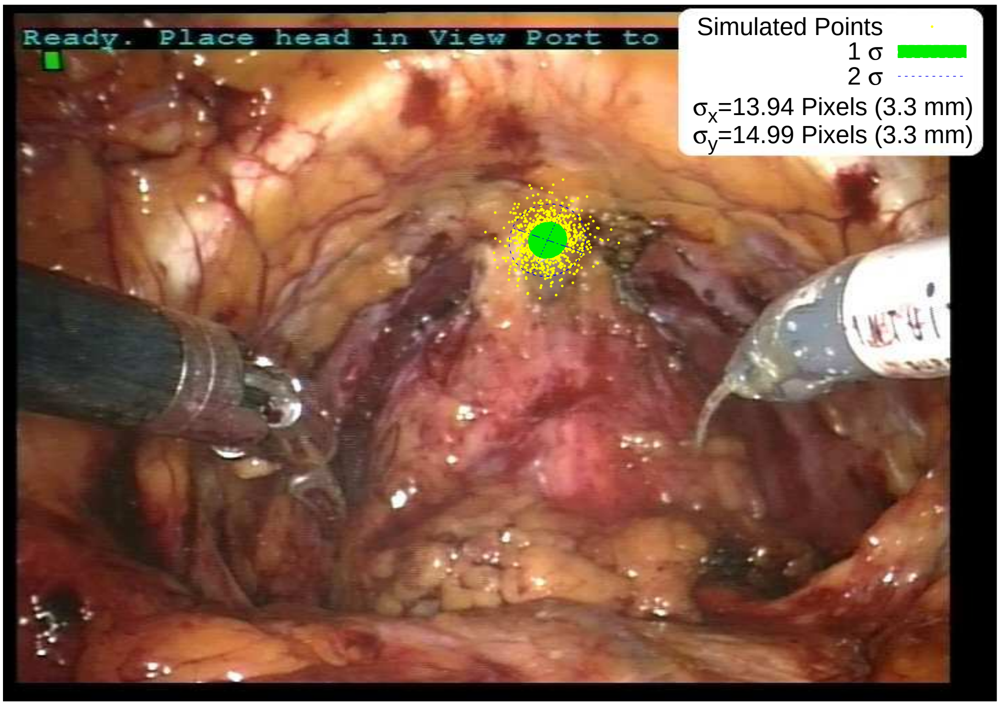
Ready. Place head in View Port to
Simulated Points
1 σ
2 σ
σx=13.94 Pixels (3.3 mm)
σy=14.99 Pixels (3.3 mm)

Figure 3.14: A landmark point (near the apex of the prostate) has been chosen and perturbed by a random Gaussian error in 3D equivalent to the bone segmentation error (5.29mm) 1000 times. The resulting projections give an indication of the on screen error due to this error source.

Chapter 4

Registration of Ultrasound to Fitted Bone Image

This chapter details the methods used to register the preoperative MRI data to the patient in theatre, that is determining $T_{M \Rightarrow O}$ (see Figure 1.3). The method we use for this is an image to image registration using ultrasound images of the patient gathered with the patient in the operating position. As discussed is Chapter 2, rather than attempting to register the ultrasound images directly to the MRI image we register the ultrasound to a CT image from a second patient (the shape model) that has been warped, using $T_{SM \Rightarrow M}$ determined in Chapter 3, to the fit the MRI image. Therefore we actually find $T_{SM(nrr) \Rightarrow O}$ in this chapter, which is identical to $T_{M \Rightarrow O}$ as the model is fitted to the MRI image.

As discussed in Chapter 2, registration of ultrasound images of bone to CT images of bone has been done previously by several authors. The method we use is derived from that of Penney et al. (2006), however there are several key differences. Whereas Penney et al. (2006) register to a CT image of the patient, so the anatomy being imaged by the CT and ultrasound is identical, we register to CT from a second patient that has been warped. This results in a pseudo CT image that is not a perfect match for the anatomy being imaged with the ultrasound probe.

We have also implemented the registration algorithm in a way that minimises the image processing required in theatre. We have a limited time period to do the registration in theatre (approximately 25 minutes), so we have attempted to streamline the algorithm as much as possible. The key difference is that rather than storing the ultrasound images as a set of discrete slices, the ultrasound images are first converted to a point cloud. By discarding points below a set intensity threshold it is possible to significantly reduce the computational load.

As discussed in Chapter 1, calculation of $T_{SM(nrr) \Rightarrow O}$ is broken into two stages. Individual ultrasound images are first converted to point clouds using $T_{UI \Rightarrow O}$. The fitted model is then registered to the ultrasound point cloud to calculate $T_{SM(nrr) \Rightarrow UI}$. In this chapter we first describe the algorithms and equipment used for these two steps. We then perform a series of experiments using the algorithms, with the aim of determining the accuracy with which $T_{SM(nrr) \Rightarrow O}$ can be found.

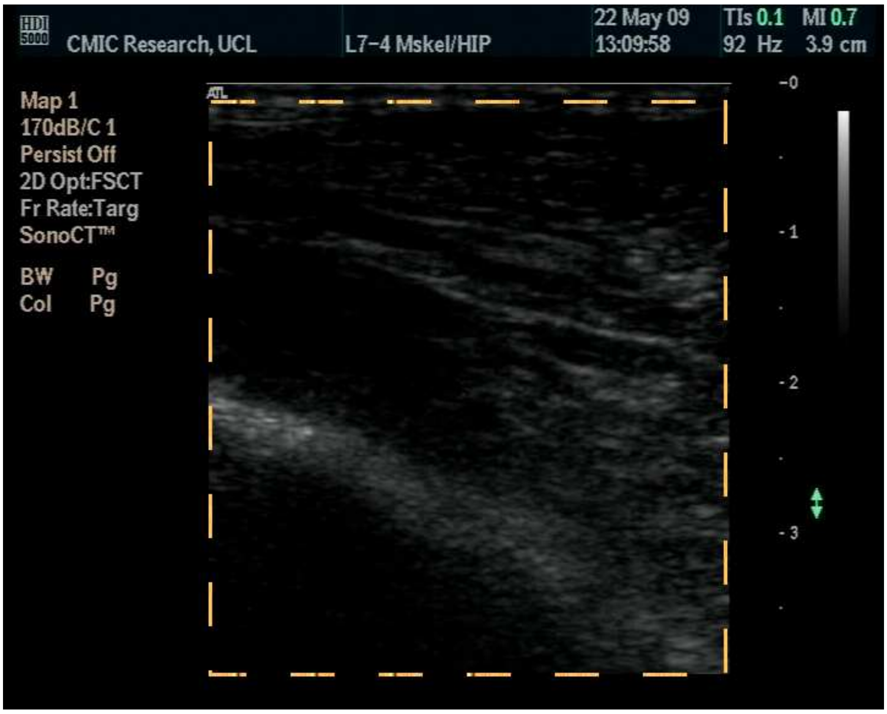

Figure 4.1: A typical ultrasound slice collected in theatre, as output from the ultrasound scanner and captured via a frame grabber. Image size is 720 by 576 pixels, the ultrasound data is in grey-scale with pixel intensity from 0 to 255. The dashed line around the data frame shows where the image is cropped in the first stage of image preprocessing. The face of the ultrasound probe is at the top of the image.

4.1 Ultrasound Image Acquisition and Processing

4.1.1 Ultrasound Image Acquisition

Ultrasound images of the patient's pelvic bone were acquired using a Phillips HDI5000 ultrasound scanner. The machine's default settings for musculoskeletal imaging of the hip were used. The images were captured using a frame grabber. This frame grabber outputs interlaced images 720 pixels wide by 576 pixels high, with pixel intensities from 0 to 255. An example of one of these images is shown in Figure 4.1.

The position of the slice relative to the optical tracking system $T_{UI \Rightarrow O}$ is determined using a set of IREDs, shown in Figure 4.2. The tracking data from the Optical tracking system is saved alongside the corresponding ultrasound image. There is a lag of approximately 375 ms between the tracking information being recorded and the image being captured. Therefore every effort was made to capture images while the probe was stationary. However the probe is hand held, so some movement is unavoidable.

4.1.2 Ultrasound Image Pre-Processing

Once the image has been acquired the following processing is performed. The images are first cropped leaving only the data window, shown on Figure 4.1. The 3 mm of tissue closest to the transducer face is also removed to prevent artefacts due to skin coupling. The image is now converted to a "bone edge probability" image, this is based on the work of Penney et al. (2006). We do this for two reasons, both

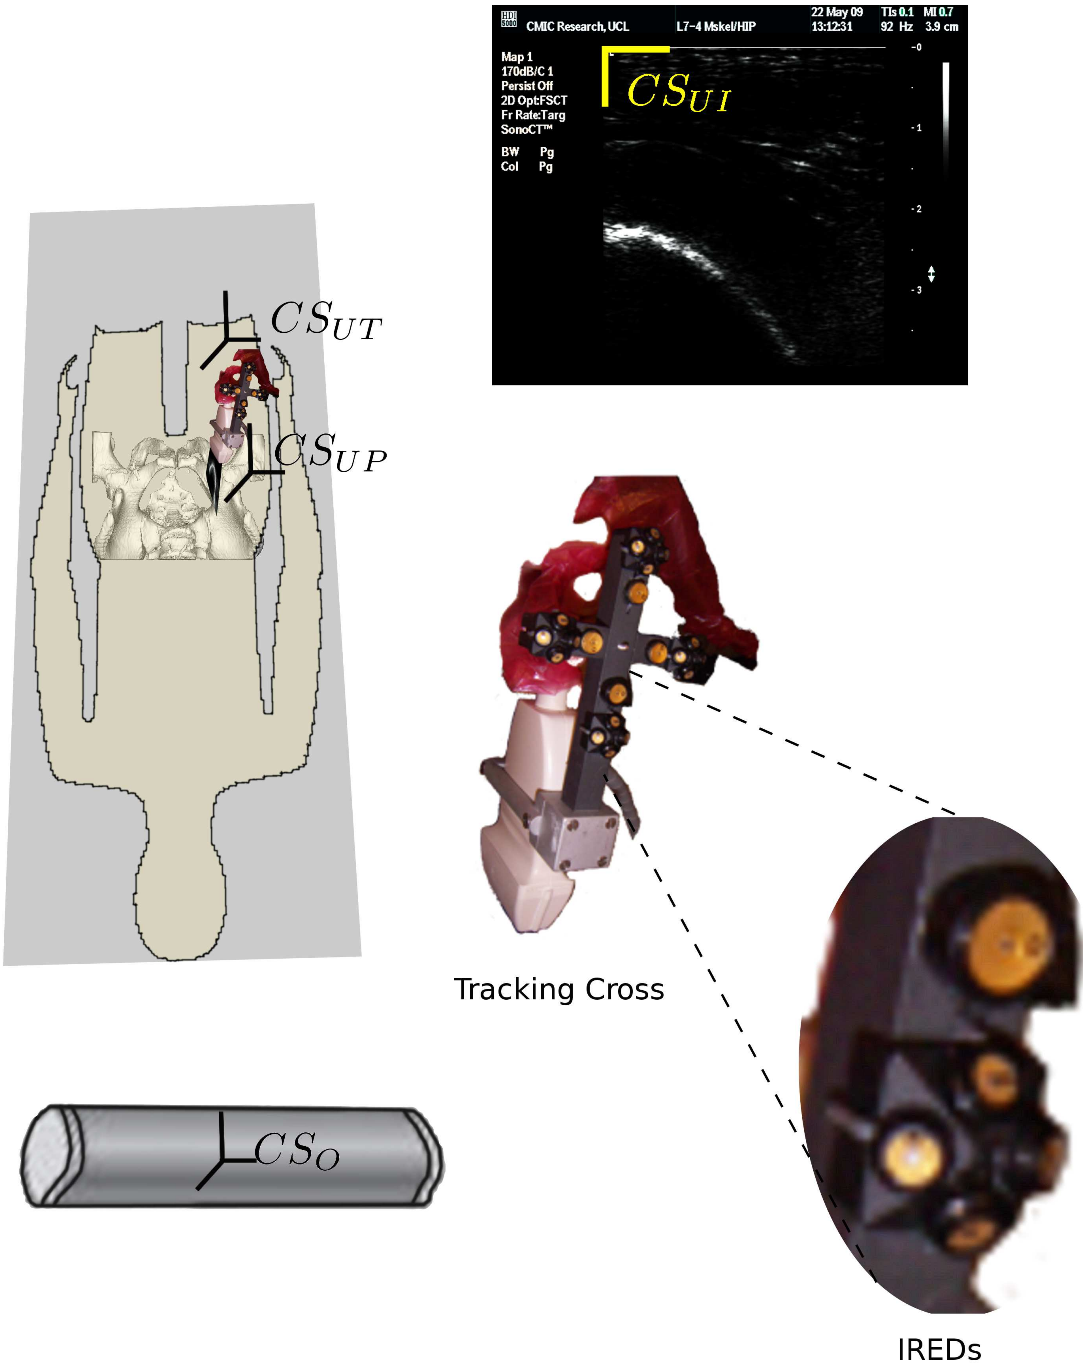
CMIC Research, UCL
L7-4 Mskel/HIP
22 May 09
13:12:31
TIs 0.1
MI 0.7
92 Hz
3.9 cm
Map 1
170dB/C 1
Persist Off
2D Opt:FSCT
Fr Rate:Targ
SonoCT™
BW Pg
Col Pg
CS_{UI}
CS_{UT}
CS_{UP}
CS_O
Tracking Cross
IREDs

Figure 4.2: The ultrasound probe, attached tracking cross and the Optotrak tracking system. 20 IREDs are attached to the cross.

due to our need for a fast running image to image registration algorithm. The first is that as image processing has been used to enhance the same features (the bone edges) the registration is closer to being an intra-modal problem. It should therefore be possible to use computationally efficient similarity measures to drive the registration. To understand the second reason an understanding of our registration algorithm is required. This will be introduced in full in Section 4.4, but the key points are outlined here.

After processing the ultrasound image will be stored not as a set of slices but as a set of 3D points with corresponding intensity values. To calculate the image similarity measure the algorithm reads through every point in the ultrasound derived bone edge probability (BEP) image, transforms it to the BEP image derived from the MRI image using an estimate of $T^{-1}_{SM(nrr)\Rightarrow UI}$ and finds the intensity of the nearest voxel in the MRI derived BEP image. The ultrasound and MRI derived BEPs are multiplied together and the sum of these products is used as the similarity measure. Points in the ultrasound image with a BEP of zero will have no effect on the registration process, the BEP product will be zero regardless of the value of the matching corresponding voxel in the MRI derived BEP image. Therefore to reduce computational time all pixels with zero intensity value are discarded during the preprocessing of the ultrasound image. Converting the images to bone edge probability images greatly reduces the number of non zero pixels. Based on the data for Cadaver-01, see Appendix A, the reduction is as follows. The cropped ultrasound images have $288 \times 413 = 118944$ pixels. On average 25400 (21%) of these are non zero. After converting to a bone edge probability image each image has an average of 3312 (3% of the original) non zero pixels. These reductions in the number of non zero points will translate directly into a faster registration process.

The conversion to bone edge probability images follows the procedure laid out by Penney et al. (2006). This is described below and shown pictorially in Figure 4.3. Figure 4.3 shows the procedure on two data sets, one taken from a plastic pelvis scanned in a water bath (the phantom data) and one taken from a patient scanned in theatre.

An “artefact threshold” intensity value is first defined. How to choose the threshold value will be discussed later in this chapter, see Section 4.4.3. The ultrasound image is scanned column by column upwards from the bottom of the image (distant from the transducer face). At the first pixel that exceeds the value of the artefact threshold the scan is stopped and the row number ($N_{Art,i}$) recorded, where i is the column number and row number 0 corresponds to the top of the image. If no pixels in the column exceed the threshold then $N_{Art,i}$ is set to zero. Two images are now created. The first is a “distance to artefact” image (row 2 of Figure 4.3). The intensity of each pixel in this image is equal to $N_{Art,i}$ for the corresponding column minus the pixels own row. Pixels with a negative value for the distance to artefact are set to zero. The term artefact is used to indicate that pixels more distant from the scan face than the bright pixels so detected are likely to be due to scanning artefacts rather than echoes off genuine anatomy. This is due to the almost total reflection of ultrasound from the bone surface. The second image is a masked and blurred version of the raw intensity image. The raw ultrasound image is first smoothed with a Gaussian kernel to reduce speckle artefacts, then pixels with a distance to artefact of 0 value are set to 0. The result of this is two metrics for every pixel in the image. These metrics are used to convert the images to bone edge probability images using a 2D lookup table.

The look up table is produced beforehand using a training set of manually segmented images, representative of the images to be processed in theatre. Our look table is constructed using 32 images taken from the phantom data set. The images are first manually segmented by picking a set of points along the bone surface. A spline is fitted to the points and pixels falling on the spline are assigned an edge probability of 1. This forms a binary segmentation image with pixels not on the the edge having pixel intensity of 0 and those on it a pixel intensity of 255. We then blur the this segmentation image. We do this because we do not expect our manual segmentation to be perfect, therefore it is realistic to say that a pixel near the manually segmented edge has a higher probability of being the edge than a pixel distant to the edge. Blurring the manual segmentation also has the effect of increasing the number of pixels used in building the look up tables, allowing fewer manually segmented images to be used to build the look up tables.

To build the look up tables the raw ultrasound image is now processed using the procedure shown in Figure 4.3. Each pixel in the image now has three metrics associated with it, the distance to artefact, the masked and blurred intensity, and the bone edge probability derived from the manual segmentation. The lookup table is simply a 2 dimensional histogram of the edge probabilities. We used 32 bins each for the distance to artefact and intensity when constructing the look up tables. Figure 4.4 shows two example look up tables made using manual segmentations of the plastic phantom data. These are the look up tables used in Figure 4.3.

4.2 Ultrasound Image to 3D Points

The ultrasound images are now stored as 2D pixel coordinates with corresponding bone edge probabilities. The next stage is to transform these 2D points to 3D points in the coordinate system of the Optotrak. This is done using $T_{UI \Rightarrow O}$ from Figure 1.3. As detailed in Figure 1.3 $T_{UI \Rightarrow O}$ is the product of the ultrasound calibration transforms ($T_{UI \Rightarrow UP}$ and $T_{UP \Rightarrow UT}$) and the ultrasound tracking transform ($T_{UT \Rightarrow O}$). As the tracking transform is required by the calibration algorithm we will first detail the ultrasound tracking process. Following this we detail the ultrasound calibration process.

4.2.1 Ultrasound Probe Tracking

The ultrasound probe is tracked using a set of 20 IREDs rigidly attached to the use probe on an aluminium cross. This is shown in Figure 4.5. The optical tracking system triangulates the position of each of the 20 IREDs and outputs them to a computer. The 20 IRED positions captured for a given ultrasound frame (i) can be registered to the IRED positions for any other frame (j) using well established point to point registration methods. Here we use an orthogonal Procustes formulation and find the best fit registration using singular value decomposition [Fitzpatrick et al. (2000)]. This gives a rigid body transform from one point set to the other, $T_{Frame_i \Rightarrow Frame_j}$. In order to transform the points into the Optotrak's coordinate system it is necessary to know the rigid body transform from one set of IRED coordinates to the Optotrak's coordinate system, we call this frame the reference frame, $Frame_{Ref}$, and use it to define the tracking rigid body's coordinate system. $T_{UT \Rightarrow O}$ can then be defined as in Equation 4.1.

Method	Phantom Data	Patient Data
Cropped Ultrasound Slice	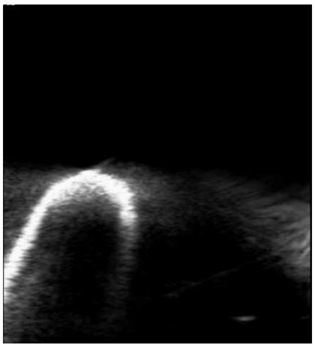	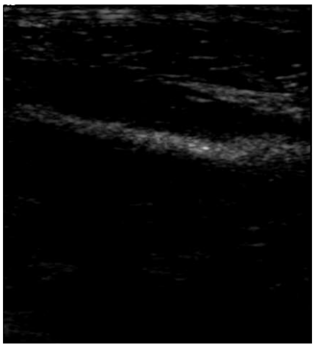
The image is scanned upwards column by column. The distance to artefact is the row number of the first pixel that exceeds the artefact threshold minus the pixels row number.	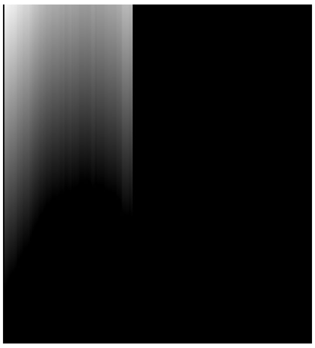	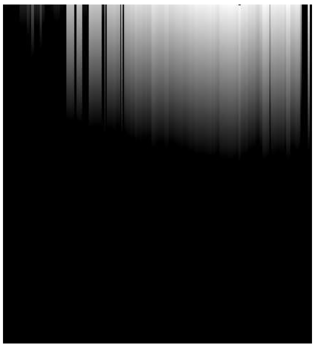
Pixels with zero distance to artefact are masked and the remaining unmasked pixels are blurred.	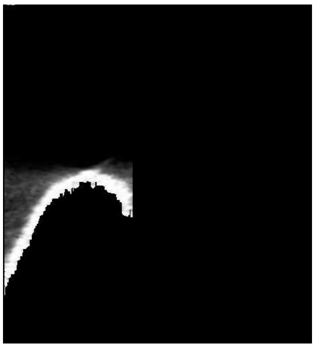	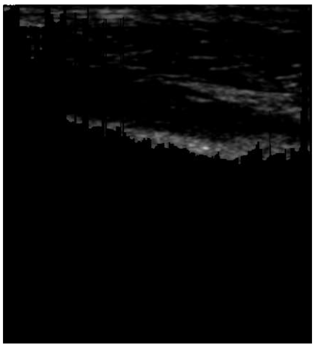
The blurred intensities and distance to artefact values are sorted into 32 bins each and combined using a look table, see Figure 4.4, to give an edge probability image, shown here.	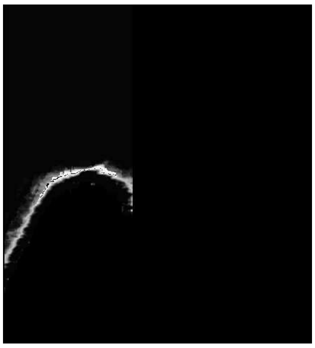	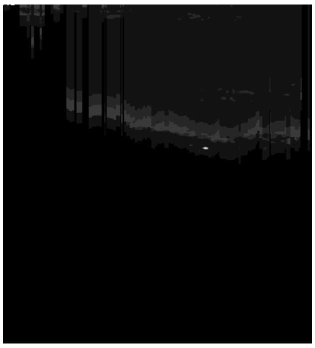

Figure 4.3: Details of the processing steps used to convert an ultrasound slice to a bone edge probability image. The ultrasound transducer face is at the top of the images.

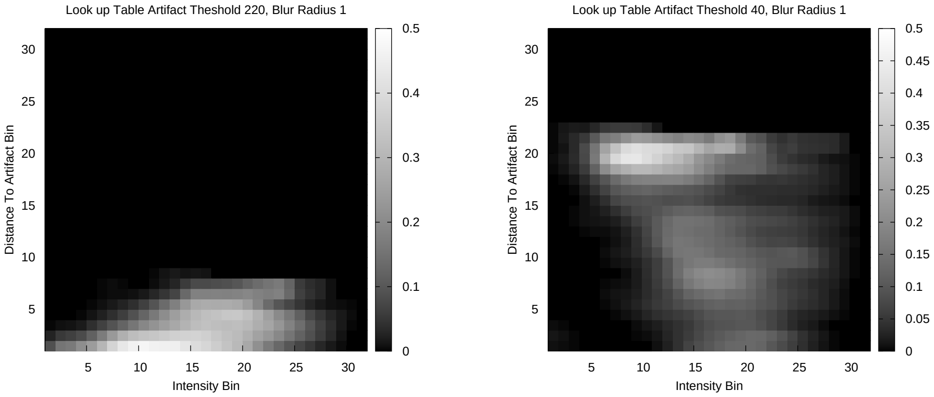

Figure 4.4: The probability look up tables used in the processing of the ultrasound slices seen in Table 4.3. Both tables use the same set of manually segmented images of the plastic phantom. By changing the artefact threshold used we can get very different look up tables. The table on the right uses an artefact threshold of 40, while the one on the left uses an artefact threshold of 220. The table on the left does a good job of segmenting the phantom data, where the bone edges were very bright and there were no reflections from soft tissue. The table on the right works better for segmenting patient data where the bone edges are not as bright and there are significant amounts of soft tissue.

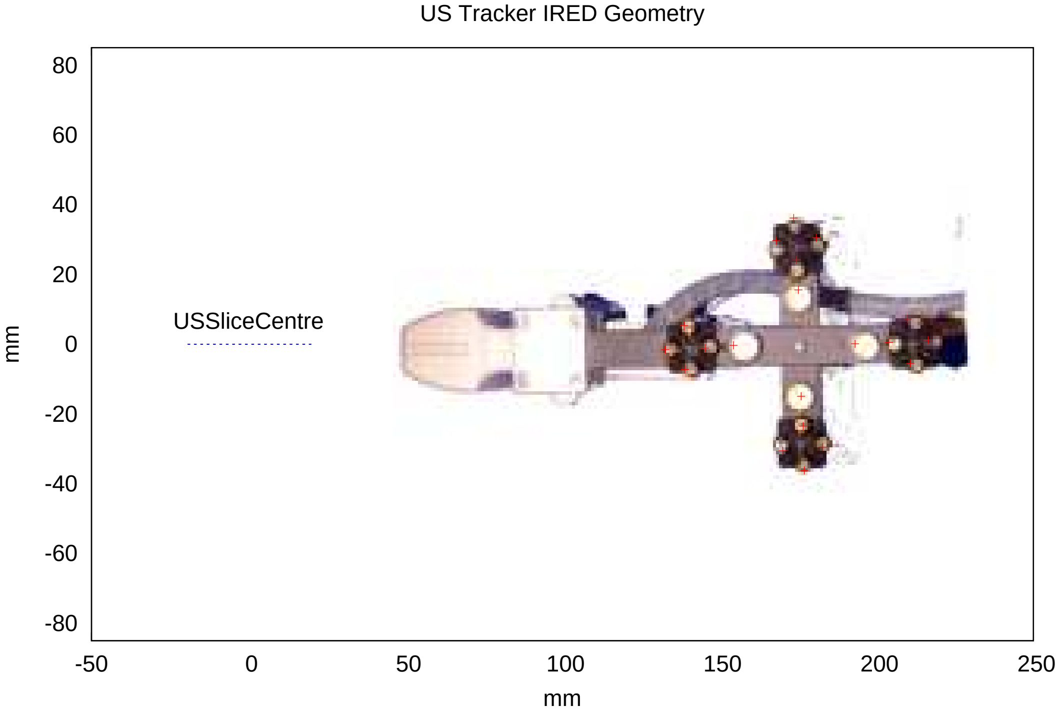

Figure 4.5: Geometry of the tracking cross used to localise the ultrasound slices.

$$T_{UT \Rightarrow O}[Frame_i] = T_{UT \Rightarrow O}[Frame_{Ref}] \times T_{Frame_i \Rightarrow Frame_{Ref}} \tag{4.1}$$

We define $T_{UT \Rightarrow O}[Frame_{Ref}]$ as the identity matrix, then $T_{UT \Rightarrow O}$ becomes simply $T_{Frame_i \Rightarrow Frame_{Ref}}$

Defining the Reference Frame

How we define the reference object will impact on the tracking accuracy of the system. Though it would seem trivial to define the reference object using an arbitrary frame, as being a rigid body the geometry of the IREDs should not change, this may not yield accurate results. The apparent position of an IRED tracked using the Optotrak will vary depending on its position and angle relative to the tracking cameras' len's surface normals, [Wiles et al. (2004)]. In order to maximise tracking range the IREDs on the tracking cross are arranged at a range of different angles. Therefore the apparent geometry of the tracking cross will vary as the cross moves in relation to the Optotrak camera unit. A reference object defined with the cross front on to the tracking camera may give poor results when the cross is at a different angle to cameras. To attempt to overcome this we used the following procedure.

In order to quantify the performance of a reference object we define the IRED registration error (IRE), this is analogous to the fiducial registration error (FRE) [Fitzpatrick et al. (1998)] commonly used in image guided interventions. Equation 4.2 defines the IRE.

$$IRE = \sqrt{\frac{1}{N}\sum_{j=1}^{N} |T_{Frame_i \Rightarrow Frame_{Ref}} \times X_{i,j} - X_{Ref,j}|^2} \tag{4.2}$$

Where N is the number of IREDs (20 in this case), $X_{i,j}$ is the position of IRED j in frame i and $X_{Ref,j}$ the position of IRED j in the reference frame. $X_{i,j}$ and $_{Ref,j}$ are both 3D position vectors. Next we collect tracking data for a set of frames that are representative of the range of positions and angles seen during the tracking application (ie. representative of scanning a patient's pelvic bone). These are registered to a reference frame and an average value of IRE for all frames is calculated. We now make the assumption that a reference object that gives a lower average IRE will necessarily give better tracking results. This assumption will not be correct in all situations, though we think it is valid for this application. We can now create an optimised reference object for our application as follows.

1. Collect a sample of frames representative of the tracking task.
2. Register all frames to a starting estimate of the reference frame.
3. Compute the IRE for each frame using equation 4.2 and take the average.
4. Transform all the IREDs $X_{i,j}$ to the reference coordinate system using the estimated values of $T_{Frame_i \Rightarrow Frame_{Ref}}$.
5. Find the average position of each of the transformed IREDs.
6. Use these average positions as a revised estimate of $X_{Ref,j}$
7. Repeat steps 2 to 5 using the new estimate of the reference frame. The average IRE will be reduced.

8. Repeat step 6 until the value of IRE converges.

For each frame of ultrasound data we can now define $T_{Frame_i \Rightarrow Frame_{Ref}}$. This is a rigid body transform that we represent as a 4×4 matrix. These can be used together with the calibration transforms to locate a pixel in an ultrasound slice image in 3D space. Before we move onto our calibration method we will first discuss our use of a slice rejection test.

Slice Rejection Test

The accuracy with which the singular value decomposition algorithm can find $T_{Frame_i \Rightarrow Frame_{Ref}}$ will vary from frame to frame. This will depend on two factors. Firstly how accurately the Optotrak can localise a given IRED, referred to as the IRED localisation error (ILE) and analogous to the fiducial localisation error (FLE) for a fiducial registration system [Fitzpatrick et al. (1998)]. The second factor is the number and the geometry of the IREDs that are tracked. Typically this will vary from frame to frame as some IREDs are obscured or at too great an angle to the camera to be tracked. The goal of a slice rejection test is to identify ultrasound images that are poorly tracked and discard them, to prevent them biasing the calibration routine or later on the registration algorithm.

For the purposes of this slice rejection test we have assumed that the ILE is an isotropic normally distributed error with constant standard deviation for all IREDs. The effects of this assumption are dealt with in greater depth in Chapter 5.

Our slice rejection test checks two parameters. The first of these is the IRE for the frame. It has been shown that the IRE for an individual frame cannot be used to predict the accuracy of the estimated transform for that frame ($T_{Frame_i \Rightarrow Frame_{Ref}}$) by Fitzpatrick (2009). Very high values of IRE would nonetheless make us question whether an equipment failure has created an erroneous measurement. The expected value of the IRE can be determined using a result by Sibson by way of Fitzpatrick et al. (1998). This is shown in Equation 4.3.

$$\langle IRE^2 \rangle = (1 - \frac{2}{N})\langle ILE^2 \rangle \tag{4.3}$$

If we assume that the IRE is a half normal distribution, (the absolute value of a normal distribution with mean 0 and standard deviation σ) then Equation 4.4 can be used to estimate its standard deviation.

$$\sigma_{IRE} = \frac{\langle |IRE| \rangle}{\sqrt{\frac{2}{\pi}}} \tag{4.4}$$

If the tracking system is functioning correctly 99.7 % of all values of IRE should be less that $3\sigma_{IRE}$. We will reject any slices with a IRE greater than this. This is conservative and will cause a small number of good slices to be rejected, but this is preferable to including erroneous data.

The second parameter is the number and distribution of visible IREDs. Fitzpatrick et al. (1998) provide an analytic method (equation 4.5) to estimate the expected value TRE given the expected value of FLE (analogous to ILE) and the configuration of visible fiducial markers (IREDs here). d_k is the distance of the target point from the kth principal axis of the fiducial configuration. f_k is the RMS distance of the fiducial markers from the kth principal axis. TRE is simply the error at a specified point when the estimated transform $T_{Frame_i \Rightarrow Frame_{Ref}}$ is applied. Therefore for our purposes the TRE can be used as a proxy for the error in determining $T_{Frame_i \Rightarrow Frame_{Ref}}$.

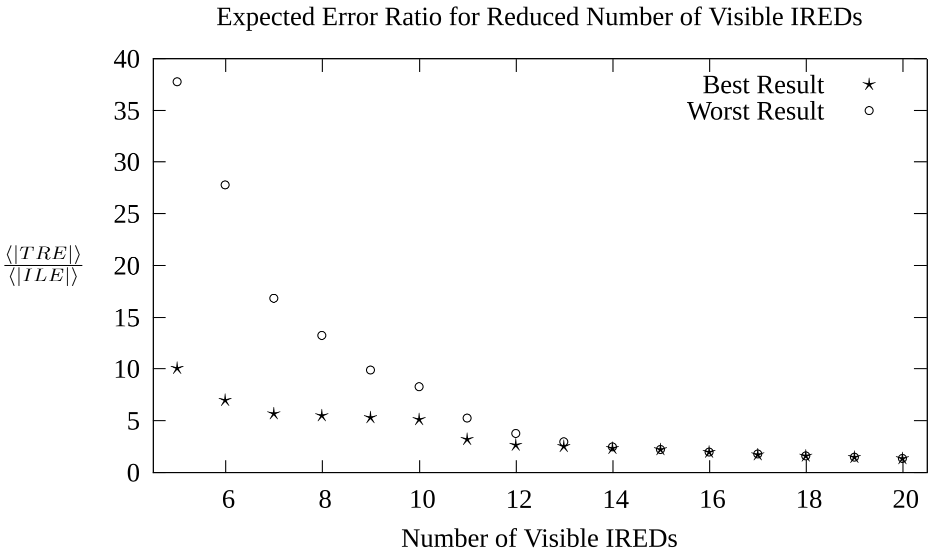

Figure 4.6: Best and worst case TRE per ILE versus number of IREDs visible for the ultrasound tracking cross. These were calculated using Equation 4.5 for all possible IRED configurations then sorted into best and worst case TREs.

$$\langle TRE^2 \rangle \approx \frac{\langle FLE^2 \rangle}{N} \left(1 + \frac{1}{3} \sum_{k=1}^{3} \frac{d_k^2}{f_k^2} \right) \tag{4.5}$$

Two implementations of the tracking software have been used in this thesis. The more recent of these provides a list of which IREDs are visible together with $T_{Frame_i \Rightarrow Frame_{Ref}}$. This enables $\langle TRE^2 \rangle$ to be calculated for each frame directly. Frames with $\langle TRE^2 \rangle$ above a present threshold were discarded. The earlier implementation of the tracking software only reported the number of visible IREDs visible, not their configuration. Therefore we performed the following analysis to determine a threshold number of visible IREDs below which frames would be discarded.

The ratio of $\langle |TRE| \rangle$ to $\langle |ILE| \rangle$ was determined for every possible combination of visible IREDs using equation 4.5. For a rigid body with 20 IREDs there are 2^{20} possible combinations. To calculate the TRE we have assumed a target point at the centre of the ultrasound slice. For each number of visible IREDs we then found the worst and best case ratio of $\langle |TRE| \rangle$ to $\langle |ILE| \rangle$, and plotted them in Figure 4.6. We elected to reject any frames with less than 11 IREDs, as below this the worst case errors become large. The relatively large increase in worst case error between 11 to 10 IREDs is due to the geometry of the cross. The tracking accuracy reduces substantially when IREDs are visible on only 2 of the four points of the cross.

Use of a Reference Rigid Body

In cases where the object being scanned (ie. the patient or the plastic phantom in this study) may move in relation to the Optotrak camera system, using the Opotrak's coordinate system for registration will lead to errors. These can be overcome by attaching a tracked "reference" rigid body to the object or patient to be scanned. Figure 4.7 shows a reference body attached to a plastic phantom pelvis. In this case $T_{M \Rightarrow O}$

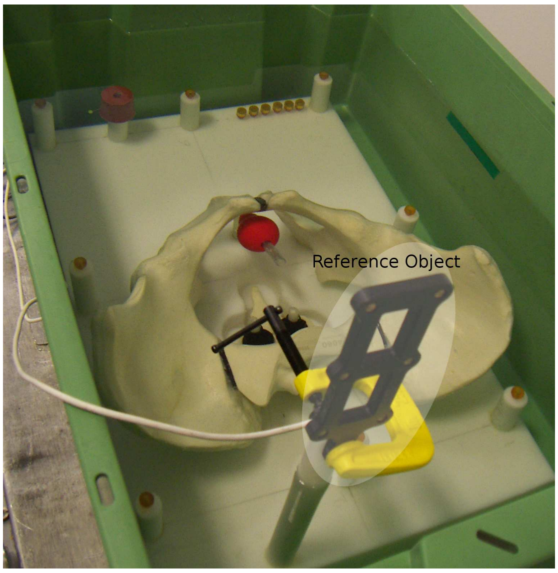

Figure 4.7: Geometry of the reference object used in acquisition of ultrasound slices from the phantom.

and $T_{UT\Rightarrow O}$ are replaced with $T_{M\Rightarrow R}$ and $T_{UT\Rightarrow R}$. $T_{UT\Rightarrow R}$ is found using equation 4.6.

$$T_{UT\Rightarrow R} = T_{UT\Rightarrow O} \times T^{-1}_{R\Rightarrow O} \tag{4.6}$$

The calculation of $T_{R\Rightarrow O}$ uses the same method as described above for the calculation of $T_{UT\Rightarrow O}$. The same approach to slice rejection was used. Slices with a reference body IRE greater than $3\sigma_{IRE}$ were discarded. Figure 4.8 shows the results of a similar analysis of $\langle|TRE|\rangle$ to $\langle|ILE|\rangle$ ratio vs visible IREDs as was done for the ultrasound tracking cross. This assumed that a 6 IRED reference body was positioned as shown in Figure 4.7 and the target point was the centroid of the prostate. Based on 4.8 we discarded any slices with less than 5 IREDs visible on the reference body.

We could not attach a reference rigid body to the patient during surgical trials. Therefore we assume that the patient does not move in relation to the Optotrak during surgery. During development work we used data from a plastic phantom data set and three cadaver data sets[1]. It was desirable to be able to move these objects relative to the Optotrak cameras, so reference rigid bodies were used.

4.2.2 Ultrasound Calibration

Ultrasound calibration was performed using an invariant point method detailed by Barratt et al. (2006). In the case of the phantom experiment we repeatedly imaged a pin head attached to the phantom and submerged in the same water bath as used for data collection. Figure 4.9 shows the location of the pinhead on the phantom.

The transform $T_{UI\Rightarrow UP}$ was first estimated using the ultrasound machine's own measurement function. This allows two points to be selected on screen and the distance between the two points measured. This is shown at the left of Figure 4.10. This gives a pixel to mm scaling which together with a definition of the image origin (in this case the top left corner of the image) defines $T_{UI\Rightarrow UP}$. $T_{UI\Rightarrow UP}$ converts the

---

[1] Details of data sets Phantom, Cadaver-01, Cadaver-02, and Cadaver-03 can be found in Appendix A.

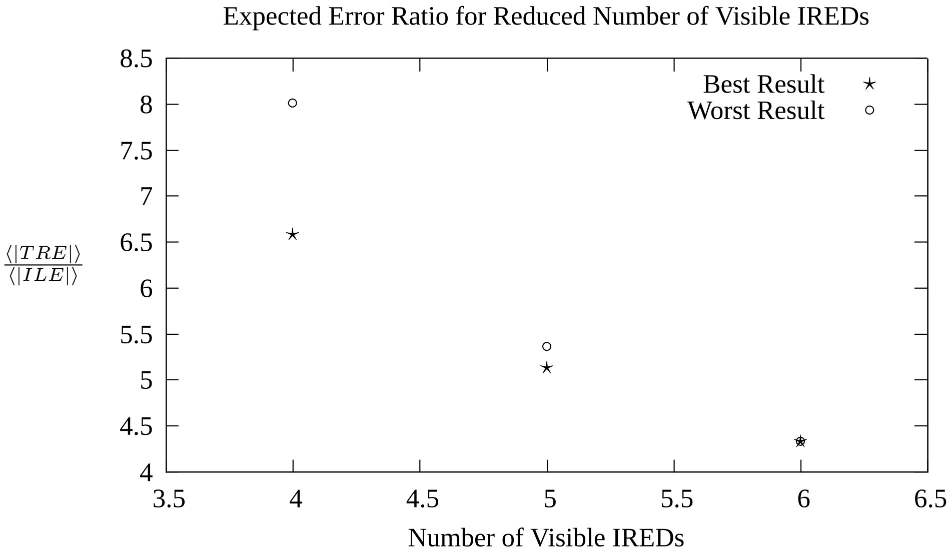

Figure 4.8: Best and worst case TRE versus number of IRED s visible for the reference object. These were calculated using Equation 4.5, using a target point at the prostate centroid. All possible IRED configurations (2^6) were tested then sorted into best and worst case TRE s.

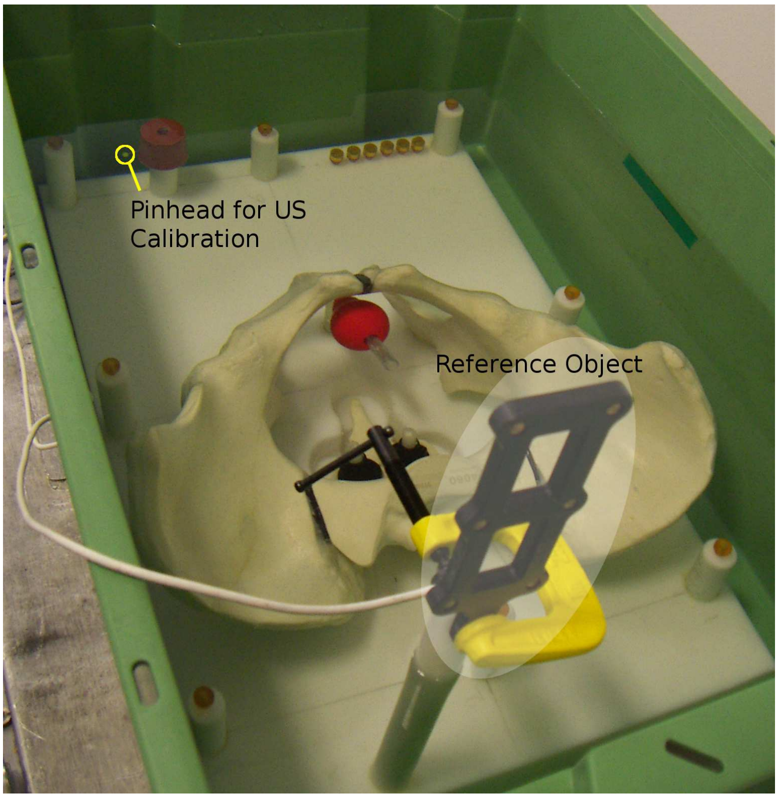

Figure 4.9: Photograph showing the pinhead object used for ultrasound calibration. Note the position relative to the reference object. The actual pinhead used for calibration is not visible in this image as it is too small. The object used was the cut off end of the pin to give small point. The visible pinhead is a plastic sphere and gave a poor ultrasound signal that was not usable for calibration.

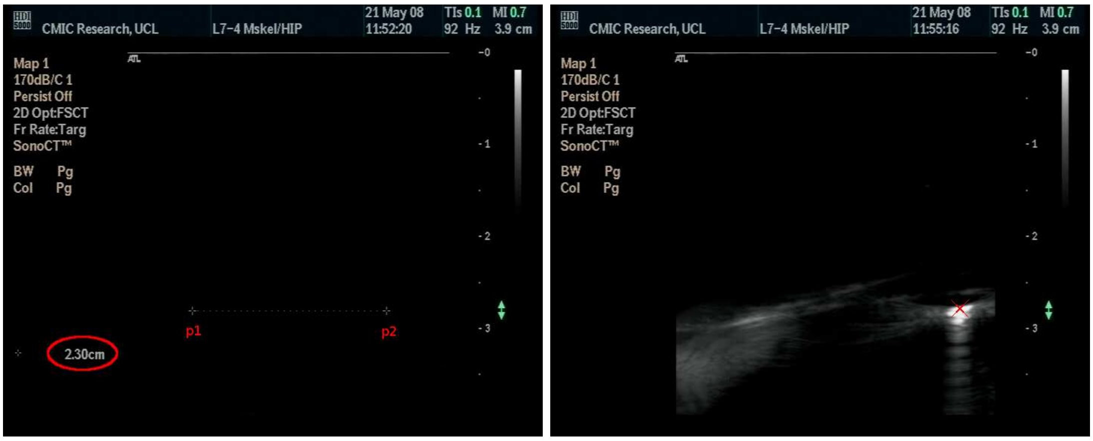
CMIC Research, UCL
L7-4 Mskel/HIP
21 May 08
11:52:20
TIs 0.1 MI 0.7
92 Hz 3.9 cm
Map 1
170dB/C 1
Persist Off
2D Opt:FSCT
Fr Rate:Targ
SonoCT™
BW Pg
Col Pg
p1
p2
2.30cm
CMIC Research, UCL
L7-4 Mskel/HIP
21 May 08
11:55:16
TIs 0.1 MI 0.7
92 Hz 3.9 cm
Map 1
170dB/C 1
Persist Off
2D Opt:FSCT
Fr Rate:Targ
SonoCT™
BW Pg
Col Pg

Figure 4.10: On the left is an example of the measurement used to determine $T_{UI \Rightarrow UP}$. The ultrasound image shows a distance which can be used to determine a pixel to mm value. The origin of the image is defined as the top left corner. On the right is a typical ultrasound image of the pinhead, with a cross indicating its manually segmented position.

pinhead coordinates in pixels to pinhead coordinates in millimetres relative to an arbitrary point defined on the probe.

The pinhead is now manually segmented from the ultrasound images, Figure 4.10 shows a typical ultrasound image of the pinhead in the water bath with a cross indicating the manual segmentation. Images that failed our previously described slice rejection test or in which the pinhead could not be properly defined were not used in the calibration. For N calibration images we now have N sets of 2D points (X_i) in mm relative to the probe face, and N estimates of $T_{Frame_i \Rightarrow Frame_{Ref}}$. We define the unchanging position of the pinhead (X_{PH}) as the origin of a "World" coordinate system. Equation 4.7 can be written for every frame.

$$X_{PH} = T_{Optotrak \Rightarrow World} \times T_{UT \Rightarrow O}[Frame_i] \times T_{UP \Rightarrow UT} \times X_i \tag{4.7}$$

From equation 4.7 we generate a system of $3N$ non linear equations from which we estimate the parameters of $T_{Optotrak \Rightarrow World}$ and $T_{UP \Rightarrow UT}$ using a Levenberg-Marquardt algorithm implemented in *MATLAB*®.

Speed of Sound Error

As formulated above $T_{UP \Rightarrow UT}$ depends on on the speed of sound in the calibration medium. If the speed of sound in the calibration medium is not the same as that in the tissue to be imaged then there will be an error in the measured position of the bone surface. The magnitude of this will depend on the size of the speed of sound mismatch and the amount of tissue that the ultrasound has to pass through to image the bone. Both Barratt et al. (2006) and Penney et al. (2006) account for this error by optimising the calibration parameters alongside the registration parameters. Their approaches rely on the assumption that the underlying anatomy being imaged is identical, so that shape variation between the two data sets will be due primarily to the this speed of sound error. In our case this is not true, as the preoperative data

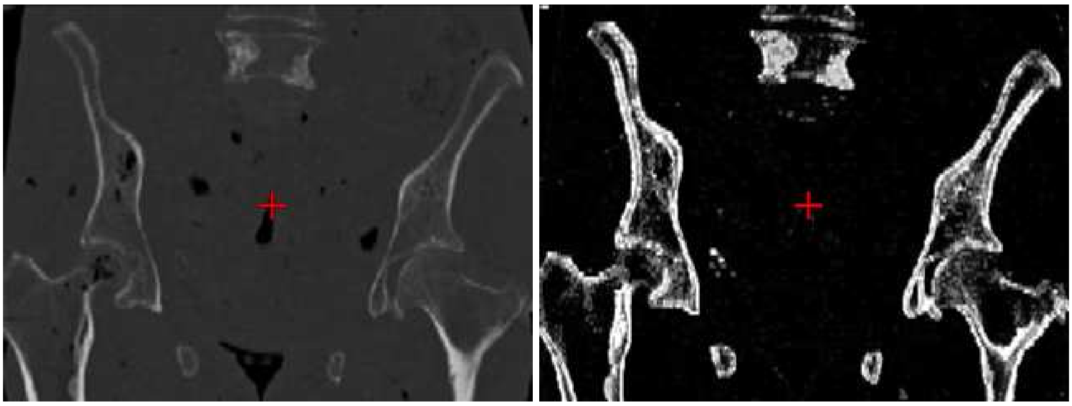

Figure 4.11: Pre-Processing of the CT image to a bone edge probability image. The image on the left is the CT image. On the right it has been converted to a bone edge probability image, using a Sobel filter, a local maximum filter, and a look up table.

is only an approximation of the shape of the patient's pelvic bone. The difference in bone shape in the two data sets will be a combination of shape modelling error and ultrasound imaging error due to the speed of sound error. Optimising the ultrasound calibration parameters to minimise the shape error will therefore yield an incorrect registration. Therefore we do no not attempt to implement these methods here. Our focus will be on minimising the speed of sound mismatch by controlling the calibration medium.

4.2.3 Transforming the Ultrasound Data

The transform $T_{UI \Rightarrow O}$ has now been defined for each frame of the ultrasound data. Pixels with a bone edge probability value greater than 0 are transformed to 3D points in the Optotrak's coordinate system using $T_{UI \Rightarrow O}$. Zero value pixels are discarded. The remaining points are stored as a list of 3D points and corresponding bone edge probabilities. In comparison Penney et al. (2006) store the images as a group of 2D points for each slice and the corresponding $T_{UI \Rightarrow O}$ for each slice. Using our method we avoid the need to apply $T_{UI \Rightarrow O}$ during the registration algorithm.

4.3 CT Image Pre-Processing

Our registration method matches the ultrasound image formed in the preceding section with a CT image that has been warped to fit the patients MRI data using the methods in Chapter 3. We follow the same methodology as Penney et al. (2006) by converting this CT image into a bone edge probability image prior to registration. The method used is similar to that used for the ultrasound images. For every voxel of the ultrasound image two metrics are calculated from the CT image. Then a look up table trained on manual bone segmentation is used to convert to the two metrics to a bone edge probability. The first metric is the edge gradient intensity found using a Sobel filter applied the CT image. The second metric is the maximum intensity found within a 3×3 voxel neighbourhood around each voxel of the CT image. We used the same look up table as Penney et al. (2006) for the conversion from these two metrics to the bone edge probability. Figure 4.11 gives an example of this image pre-processing.

4.4 Image to Image Registration

We use an image to image registration to determine $T_{SM(nrr) \Rightarrow UI}$. The registration used to determine $T_{SM(nrr) \Rightarrow UI}$ is an intra patient registration using images of similar modality (both images have been processed to bone edge probabilities). We have therefore kept the registration algorithm simple. We first define a similarity measure to be optimised, then the optimisation algorithm. We implemented in C++.

4.4.1 Similarity Measure

The similarity measure is calculated for every point ($X_{US}[i]$) in the ultrasound image (stored as a list of N points). The corresponding point ($X_{CT}[i]$) in the CT image is first found. $X_{US}[i]$ is transformed by the current estimate of $T^{-1}_{SM(nrr) \Rightarrow UI}$ to give a coordinate in the CT image. Nearest neighbour interpolation is used to return the bone edge probability ($BEP_{CT}[i]$) of the voxel nearest to this point. Interpolation of the CT image is required as the resolution of the ultrasound image is much higher than that of the CT image (0.088 mm versus 0.7 mm in plane). We tested the effect of using tri-linear interpolation instead of nearest neighbour interpolation on registration accuracy. Using tri-linear interpolation did not improve registration accuracy and has a computational cost. Therefore we used nearest neighbour interpolation. If the coordinate is outside the bounds of the CT image a zero value is returned.

The image similarity is then defined as a simple cross correlation in equation 4.8.

$$Sim_{CC} = \frac{\sum_{i=1}^{N} BEP_{CT}[i] \times BEP_{US}[i]}{N} \tag{4.8}$$

We chose this similarity measure because it can be calculated rapidly, requiring a single loop through the ultrasound points. Using this measure there is no normalisation for overlapping ultrasound points. If the point cloud ultrasound image contains a high proportion of points from a particular region then the registration will be weighted to that region. Therefore the value of the similarity measure depends on both the number and distribution of ultrasound slices. This feature can be used to our advantage. Section 4.6 will detail an experiment that examines the effect of different ultrasound slice distributions on the registration accuracy.

4.4.2 Optimisation Algorithm

We elected to use a basic gradient descent optimisation algorithm. This is a well established optimiser for rigid registration of medical images. We used a similar algorithm for the determination of $T_{SM \Rightarrow M}[nd9]$ in Chapter 3. We start with an estimate of $T_{SM(nrr) \Rightarrow UI}$ and use this to calculate Sim_{CC}. Gradients are then determined numerically in each of the 6 degrees of freedom (3 translations and 3 rotations) by perturbing $T_{SM(nrr) \Rightarrow UI}$ by a predetermined step size ($d_{x,y,z,\alpha,\beta,\gamma}$) and recalculating Sim_{CC}. This gives 12 new estimates of $T_{SM(nrr) \Rightarrow UI}$. The estimate with the highest value of Sim_{CC} is used as the new estimate of $T_{SM(nrr) \Rightarrow UI}$ and the process is repeated until a maxima of Sim_{CC} is found. The step size ($d_{x,y,z,\alpha,\beta,\gamma}$) is then halved and the process repeated until a new maxima is found. The step size if halved $N_{StepHalve}$ times. This process requires an initial step size to be defined. We found through trial and error that an initial step size of 4mm for the translation components ($d_{x,y,z}$) gave good results for our application. The ratio (slightly less than 0.25 degrees per mm) between rotation step size and translation step size was chosen so that over the region of interest (the pelvic bone) a rotation step would give the same average displacement as a translation step. Applying this ratio to the 4mm initial step size gives an initial rotation step size of 0.98 degrees (for $d_{\alpha,\beta,\gamma}$).

Target Blurring

Two versions of the preoperative image were used. For large step sizes a blurred image was used (the edge probabilities were blurred with a 4 mm Gaussian kernel). After $\frac{N_{StepHalve}}{2}$ step size reductions the edge probability image was used without blurring.

Registration Initialisation

An initial estimate of $T_{SM(nrr)\Rightarrow UI}$ is required to start the registration process. For pelvic anatomy Penney et al. (2006) describe a practical method involving a tracked pointer and the picking of skin landmarks. This may not be necessary in our case as the position of the patient's pelvis in a *daVinci*^{TM} based prostatectomy is constrained by the practicalities of fitting all the equipment in the room. Therefore an estimate of $T_{SM(nrr)\Rightarrow UI}$ can be taken from a past operation. If this is not sufficiently accurate it could be refined rapidly using suitable visualisation software, eg VTK. This will be investigated further in 6.

Registration Termination

In practice the optimisation algorithm may not always converge on the global maximum. The following experiments will determine how reliable the optimisation algorithm is. To allow for the occurrence of failed optimisation in theatre we propose running the algorithm from multiple, randomly perturbed estimates of $T_{SM(nrr)\Rightarrow UI}$ and checking for convergence of the result. To check for convergence we define a "Transformation Distance" D. D is a measure of the size of the transform and is defined by equation 4.9. $[XYZ]_{trans}$ are the translations in mm and $[XYZ]_{rot}$ are the rotations in degrees. The constant C is used to convert degrees to mm and is set so that a one degree rotation causes an average translation of 1 mm over the region of interest (the pelvis). We have used a value of $4.08\frac{mm}{\circ}$.

$$Distance = \sqrt{X_{trans}^2 + Y_{trans}^2 + Z_{trans}^2 + (CX_{rot})^2 + (CY_{rot})^2 + (CZ_{rot})^2} \tag{4.9}$$

In theatre we propose running the following procedure. We start with an initial estimate of $T_{SM(nrr)\Rightarrow UI}[init]$. If this corresponds to a local maxima in the image similarity (equation 4.8) then running the registration optimiser will have no effect. Therefore we perturb $T_{SM(nrr)\Rightarrow UI}$ by a random transform of distance $D_{perturb}$. The process for generating this random transform will be described below. The registration optimiser is then run from this perturbed starting position. This process is repeated N_{reg} times, to give N_{reg} new estimates of $T_{SM(nrr)\Rightarrow UI}$. The distance of each of these transforms from $T_{SM(nrr)\Rightarrow UI}[init]$ is calculated using equation 4.9. If $T_{SM(nrr)\Rightarrow UI}[init]$ is near the global maxima we would expect the RMS distance to be near zero. We set a threshold value $D_{threshold}$. If the RMS distance of the estimated transforms from $T_{SM(nrr)\Rightarrow UI}[init]$ is less than $D_{threshold}$ we consider the registration to have converged at the global maximum. If not we repeat the process with a new estimate of $T_{SM(nrr)\Rightarrow UI}[init]$. The new estimate of $T_{SM(nrr)\Rightarrow UI}[init]$ is chosen from the N_{reg} calculated estimates of $T_{SM(nrr)\Rightarrow UI}$ on the basis of their optimised similarity measure values. The transform which gives the highest similarity measure is used.

The number of repeats N_{reg} is a function of the failure rate of the optimisation algorithm. The perturbation distance $D_{perturb}$ is a function of the capture range of the algorithm. The threshold distance $D_{threshold}$ is a function of the accuracy of the optimisation algorithm. We will attempt to quantify these values in the following sections in order to design a robust registration process. We now describe a process for the perturbation of $T_{SM(nrr)\Rightarrow UI}[init]$.

Perturbation of Rigid Transforms

We wish to define a method for randomly perturbing a rigid transform by a set distance D. These random perturbations will be used to check for convergence of the registration algorithm.

The distance $D_{perturb}$ (equation 4.9) defines a six dimensional hypersphere of radius D. The intent here is to randomly sample from points uniformly spread on the surface of this hypersphere. We use the following method to do this. Values of each of the transformation parameters (3 translations and 3 rotations) are independently randomly sampled from a normal distribution with mean zero and standard deviation 1.0 to give $[XYZ]_{trans}^{normal}$ and $[XYZ]_{rot}^{normal}$. Equation 4.9 is then used with $C = 1.0$ to calculate D^{normal}. The transformation parameters are then scaled using $\frac{D_{perturb}}{D^{normal}}$, ie. $X_{trans} = X_{trans}^{normal} \times \frac{D_{perturb}}{D^{normal}}$. Finally the rotational components are divided by $C = 4.08$.

4.4.3 Summary of Registration Parameters

The registration algorithm as described has many parameters and their selection will influence the success of the algorithm. The parameters are summarised in Table 4.1.

The performance of the proposed registration algorithm used will depend on the value used for the parameters listed in Table 4.1. The remainder of this Chapter is divided into 5 experiments. The first four investigate the sensitivity of the algorithm to four of the parameters , ultrasound processing, number and distribution of ultrasound slices, number of step size reductions, and the capture range. The chapter is concluded with an experiment to validate the performance of the complete algorithm on 4 sets of data.

4.5 Selection of Ultrasound Image Processing Parameters

4.5.1 Aim

Preprocessing of the ultrasound image from ultrasound intensity to bone edge probability requires the selection of two parameters. These are the artefact intensity threshold and the degree of blurring applied to the ultrasound image. This experiment seeks to determine what effect the value of these two parameters has on registration accuracy.

4.5.2 Data

For testing of the registration algorithm we designed and constructed a phantom. Figure 4.12 shows the phantom. The phantom consists of a plastic pelvic bone and prostate rigidly mounted on a thick plastic base. 8 fiducial markers are also rigidly attached to the perimeter of the rigid base. We collected a CT image of the phantom and 654 ultrasound slices of the phantom immersed in a water bath. Details of the data collected are given in Appendix A. The fiducial markers were locating in both the CT image and in the coordinates of the Optotrak camera. A tracked pointer was used to locate the fiducial markers in the

Ultrasound Processing	
Artefact Distance Threshold	3mm
Artefact Intensity Threshold	0 to 250
Ultrasound Blur Radius	0 to 8 mm
Edge Probability Threshold	-1 , 0
Slice Rejection IRED threshold	11 (ultrasound cross) / 5 (Reference)
Calibration Matrix	Fixed
CT Model Processing (Adapted with change from Penney et al. (2006)	
Sobel Filter	On
Maximum Value Mask Size	3×3
Registration	
Number of Ultrasound Slices	10 to 600
Distribution of Ultrasound Slices	Even or Weighted
Initial Step Size	4mm
Number Of Step Size Reductions	4 to 10
Ultrasound In Plane Resolution	0.088 mm
CT In Plane Resolution	0.7 mm
CT Image Interpolation	Nearest Neighbour
CT Blurring	4 mm Gaussian Kernel
CT Out of Plane Resolution	2 mm
Registration Starting Distance	10 to 60 mm
Degrees to mm Conversion	$0.25°/mm$
Convergence / Termination Criteria	Yes
Similarity Measure	Cross Correlation

Table 4.1: Summary of the parameters used by the registration algorithm.

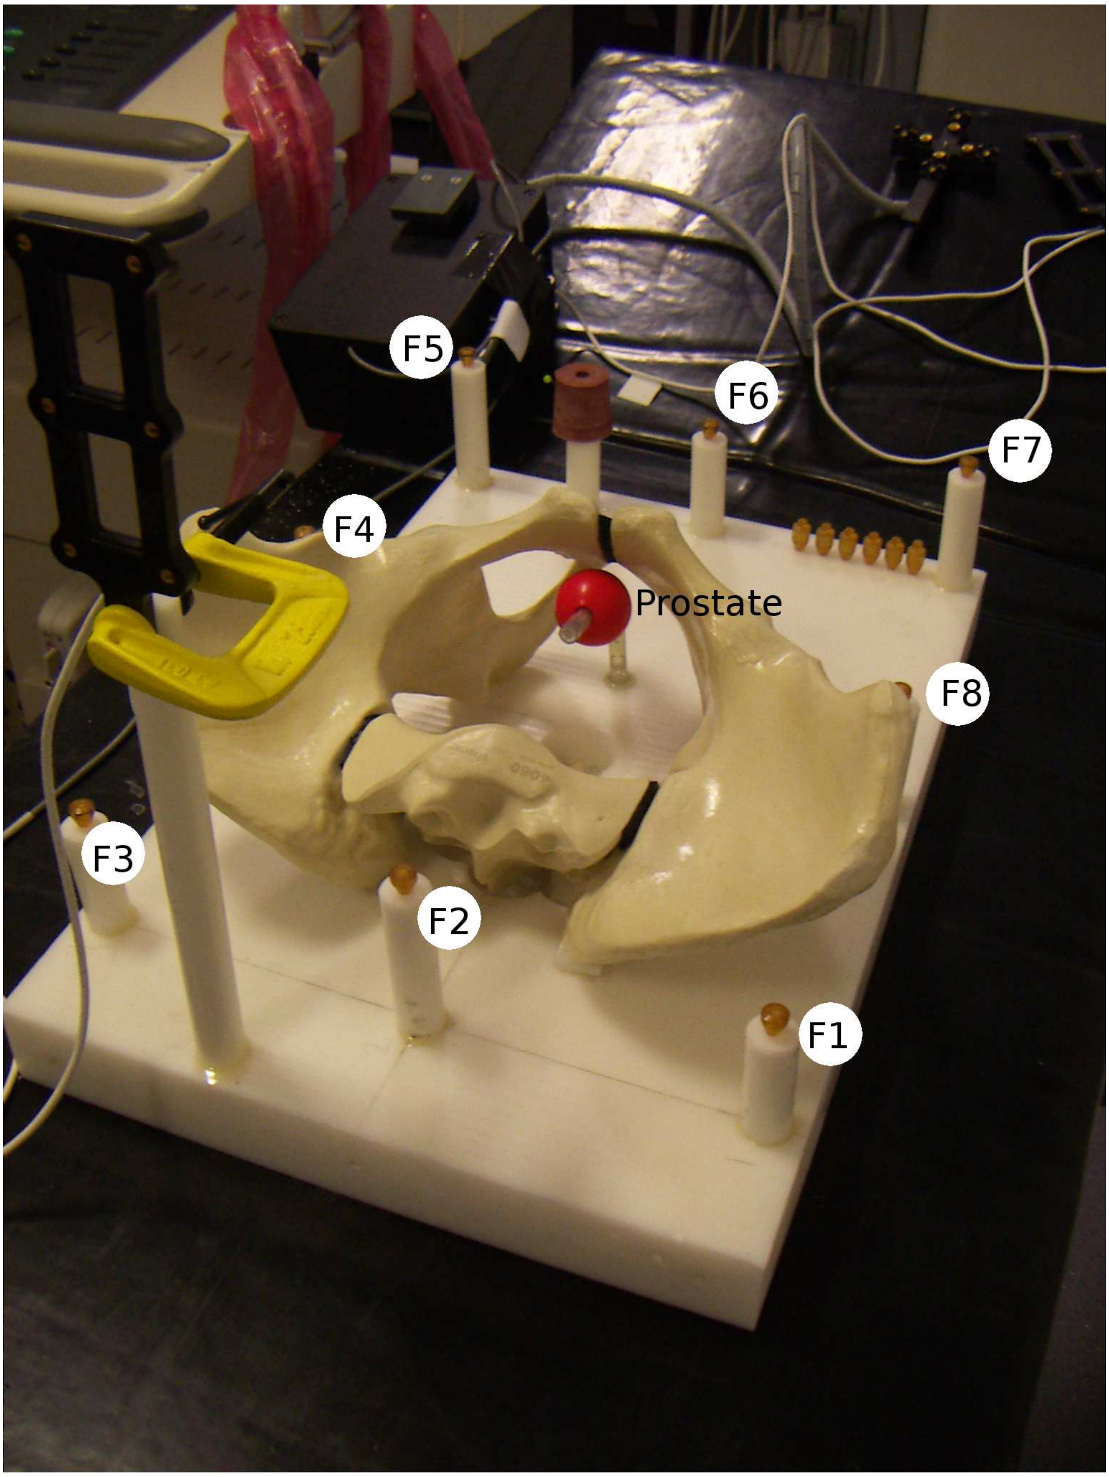

Figure 4.12: The plastic phantom constructed for registration experiments. A plastic pelvis and prostate are rigidly mounted on a thick plastic base. 8 fiducial markers are also rigidly attached to the plastic base. The fiducial markers are used for gold standard registrations between different image sets (CT and ultrasound). The reference rigid body clamped on is removed for CT imaging. Its purpose is to track motion occurring during ultrasound scanning.

Optotrak coordinate system. A point based registration between the two point sets was then used to get a gold standard version of $T_{SM(nrr)\Rightarrow O}[GS]$. The 8 fiducial markers are located so that their centroid is at the prostate, to minimise errors in the gold standard transform. To allow for movement of the phantom during ultrasound scanning a reference body was attached to the phantom.

Ultrasound and CT images of the phantom were used for this experiment. 32 of the ultrasound images were manually segmented and these were used for building the look tables used by the ultrasound processing algorithm.

4.5.3 Method

Two experiments were performed. The first was to look at the effect of artefact intensity threshold and blur radius on TRE at the prostate. The second was to try and compare the automatic image processing with a manual segmentation.

TRE at Prostate

The CT data was processed to a bone edge probability image (BEP_{CT}) as per Section 4.3. The processing of the ultrasound data was performed as described in Section 4.1 to give an ultrasound derived bone edge probability image (BEP_{US}). This is a cloud of 3D points in Optotrak coordinates found using $T_{UI\Rightarrow O}$. Ultrasound probe calibration was performed on a pinhead attached to the phantom and immersed in the same water bath, so there should be no speed of sound error present in $T_{UI\Rightarrow O}$. Ultrasound image processing was repeated using different values of artefact intensity threshold and blur radius. Blur radius was varied from 0 to 8 mm in 1 mm steps and artefact intensity threshold was varied from 0 to 250 in steps of 10, ie $26 \times 9 = 234$ versions of BEP_{US} were tested. The ultrasound images are 256 level grey scale images so the maximum possible artefact intensity threshold is 255. A threshold of 255 would result in all ultrasound data being discarded.

The registration algorithm detailed in Section 4.4 was used to register each version of BEP_{US} to BEP_{CT}. The CT was put into the Optotrak's coordinate system using $T_{SM(nrr)\Rightarrow O}[GS]$. The ultrasound images BEP_{US} were then perturbed using a random transform of distance 8 mm as described in Section 4.4.2. The registration algorithm was then used to get an estimate of $T_{SM(nrr)\Rightarrow UI}$ and hence $T_{SM(nrr)\Rightarrow O}[Estimate]$. The RMS distance between the nominal prostate points (see section 1.6.1) defined by $T_{SM(nrr)\Rightarrow O}[GS]$ and $T_{SM(nrr)\Rightarrow O}[Estimate]$ was then calculated.

Comparison with Manual Segmentation

An alternative to using the automatic processing of the ultrasound image would be to manually segment the ultrasound images. It would be interesting to know how our method compares with registration using a manual segmentation. Rather than manually segmenting all the ultrasound images (over 600) to find the TRE directly, we attempt here to find it indirectly by correlating the TRE with a separate image similarity metric. This is done using only the 32 images for which we have a manual segmentation.

For each of the 32 manually segmented images (US_i) a look up table for bone edge probability is constructed from the other 31 images (ie. leave one out testing). This look up table is then used to process the image US_i to a bone edge probability image BEP_i as per Section 4.1. This bone edge probability image is then compared to the manual segmentation MAN_i using a normalised cross correlation similarity measure, defined by Equation 4.10. N is the number of pixels in each image, $MAN_i[j]$ refers to the jth pixel of the manual segmentation of the ith image. $BEP_i[j]$ refers to the corresponding pixel in the bone edge probability image.

$$Similarity = \frac{\sum_{j=1}^{N} MAN_i[j] \times BEP_i[j]}{\sqrt{\sum_{j=1}^{N} MAN_i[j]^2 \times \sum_{j=1}^{N} BEP_i[j]^2}} \tag{4.10}$$

An average value of image similarity for all 32 images is calculated. Calculation of the average image similarity is repeated for the values of artefact threshold and blur radius as in the preceding experiment. The average similarity can then be plotted against the TRE calculated in the preceding experiment. As the similarity metric will be one if the images are identical then it may be possible to extrapolate a TRE for a theoretical manual segmentation of the ultrasound image.

4.5.4 Results

Figure 4.13 shows the TRE at the prostate versus artefact threshold and blur radius. For the phantom data it can be seen that the selection of the artefact threshold can cause the TRE to vary significantly, (by up to 1.24 mm). Selection of blur radius does not have any significant effect on the registration accuracy. As larger blur radii require longer processing time it therefore makes sense to use little or no blurring. The minimum value for TRE occurs at 240. Above this limit the TRE increases dramatically, presumably as there is insufficient bone surface left to perform a registration.

Figure 4.14 shows the image similarity metric plotted against the TRE calculated at the prostate. Each data point represents a look up table using a particular combination of blur radius and artefact threshold. The TRE is calculated as above, by registering the processed ultrasound to the CT image. The image similarity measure is calculated by comparing the processed individual slices to a manual segmentation of the slice. It was found that the relationship between the image similarity metric and the TRE depended on the artefact threshold used. Therefore the individual data points have been grouped depending on their artefact threshold.

For artefact threshold below 150, increasing the artefact threshold tends to increase the image similarity metric and reduce the TRE. For artefact thresholds above 150, increasing the artefact threshold continues to reduce the TRE, however the image similarity metric reduces. For this reason it was not possible to extrapolate a TRE for a similarity of 1.

The change in the correlation direction raises the question of whether a segmentation using a high artefact threshold will outperform a manual segmentation. Using a high threshold will exclude all but the strongest bone edges, whereas a manual segmentation will include weaker edges due to the human operator extrapolating a surface between two obvious surfaces. In cases where many ultrasound slices are available (as is the case here) rejection of all but the strongest edges may lead to better registration results.

There is some evidence that this may be the case. A comparison between the results of Penney et al. (2006) and Barratt et al. (2006), see Table 2.6. These papers used the same data sets, Penney et al. (2006)

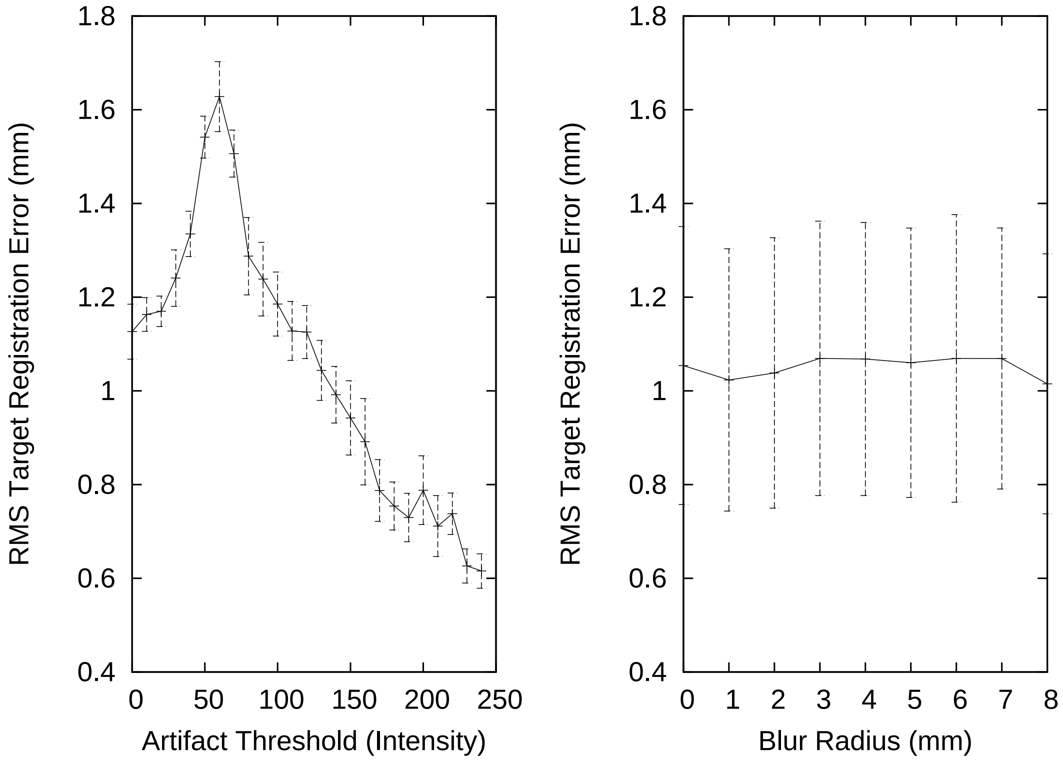

Figure 4.13: Plots of the registration TRE measured at the prostate surface, for a range of artefact threshold intensities and blur radii. In each case the error bars show 1 standard deviation of the results at that point, ie for the threshold value plot the error bars represent the variation in TRE due to the blur radius at that point. It is clear that for the phantom data the selection of the right artefact threshold is important, while the blur radius has no significant effect on the results. The minimum TRE was achieved using an artefact threshold of 240. At a threshold of 250 the TRE increased dramatically (to over 5mm), therefore this was not plotted. The following registration parameters were used: initial step size 4 mm, Terminal step size 0.25 mm, Registration start distance 8 mm.

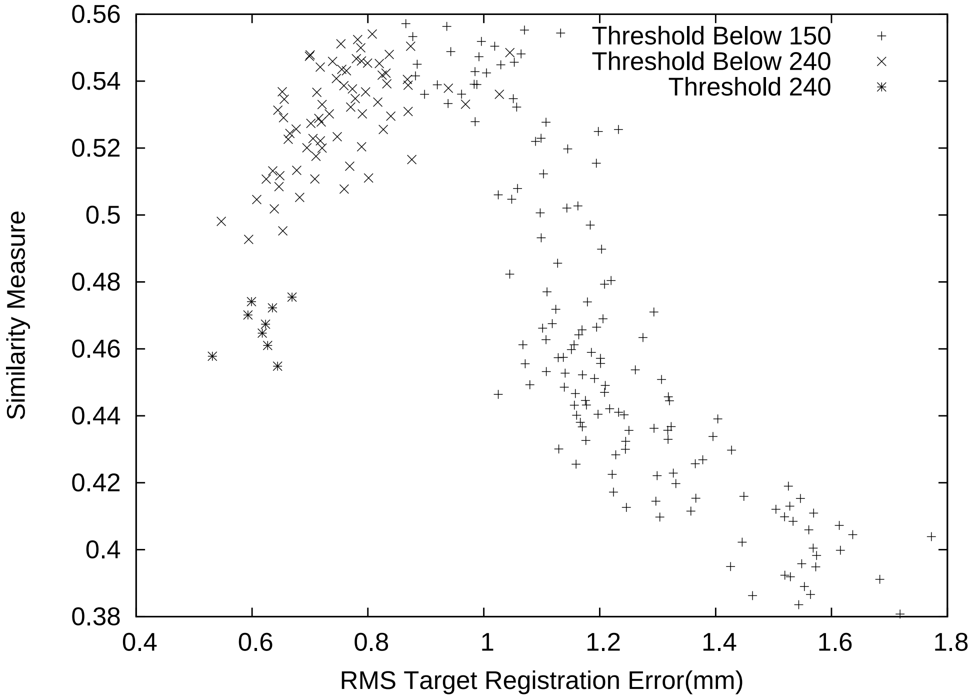

Figure 4.14: Registration TRE versus image similarity for phantom data. Below an artefact threshold of 150 there is a correlation between TRE and image similarity. TRE reduces as the similarity measure increases. Above this threshold the correlation changes direction, with TRE reducing as the similarity measure reduces. Results for thresholds greater than 240 are not plotted as they give a very high TRE.

used automatic segmentation while Barratt et al. (2006) used manually segmented data. The results were presented as registration TRE across the entire pelvic bone rather than just at the prostate. Penney et al. (2006) give an error of 1.63 mm, while Barratt et al. (2006) get an error of 2.17 mm, ie. automatic segmentation outperforms manual segmentation in this case. It should be noted that there are other differences in their approaches that will also contribute to their differences in the error metric. A more detailed comparison of the two methods on the cadaver data may make for interesting future work.

4.6 Number and Distribution of Ultrasound Slices

4.6.1 Aim

The ultrasound derived bone edge probability image BEP_{US} is defined as a set of discrete points. By including more points from some regions of the pelvis than from others it is possible to weight the similarity metric described by equation 4.8. This experiment aims to show what effect both the number of ultrasound slices used and their distribution will have on the registration accuracy.

4.6.2 Data

The CT image of the phantom (described in Section 4.5) was used here. Rather than registering the ultrasound data to this we generated “simulated” ultrasound data. Using the real ultrasound data would introduce tracking and calibration errors present in the determination of $T_{UI \Rightarrow O}$. To avoid this we created simulated ultrasound by re-slicing and interpolating the CT data along the planes defined by the ultrasound data, transformed using the known gold standard transform $T_{SM(nrr) \Rightarrow O}[GS]^{-1}$. The result is two different representations of the same image, one is a voxel based image, while the other is a point cloud.

4.6.3 Method

A slice classification algorithm was first used to classify the simulated ultrasound slices into regions. We manually defined three regions of the pelvis as 3D ellipsoids. These regions are defined in the coordinate system of the CT data. The simulated ultrasound slices are transformed to the CT image using $T_{SM(nrr) \Rightarrow O}[GS]^{-1}$ and classified based on which ellipsoid their centre point falls within. We started with 654 ultrasound slices, 175 of these were classified as pubis slices, 228 were from the left iliac, and 247 were from the right iliac. 4 slices did not fall within any of the classification ellipsoids so were discarded.

The total number of ultrasound slices used was varied from 20 to 600. The slices used were randomly selected from the available slices. Three slice selection methods were used. The first randomly selected slices from all subgroups of slices. The second used only slices drawn from the iliac regions. The third used slices drawn only from the pubic region.

The CT and simulated ultrasound data were brought into alignment using $T_{SM(nrr) \Rightarrow O}[GS]^{-1}$. The ultrasound data was then perturbed using a random transform of size 10 mm as defined in Section 4.4.2. The registration algorithm described in Section 4.4 was then used to estimate $T_{SM(nrr) \Rightarrow UI}$ and hence $T_{SM(nrr) \Rightarrow O}[est.]^{-1}$. $T_{SM(nrr) \Rightarrow O}[est]^{-1}$ and $T_{SM(nrr) \Rightarrow O}[GS]^{-1}$ were used to measure TRE at the prostate for each data point. The registration optimisation was repeated 20 times for each data point.

4.6.4 Results

Figure 4.15 shows the relationship between the number of slices used, their distribution and the resulting TRE at the prostate.

Registration using a subset of slices from across the pelvis is successful providing more than 100 slices are used. Registration using just iliac slices yields notably poorer results, they also show an increase in TRE as the number of slices is reduced. The results using just pubis data were very poor, with TRE around 20 mm. For this reason they are not shown in Figure 4.15.

4.7 Number of Step Size Reductions

4.7.1 Aim

The registration algorithm defined in Section 4.4 proceeds by halving the step size each time a local maxima is found. This experiment seeks to determine what effect the number of step size reductions (and hence the terminal step size for a constant initial step size of 4mm) has on the accuracy of the registration algorithm.

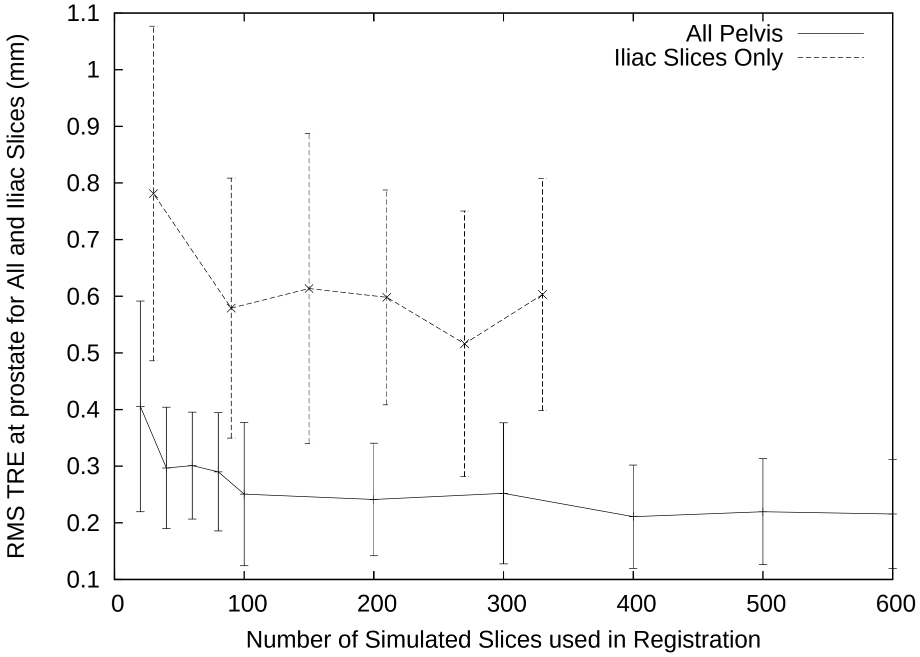

Figure 4.15: Registration of CT to simulated CT data. TRE calculated at points on prostate surface. For each number of slices the registration was repeated 20 times. Registrations with TRE greater than 2 mm were classed as failures and rejected. (This gave 3 failures at All pelvis 60 slices, 2 at All Pelvis 40 slices, Otherwise no data point had more than 1 outlier). The average and standard deviation of the remaining TREs were computed. The error bars represent +/- 1 standard deviation of the successful registrations. When using slices spread across the pelvis (All Pelvis) the result appears fairly stable as long as more than 100 slices are used. Below this and the TRE steadily climbs. Registration using just the iliac data appears to work but gives a worse result than using all data. Registration using just the pubis data is not shown as the TREs were very high, averaging 20 mm.

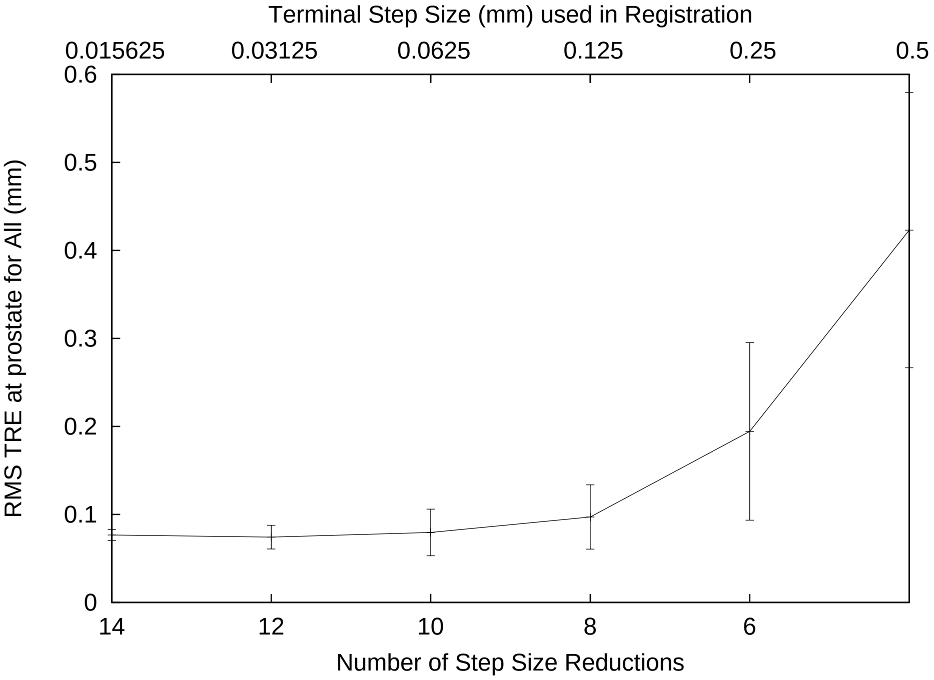

Figure 4.16: Effect of number of step size reductions on registration accuracy. Registration TRE of points on prostate surface using synthetic (ideal) data. For each number of reductions, the registration was repeated 10 times. Failed registrations (with a TRE greater than 2 mm) were removed (1 failed registration at 4 step size reductions). The average and standard deviation of the remaining TREs was computed. The error bars represent +/- 1 standard deviation of the remaining data. Beyond 8 step size reductions there is no improvement in average result (around 0.075 mm), however the standard deviation continues to reduce.

4.7.2 Data

The phantom data was used with simulated ultrasound, as per Section 4.6.

4.7.3 Method

The simulated ultrasound was registered to the CT using all available ultrasound slices and different numbers of step size reductions, from 4 to 14. At each data point the registration was repeated from different random starting points using the same process as in 4.6.

4.7.4 Results

Figure 4.16 shows the results. More step size reductions will give a better registration. The improvement (in reduction of standard deviation, rather than average) continues beyond 8 step size reductions. It is interesting to compare the terminal step size with the resolution of the CT image used ($0.9 \times 0.9 \times 0.5mm^3$). The CT image resolution is not having a significant effect on the registration accuracy in this case. The terminal step size has a larger effect. The benefits gained by using more than 8 step size reductions are marginal in the context of the system error. As there is a computational cost in using more step size reductions we only propose to use 8 step size reductions in practice.

4.8 Optimiser Capture Range

4.8.1 Aim

The experiments to date have only used registrations from starting estimates of $T_{SM(nrr)\Rightarrow UI}$ close to the known gold standard registration. We have not yet determined how close our initialisation estimate of $T_{SM(nrr)\Rightarrow UI}$ will be in theatre, this will be discussed more in Chapter 6. For a similar problem Penney et al. (2006) used an initial estimate of $T_{SM(nrr)\Rightarrow UI}$ based on picking three skin surface points. They generated starting estimatefor $T_{SM(nrr)\Rightarrow UI}$ based on the assumption that these skin points were localised with an error of 20 mm. We measured the distance (see equation 4.9) of 300 of their starting estimates from the known gold standard value of $T_{SM(nrr)\Rightarrow UI}$. The average starting distance was 30 mm. If we assume that we will achieve similar initialisation then our algorithm must be able to reach the global maximum from initialisations at the this distance. The aim of this experiment is to determine the capture range of the algorithm.

4.8.2 Data

The same “simulated ultrasound” method as used in Sections 4.6 and 4.7 was used to generate ultrasound slices from voxel based CT images. In this case rather than using the phantom data sets we used the 21 CT data sets used in Chapter 3. These were first registered to the phantom so that the ultrasound slice positions from the phantom scan could be used to re-slice the data. The CT images were first processed to bone edge probability images using the method in section 4.3. This gave 21 data sets of voxel CT images and corresponding slice based images of the same data.

4.8.3 Method

The registration algorithm described in Section 4.4 was used to register the simulated ultrasound to the matching CT data from starting estimates $T_{SM(nrr)\Rightarrow UI}[init]$. The starting estimates were generated by perturbing the known gold standard $T_{SM(nrr)\Rightarrow UI}[GS]$ by a distance D_{init}. The value of D_{init} was varied from 10 to 60 mm. At each distance the registration was repeated 10 times for each of the 21 data sets. The convergence checking termination criteria described in Section 4.4 was used with the following parameters. N_{reg} was 10, $D_{perturb}$ was 10 mm, and D_{thresh} was 2 mm. The method for perturbing the transformations is described in Section 4.4. The results were measured as a registration failure rate (the percentage of registrations with a TRE at the prostate greater than 2mm) and as the RMS TRE for successful registrations. TRE and failure rate were calculated both for individual runs of the gradient descent optimiser and for the global optimiser which uses repeat runs of the gradient descent optimiser.

4.8.4 Results

Figure 4.17 plots the results in terms of TRE and failure rate at start distances up to 60mm. On the left is shown the performance of the gradient descent optimiser on its own. This only operates reliably at a start distance of 10mm or less. This is unlikely to be possible in practice. On the right the performance of the full algorithm described in Section 4.4 is shown. By repeating the gradient descent optimisation and checking for convergence of the resulting transforms it is possible to get satisfactory registration results up to a start distance of 40mm. This is probably adequate for our application.

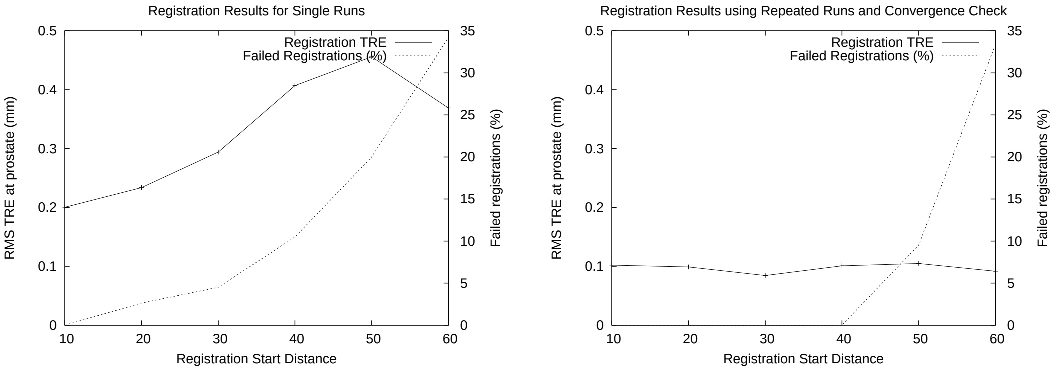

Figure 4.17: Registration TRE (left axes) and number of failed registrations (right axes) for start distances up to 60 mm using re-sliced CT data. For each starting distance registration using the repeated run and convergence check algorithm was performed on 21 different CT data sets. The chart on the left shows the performance of individual registration optimisations using the gradient descent optimiser. On the right is the performance of the full algorithm which repeats the registration 10 times and checks for convergence of the results. A failed registration was defined as having a TRE greater than 2 mm.

4.9 Registration Using Real Ultrasound Data

4.9.1 Aim

With the registration parameters determined as in the preceding sections, the algorithm is now tested on real ultrasound and CT data, both on its own and in combination with the shape model fitting presented in Chapter 3. This experiment determines the accuracy with which the $T_{M \Rightarrow O}$, see Figure 1.3, can be found. In addition the experiment confirms the effects of the error in $T_{SM \Rightarrow M}$ on registration accuracy.

4.9.2 Data

Four data sets were used (Phantom, Cadaver-01, Cadaver-02, and Cadaver-03, see Appendix A). These have CT data for the pelvis and matching ultrasound slices. For each data set a gold standard registration between the two modalities was established using fiducial markers. In the case of the cadaver data, the gold standard registrations were determined using bone implanted fiducial markers, see Penney et al. (2006) for details. For the phantom data the gold standard registration was performed using a set of eight fiducial markers around the pelvic bone, described in Section 4.5.

4.9.3 Method

The ultrasound data was first transformed to the CT data, using the known gold standard transforms $T^{-1}_{SM(nrr) \Rightarrow UI}$. All data sets were then transformed into the coordinate system of the shape model from Chapter 3. For each of the four data sets $T_{SM \Rightarrow M}[nd9]^{-1}$ was found by aligning the CT data to one of the pelves (SM-K) used in the shape model building in Chapter 3. $T_{SM \Rightarrow M}[nd9]^{-1}$ was then used to transform both the CT and ultrasound data to the model's coordinate system. The CT image was re-sampled to the voxel sizes of the shape model using tri-linear interpolation. The ultrasound data set did not require interpolation as the data is stored as coordinates in 3D space.

Region	Phantom	Cadaver-01	Cadaver-02	Cadaver-03
Total Slices Collected	721	282	482	342
Used for LookUp Table Construction	32	N/A	N/A	N/A
Rejected Due to Lack of visible IREDs	35	81	164	10
Pubis	175	82	81	39
Left Iliac	228	70	176	152
Right Iliac	247	46	59	140
Not Classified	4	3	2	1
Total Good	650	198	316	331

Table 4.2: The results of slice rejection and classification of the four data sets used. Slices with less than 11 IRED s visible on the ultrasound cross or less than 5 IRED s visible on the reference object were rejected. Slices were then classified into regions based on the position of their centres relative to regions manually defined on the Phantomdata set. Slices that fell outside all regions are listed as Not Classified and were not used.

Conversion of the phantom ultrasound data to probability images was done using an artefact threshold of 240 and a blur radius of 1 pixel as suggested by the analysis in Section 4.5. Conversion of the cadaver data was done using the same parameters and lookup tables as Penney et al. (2006). The ultrasound data for each data set was classified into regions based on the position of the slice centres relative to the regions manually defined in the phantom data set. Table 4.2 shows the results of this slice classification. Conversion of CT images to bone edge probability images was done using a Sobel filter and a local maximum as described by Penney et al. (2006).

For each data set, the ultrasound data was registered to the CT data using the algorithm described in Section 4.4. Four different variations were trialled. Two methods of processing the CT data and two different distributions of ultrasound slices were used.

The CT data was first transformed to bone edge probability images directly. The registration error resulting from this comes from the registration method only, ie. it is an estimate of the accuracy with which we can estimate $T_{SM(nrr)\Rightarrow O}$. However we can also fit a shape model to each of the four data sets using the methods outlined in Chapter 3. This fitted model can then also be transformed to a bone edge probability image. Registering the ultrasound data to this shape model derived bone edge probability image gives an estimate of the combined error in estimating $T_{SM\Rightarrow M}$ and $T_{SM(nrr)\Rightarrow O}$, and hence $T_{M\Rightarrow O}$.

The three cadaver data sets used here are from female cadavers. Similar data for male cadavers was not available. The shapes of the male and female pelvis is similar enough that the registration errors (error in $T_{SM(nrr)\Rightarrow O}$) should be very nearly the same as those for a male pelvis. Hence the use of female pelves to estimate this error is valid. Whether it is also valid to estimate $T_{SM \Rightarrow M}$ by fitting the male shape model to the female data is a more complex question. This subject is discussed further in Appendix B. To account for this here we relaxed the shape constraints of the model, allowing the shape parameters to fall beyond the 3 standard deviation limit used when fitting to male MRI data. The shape model fitting here is intra-modal (CT model to CT patient) rather than inter-model (CT model to MRI) patient. In this case it was found that it was not necessary to constrain the shape model in order to achieve good results.

As shown in Section 4.6 the number and distribution of ultrasound slices used will impact on the registration results. The number and distribution of slices was therefore controlled. Two alternative strategies were tried, one used slices evenly spread across the pelvis, the other used relatively more slices from the central pubic region. During development of the registration method it was noted that weighting the algorithm to fit the pubic region more accurately gave a lower TRE at the prostate on average. This weighting can be done quite simply by using more slices from the pubic region, as classified in Table 4.2. Similarly it was noted in Section 4.6 that slices from the iliac regions were necessary for a successful registration. In practice it will be necessary to implement a system to control the distribution of slices used for registration in theatre. This will be discussed further in Chapter 6.

Here 2 distributions of slices are compared, one has an approximately even distribution of slices with 168 slices from each iliac region, and 84 from the pubis, whilst the other weights the registration to the pubis by having 180 pubis slices and 120 slices from each iliac region. In each case the same total number of slices (420) was used. The actual ultrasound slices used were randomly chosen at the start of each registration. In cases where the number of slices used was greater than the available slices (see Table 4.2) some slices (randomly chosen) were used multiple times as necessary. Whilst this will not add new information to the registration it will still weight the registration to that region.

For each data set, each CT processing method, and each ultrasound slice distribution the ultrasound to CT registration was repeated 100 times using the parameters shown in Table 4.3. Random selection of ultrasound slices was performed for each repeat registration. Each registration was started using an initial estimate for $T_{SM(nrr) \Rightarrow O}$ that was based on perturbing the known gold standard $T_{SM(nrr) \Rightarrow O}$ by a random transform of distance 10 mm (as defined by equation 4.9).

For each registration the error was determined by calculating the distance error at the "nominal prostate" points defined in Section 1.6.1. They were first transformed to the Optotrak's coordinate system using the known gold standard transform $T_{SM(nrr) \Rightarrow O}[GS]$. The transform estimated using the registration algorithm was then used, and the difference in the point positions measured. Registrations giving a TRE at the prostate greater than 10 mm were classed as registration failures. The results are presented as the RMS TRE at the prostate for successful registrations and the percentage of failed registrations.

4.9.4 Results

The results, expressed as the RMS TRE for successful registrations and a failure (TRE greater than 10mm) are shown in Table 4.4. These results are for individual runs of the gradient descent optimiser,

Parameter	Phantom	Cadaver-01	Cadaver-02	Cadaver-03
Number Of Ultrasound Slices	650	198	316	331
Number Of Non Zero Bone Edge Probabilities	7547242	931274	4281129	1802785
Initial Step Size	4 mm	4 mm	4 mm	4 mm
Ultrasound Slice Resolution (In Plane)	0.088 mm	0.12 mm	0.12 mm	0.12 mm
CT XY Resolution	0.70 mm	0.70 mm	0.70 mm	0.70 mm
CT Z Resolution	2 mm	2 mm	2 mm	2 mm
Registration Start Distance	10 mm	10 mm	10 mm	10 mm
Degrees to mm	0.25 Deg/mm	0.25 Deg/mm	0.25 Deg/mm	0.25 Deg/mm

Table 4.3: Parameters used for registration used in validation experiments.

		Data Set				
Method	Slice Distribution	Cad. 01	Cad. 02	Cad. 03	All Cads.	Phantom
Reg. Only	168:84:168	3.64 (10%)	3.19 (25%)	6.71 (17%)	4.79 (17%)	1.62 (0%)
Reg. Only	120:180:120	3.70 (8%)	4.44 (34%)	6.54 (18%)	5.03 (20%)	1.70 (0%)
Reg. & Mod.	168:84:168	7.85 (56%)	7.08 (31%)	7.85 (30%)	7.57 (39%)	4.32 (0%)
Reg. & Mod.	120:180:120	7.42 (30%)	7.68 (25%)	6.99 (15%)	7.35 (23%)	3.96 (0%)
Model Only $\sqrt{R4^2 - R2^2}$	120:180:120	6.43	6.27	2.47	5.36	3.56

Table 4.4: The registration results for the 4 data sets using bone edge images derived from the actual CT data (Registration Only) and a shape model fitted to the CT data, model (Registration and Model). The figures are expressed as a RMS TRE in mm at a nominal prostate and a failure rate in percent (RMS TRE greater than 10mm). These results are for individual runs of the gradient descent optimiser, not the repeated run and convergence check method discussed in Section 4.4. The final row is an estimate of the error induced by using a shape model to segment the pelvic bone.

not the repeated run and convergence check method discussed in Section 4.4.

The last row of Table 4.4 is an estimate of the registration error caused by segmenting the data with a shape model. In Chapter 3 we predicted that this error would be normally distributed with an RMS error of 5.3 mm. The results in the last row of Table 4.4 agree with this estimate.

Weighting the registration to the pubic region had variable results for the registration error only tests. When using the shape model however the pubic weighting was generally beneficial.

4.10 Discussion

4.10.1 Optimiser Performance and Registration Run Time

The results in table 4.4 indicate that when used on real data the gradient descent registration optimisation algorithm will fail 23% of the time. To allow for this in theatre we propose running the algorithm repeatedly and checking for convergence of the result, as discussed in Section 4.4. Based on our results it is likely that the optimiser will have to be run 20 to 30 times. At present the optimisation algorithm has an average run time of approximately 10 minutes on a desktop PC. Therefore we would require up to 300 minutes to perform the registration. To be used in theatre the algorithm must run in less than 20 minutes. For clinical implementation it will be necessary to increase the speed of the algorithm. This can be done in two ways. Firstly we will investigate faster and/or more robust optimisation algorithm. Using a more robust optimiser would reduce the need to run repeated registrations. Secondly, we could reimplement the algorithm in parallel take advantage of either a multi-processor machine or a graphics processing unit. Though the speed increase required is very large, the algorithm as currently implemented is very simplistic, so it is likely that using these approaches could yield a sufficiently fast registration method.

4.10.2 Registration Error

Two errors can be determined from the experiment in Section 4.9. One is due to the performance of the registration algorithm used on images with well defined bone edges and no ultrasound calibration errors. This is the TRE for the registration only Phantom data in Table 4.4. However, the bone edges in the ultrasound images of the phantom (a plastic/water interface) are significantly better defined than they are in ultrasound images of human pelves (a bone/soft-tissue interface). Therefore the bone edge probability images from the cadaver data are less accurate than from the phantom data. Refer back to Figure 4.3 for an example of this. Furthermore there is an additional error due to the speed of sound mismatch between the ultrasound calibration medium and human tissue. The registration error shown in rows 1 and 2 of Table 4.4 for all phantom data (4.79 and 5.03 mm) is a combination of three errors. These are the registration algorithm accuracy, the error in converting the ultrasound images to bone edge probability images, and finally the the speed of sound mismatch for the ultrasound calibration. This is the error we would expect to achieve in theatre, using our current methods.

4.10.3 Conclusion

In this chapter we have detailed our implementation of an ultrasound to CT registration algorithm. The algorithm has been optimised to account for the fact that the CT image is an approximation. The CT image is formed by warping another patient's CT to fit the prospective RARP patient's MRI. The algorithm requires significant optimisation before it will be ready for clinical use. We have also assessed the registration error for the algorithm and determined the accuracy of the transform $T_{M \Rightarrow O}$.

The error due to ultrasound to CT image registration (5.03 mm) can be projected on screen using the procedure outlined in Section 1.6.2. Figure 4.18 provides a visualisation of this error overlaid on a typical surgical scene.

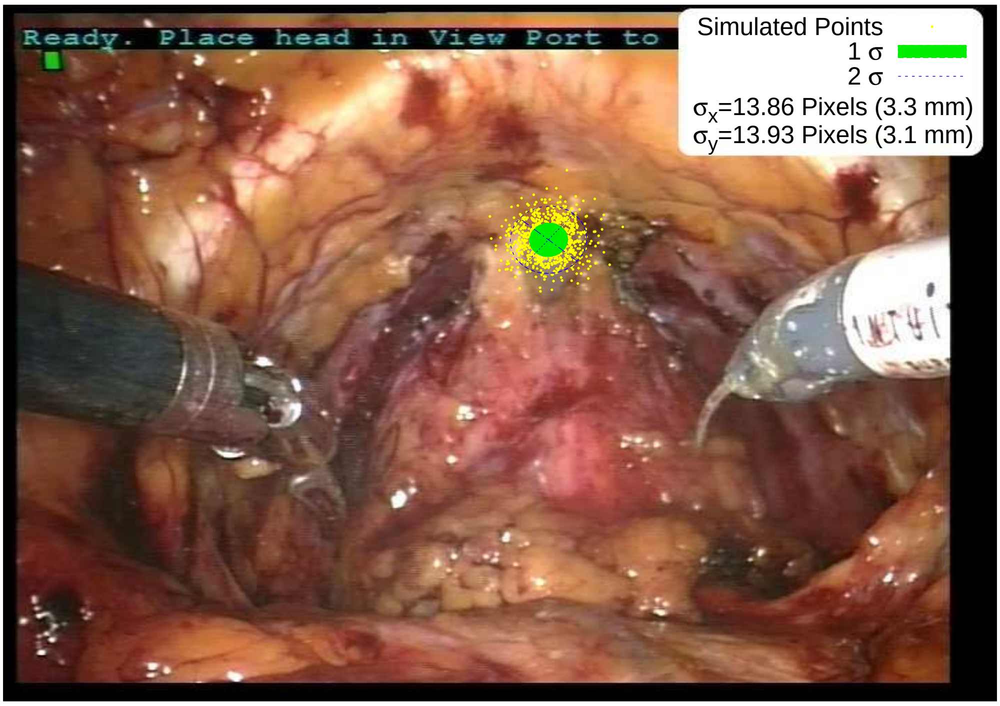
Ready. Place head in View Port to
Simulated Points
1 σ
2 σ
σx=13.86 Pixels (3.3 mm)
σy=13.93 Pixels (3.1 mm)

Figure 4.18: A landmark point (near the apex of the prostate) has been chosen and perturbed by a random Gaussian error in 3D equivalent to the registration error (5.03 mm) 1000 times. The resulting projections give an indication of the on screen error due to this error source.

Chapter 5

Endoscope Calibration and Tracking using an Optical Tracker

The preceding two chapters have placed the patient's preoperative image into the coordinate system of the Optotrak system. To complete the image guided surgical system we now find the transforms to put the image onto the endoscope screen, $T_{O \Rightarrow ES}$ from Figure 1.4. $T_{O \Rightarrow ES}$ is the product of three transforms. The endoscope's intrinsic projection charateristics $T_{EL \Rightarrow ES}$, the tracked endoscope's extrinsic transform $T_{ET \Rightarrow EL}$, and the endoscope's tracking matrix $T_{O \Rightarrow ET}$.

In this chapter we first define our methods for estimating each of these transforms. We then present a novel method for determining the error in estimating each of these transforms. Following this we present the results of our endoscope calibration and a series of experiments to test potential improvements in the tracking system.

5.1 Endoscope Tracking $T_{O \Rightarrow ET}$

The *daVinci*^{TM} endoscope is tracked using tracking markers (IREDs) mounted on the endoscope and the same Optotrak Certus tracking system as was used in the registration of the preoperative data in the preceding chapter. The Optotrak was used in preference to the *daVinci*^{TM} 's own kinematic data for three reasons. Firstly it avoids the need for an additional transform from the *daVinci*^{TM} 's coordinate system to the Optotrak's coordinate system, avoiding an additional calibration step and the accompanying errors. Secondly the system remains independent of the *daVinci*^{TM} and could be applied to any endoscopic system. Finally, results in the literature for endoscope tracking using the *daVinci*^{TM} kinematics [Mourgues and Coste-Maniére (2002)] are not encouraging.

Figures 5.1 and 5.2 show the tracking collar used and the Optotrak Certus system in theatre. The Optotrak is used to triangulate the position of each IRED on the tracking collar. The IRED positions define a rigid body which can be registered to a reference tracking collar rigid body to determine $T_{O \Rightarrow ET}$. The tracking problem is formulated as a Procustes point alignment and the optimum transform found using SVD. This is the same method as was used for the ultrasound cross, described in Section 4.2.1. As described by West and Maurer Jr (2004) the design goal of such a tracking system is to place as many tracking points as possible, spread as far apart as possible, and with their centroid positioned at the point

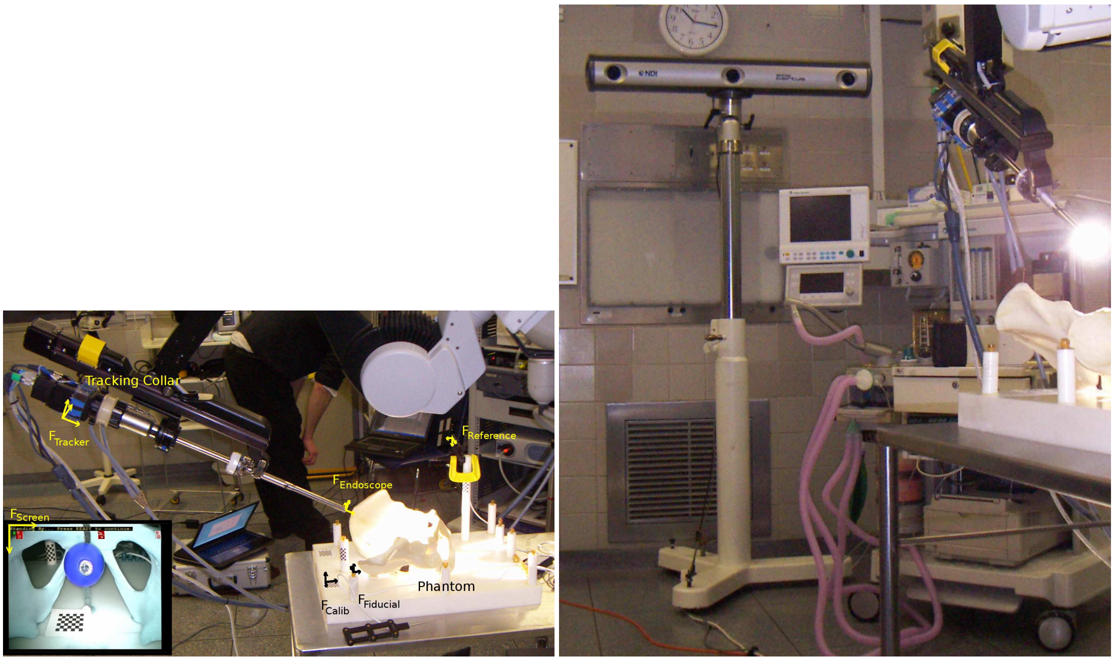

Figure 5.1: The tracking collar mounted on the *daVinci*^{TM} endoscope and the Optotrak camera system in theatre.

to be tracked. The on screen error due to the endoscope tracking accuracy will depend on the position of the object being viewed and the projection characteristics of the endoscope. This makes it difficult to compare different results from the literature. One of the simplest ways to model the endoscope tracking error is to use an anatomical point beyond the end of the endoscope as the target point. The TRE for this point can be determined using Equation 5.1 from Fitzpatrick et al. (1998). This approach was used by Shahidi et al. (2002). Here we can use the same "prostate apex" point introduced in Section 1.6.2. The geometry of the resulting tracking system is shown in Figure 5.2. The design shown in Figure 5.2 fails to meet this design goal for several reasons. The number of IREDs is restricted by the quantity of wiring required to be routed over the endoscope. The spread of the IREDs around the axis of the scope is restricted by the need to maintain clearance to the *daVinci*^{TM} arm. The most significant design problem is the distance from the tracking IREDs to the target point along the axis of the endoscope. This is necessary to maintain a line of sight on the IREDs from the Certus cameras.

The relative performance of the tracking set up can be determined by using equation 5.1 [Fitzpatrick et al. (1998)]. FLE refers to to the localisation error of a single IRED, TRE is the error at the target point, d_k is the distance of the target point from the k_{th} principal axis of the visible IREDs, and f_k is the moment of the visible IREDs about that axis.

$$\frac{\langle TRE^2(r)\rangle}{\langle FLE^2\rangle} \approx \frac{1}{N}\left(1+\frac{1}{3}\sum_{k=1}^{3}\frac{d_k^2}{f_k^2}\right) \tag{5.1}$$

The same method as in Section 4.2.1 of calculating the ratio of expected TRE to expected FLE for all possible combinations of visible IREDs and plotting the best and worst cases was used to generate Figure 5.3. Provided at least 10 IREDs are visible this ratio is less than 7. This can be compared with the same

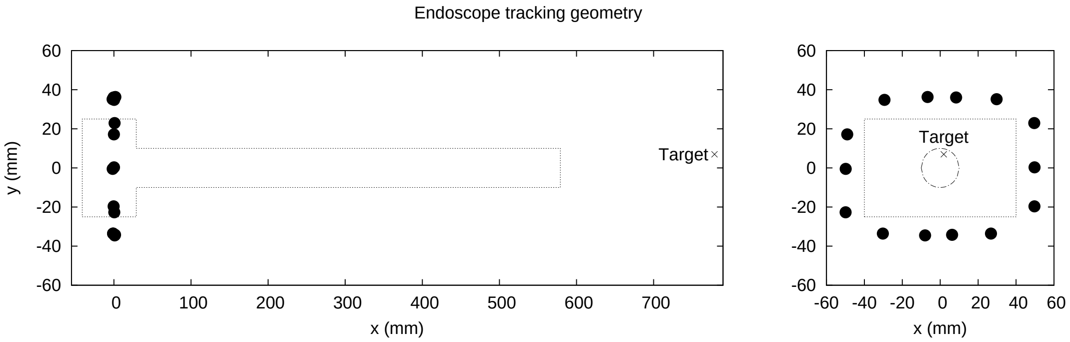

Figure 5.2: A schematic of the tracking collar and endoscope showing the approximate dimensions.

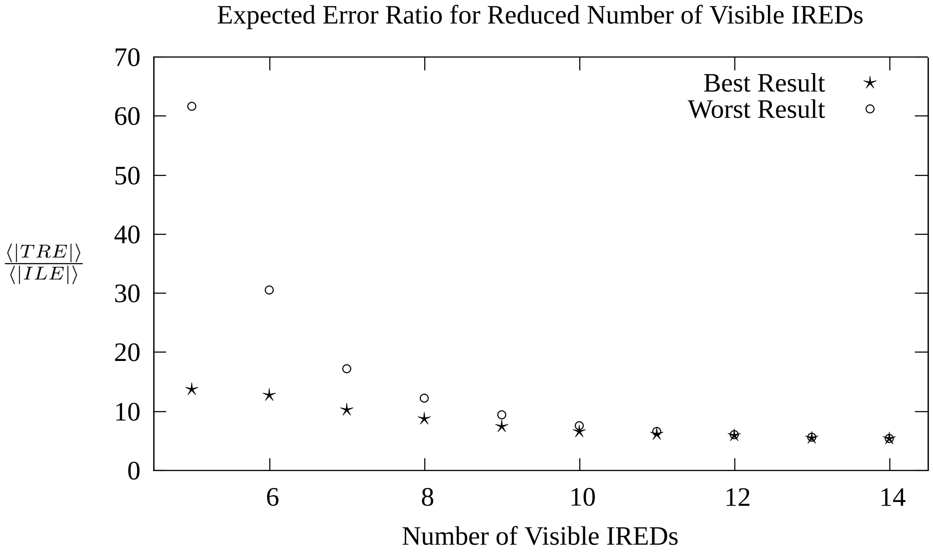

Figure 5.3: Ratio of the expected values of the absolute TRE and FLE for the endoscope tracking system. The ratio is calculated for all possible combinations of visible IREDs using equation 5.1. The best and worst case values are shown for each number of visible IREDs.

plot for the ultrasound tracking cross used in the preceding chapter, see Figure 5.3, where values closer to 2 were the norm.

If we were to assume (as is commonly done in the literature [Shahidi et al. (2002),King et al. (1999)]), that the expected value of the FLE is 0.2 mm (from the Optotrak specification, NDI (1992)) then the expected value of the TRE will be somewhere between 2 and 6 mm, assuming that at least six IREDs remain visible[1]. This error is similar to the registration error of 5 mm identified in the preceding chapter, corresponding to an on screen error of around 18 pixels (see Figure 4.18). If we were to be optimistic and use the figure of 0.02 mm FLE from Barnes et al. (2007), then the expected TRE would be less than 0.6 mm. It will be shown in this chapter that both these figures are significantly less than the errors seen in practice. Equation 5.1 assumes anisotropic, homogeneous, and normally distributed error, which is not the case when using an optical tracking system.

---

[1] It was found that in practice many of the 14 IREDs were obscured by wiring and surgical equipment. The assumption of six IREDs is based on our experience in theatre

Defining the Reference Tracking Body

As discussed for the case of the ultrasound cross in Section 4.2.1, the apparent shape of the tracking collar (the positions of the IREDs) will vary across the tracking volume due to the effect of IRED angle on tracking accuracy. Therefore a reference tracking collar defined "front on" to the Optotrak cameras may not perform well in practice, when the collar will be at various angles to the Optotrak. To allow for this the same iterative method for defining the reference frame as in Section 4.2.1 was used. For a given sequence of tracking frames, the reference IRED pattern was found by registering all frames to the first frame in the sequence, then averaging the registered IRED positions. This average frame was used as a new estimate of the reference frame. This process is repeated until the average RMS registration error converges.

5.1.1 Calibration - Calculation of $T_{EndoTrkToEndoLens}$ and $T_{Intrinsic}$

Endoscope Model

The endoscope camera was modelled as a pinhole camera, based on the model of Heikkila and Silven (1997). A point defined in the 3D coordinate system of the Optotrak ($X_{Optotrak}$) is first transformed into a coordinate system with an origin at the endoscope lens and z axis along the central ray ($X_{EndoLens}$). This transformation uses the tracking transform and the endoscope's extrinsic matrix, equation 5.2.

$$X_{EndoLens} = T_{ET \Rightarrow EL} \times T_{O \Rightarrow ET} \times X_{Optotrak} \tag{5.2}$$

The normalised image projection is now defined as Equation 5.3.

$$X_n = \begin{bmatrix} x \\ y \end{bmatrix} = \begin{bmatrix} \frac{x_{EndoLens}}{z_{EndoLens}} \\ \frac{y_{EndoLens}}{z_{EndoLens}} \end{bmatrix} \tag{5.3}$$

Radial and tangential lens distortion up to 4th order even radial terms were modelled, the normalised coordinates after distortion are given by Equation 5.4.

$$\begin{gathered} X_d = \begin{bmatrix} x_d \\ y_d \end{bmatrix} = (1 + D_{rad}(1)r^2 + D_{rad}(2)r^4)X_n + D_{tan} \\ \text{where} \quad r^2 = x^2 + y^2 \\ \text{and} \quad D_{tan} = \begin{bmatrix} 2D_{tan}(1)xy + D_{tan}(2)(r^2 + 2x^2) \\ D_{tan}(1)(r^2 + 2y^2) + 2D_{tan}(2)xy \end{bmatrix} \end{gathered} \tag{5.4}$$

where $D_{rad}(1,2)$ are the radial distortion coefficients and $D_{tan}(1,2)$ the tangential distortion coefficients. The distortion corrected normalised points are then projected to on screen pixel coordinates using Equation 5.5.

$$\begin{bmatrix} x_{EndoScreen} \\ y_{EndoScreen} \\ 1 \end{bmatrix} = \begin{bmatrix} fc(1) & 0 & cc(1) \\ 0 & fc(2) & cc(2) \\ 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} x_d \\ y_d \\ 1 \end{bmatrix} \tag{5.5}$$

where $fc(1,2)$ are the endoscope focal lengths and $cc(1,2)$ the principal points. This gives an intrinsic camera model with 8 parameters plus a 6 dimensional extrinsic transform for each view.

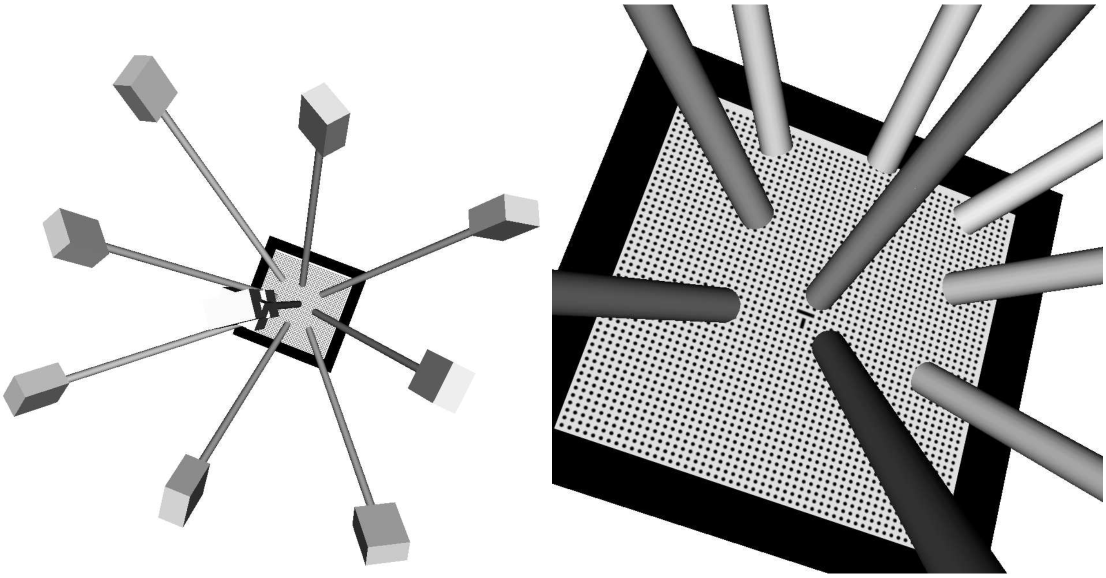

Figure 5.4: An example of the endoscope poses and calibration grid used for calibration.

Endoscope Calibration

Calculation of the endoscope's intrinsic and extrinsic parameters was done using a set of images of a regular grid designed by Wengert et al. (2006). This is visible in Figure 5.4. Determination of the extrinsic matrix is as proposed by Tsai and Lenz (June 1989) and was done using a *MATLAB*®implementation made available by Wengert et al. (2006). For accurate calibration the images of the grid must be taken from a range of view points. For tracking the *daVinci*^{TM} endoscope the possible range of views is limited by the possible motion of the endoscope and by the field of view (or characterised volume) of the Optotrak camera system. Figure 5.4 shows an (ideal) example of the pattern of endoscope orientations used for calibration. They are evenly spread on a grid spanning an angle of 60 degrees.

An automatic feature extraction algorithm supplied by Wengert et al. (2006) is first used to extract the grid points on each image. Optimisation of the global intrinsic parameters is then performed using a gradient descent algorithm, implemented in *MATLAB*® in the "Camera Calibration Toolbox", freely available on line. For N calibration images this algorithm outputs the global (common to all images) intrinsic parameters ($T_{EL\Rightarrow ES}$) and N transforms ($T_{Grid\Rightarrow Lens}[i]$). These transforms are estimated based on the visible orientation of the grid. For each image we also have the tracking transform ($T_{O\Rightarrow ET}[i]$) determined from the measured IRED positions. Equation 5.6 defines the relationship between these matrices.

$$T_{Grid\Rightarrow Lens}[i] = T_{ET\Rightarrow EL} \times T_{O\Rightarrow ET}[i]^{-1} \times T_{Grid\Rightarrow Optotrak} \tag{5.6}$$

These form a set of N equations which can be solved for $T_{ET\Rightarrow EL}$ and $T_{Grid\Rightarrow Optotrak}$. A *MATLAB*® implementation (Wengert et al. (2006)) of the method proposed by Tsai and Lenz (June 1989) was used to solve this.

5.2 Determination of Calibration and Tracking Errors

The accuracy of the preceding endoscope calibration can be assessed by measuring the back projection errors. The back projection error is the average difference (in pixels) between the location of a point on the calibration grid determined using the estimated transforms $T_{O \Rightarrow ES}$ and $T_{Grid \Rightarrow Optotrak}$ and the point on screen measured from the image. The average back projection error however does not provide an estimate of the tracking or calibration error. If only a small number of views are used for the calibration it is possible to have a low back projection error even if the estimates of $T_{EL \Rightarrow ES}$ and $T_{ET \Rightarrow EL}$ are incorrect. The optimisation algorithm can find a solution that minimises the back projection errors for the limited views used in calibration but is in fact not valid for other views. This problem can be detected in practice by using additional views that were not used in the calibration to measure the back projection error. The back projection error found will be a function of three errors. The tracking error (error in $T_{O \Rightarrow ET}$), the intrinsic calibration error (error in $T_{EL \Rightarrow ES}$) and the extrinsic calibration error (error in $T_{ET \Rightarrow EL}$). To properly understand the performance of the image guidance system it is necessary to estimate these errors individually. Estimation of these errors can be done if the endoscope calibration is performed on an endoscope where the four transforms (the three endoscope transforms plus $T_{Grid \Rightarrow Optotrak}$) are known. This cannot realistically be done for the *daVinci*^{TM} endoscope, hence we propose a Monte Carlo simulation of endoscope calibration and tracking.

5.2.1 Monte Carlo Modelling of Calibration

We start with the assumption that there will be three main sources of error in the endoscope calibration procedure. These are the endoscope tracking error due to the error in localising the individual IREDs, the point extraction error for determining the on screen location of the calibration grid centres, and potentially a scaling error in the printing of the calibration grid. We have not attempted to model any non rigid motion between the endoscope lens and the tracking collar. Examination of the endoscope suggests that this non rigid motion should be negligible, but this has not been proven. Modelling of the calibration process is then done as follows.

A set of N versions of $T_{O \Rightarrow ET}[i]$ are created. It is important that these are representative of endoscope poses that can be achieved in practice. In our case we used transforms recorded during an actual endoscope calibration. A set of grid points X_{Grid} are first defined in 3D space. We use a grid of 51 by 51 points with a centre distance of 4mm, which is the same as the grid used for calibration as shown in Figure 5.4. If a scaling error is being modelled the grid is first scaled as necessary. This is to model any inaccuracy in printing the grid. In our case we are printing the grid using a standard office laser printer so it is unlikely to be error free. The grid measures 200 by 200 mm, and we can measure its size to within 1mm easily. Therefore it makes sense to model scaling errors of $\pm 0.5\%$. The grid points are transformed to Optotrak coordinates using $T_{Grid \Rightarrow Optotrak}$, then to the endoscope lens coordinate system using $T_{O \Rightarrow ET}$ and $T_{ET \Rightarrow EL}$. We used values for these three transforms taken from an actual calibration of the endoscope. $T_{EL \Rightarrow ES}$ is then used to transform the points to on screen points. At this point we can add an error corresponding to the error in detecting the centre of the on screen grid point. Shahidi et al. (2002) cite a sub-pixel point extraction error for a similar point extraction algorithm.

Therefore we add normally distributed point extraction error with a standard deviation of up to 1 pixel at this point. The point extraction algorithm used [Wengert et al. (2006)] does not successfully extract every visible point, so some points are also removed at this stage. The removal of points is based on that achieved during an actual calibration. This gives a set of on screen points for each simulated image. These are then passed, along with unscaled points in 3D, to the calibration algorithm which estimates $T_{EL \Rightarrow ES}$ and $T_{Grid \Rightarrow Lens}[i]$. The next stage is to generate a matching set of estimates for $T_{O \Rightarrow ET}[i]$.

$T_{O \Rightarrow ET}[i]$ is estimated for each view as follows. The reference version of the tracking collar IREDs are first put into Optotrak coordinates using the known gold standard transforms $T_{O \Rightarrow ET}[i]^{-1}$. As shown in Figure 5.3 the tracking error will depend on the number and configuration or IREDs visible. Therefore some of the IREDs are now removed. The removal of IREDs is based on the IRED visibility states observed during an actual calibration procedure. The remaining IREDs are now perturbed by an error function. The perturbed IREDs are registered back to the reference frame to give an estimate of $T_{O \Rightarrow ET}[i]$. The error function used for the perturbation of the IREDs will have a great impact upon the the results. Determining an appropriate error function is a key part of the novelty of this work and is discussed now.

5.3 Determination of IRED Localisation Error

The ILE for an Optotrak tracking system is known to be anisotropic and dependent on the position and orientation of the individual IREDs [Wiles et al. (2004)]. It has also been shown to markedly increase when tracking moving IREDs [Barnes et al. (2007)]. For endoscope calibration the endoscope is held static during data collection, so the error due to IRED motion can be ignored. Wiles et al. (2008) present an analytic method for modelling the accuracy of a point based registration with anisotropic normally distributed error. However the ILE is correlated to IRED position and IRED angle. The calibration process uses a set of discrete views with the IREDs in different orientations and positions so it is unlikely that the ILE will be normally distributed. Therefore we require a method to model the ILE that does not assume a normal distribution. We do not know at this stage what the distribution of ILE will be, so we propose a method to measure the ILE that occurs in an actual calibration process. After the ILE is estimated we can attempt to fit an analytic model to it. However, this is not necessary as long as the IRED positions and angles used to estimate the ILE are representative of those encountered in the tracking process to be simulated. In this case we can use the measured ILEs directly as a look up table in the simulation.

The obvious method for determining ILE is to track the IREDs with a separate, more accurate system, such as is done by Barnes et al. (2007). The IRED positions measured by the Optotrak can then be compared to a set of known gold standard IRED positions. For the Optotrak system however this method has a major shortcoming. As the ILE is dependent on the position and orientation of the IREDs the measurement must be representative of the positions occurring in use. In general this is not practical. We therefore propose a method that estimates the ILE using the IRED positions recorded in use.

A set of measured marker locations $X_m[i]$ are collected from an actual endoscope calibration. Direct measurement of the marker localisation error vectors $E_{ML}[i]$ is not possible as the actual marker locations remain unknown. However, their position can be estimated using the remaining markers on the rigid body. These can be registered to a reference rigid body, in this case using the orthogonal Procustes formulation and singular value decomposition, Fitzpatrick et al. (2000). The inverse of the registration transform can then be used to give an estimated position vector $X_R[i]$ for the marker. The measured error vector $E_m[i]$ is defined by equation 5.7.

$$E_m[i] = X_m[i] - X_R[i] \tag{5.7}$$

The marker position $X_R[i]$ estimated by registration, however, is not the actual marker position, which remains unknown. The measured error is therefore the combination of the ILE ($E_{ILE}[i]$) and the marker registration error $E_R[i]$, equation 5.8.

$$E_m[i] = E_{ILE}[i] + E_R[i] \tag{5.8}$$

$E_R[i]$ is unknown, however it is analogous to the TRE which has been studied by Fitzpatrick et al. (1998) and was used in equation 5.1. Equation 5.1 however assumes isotropic, normally distributed ILE. A more recent method by Wiles et al. (2008) allows anisotropic, though still normally distributed, ILE. Equation 5.9 presents a method for calculating the covariance matrix of the expected value of TRE for anisotropic normally distributed ILE. Whilst this will not be entirely accurate in this application, we rely on the assumption that the estimated value of TRE is small in comparison to the measured value of E_m, so the error due to the assumption of normally distributed error is small.

$$\begin{gathered}(\Sigma_{TRE})_{ij} \approx \epsilon^2 \left(\frac{\sigma_{ij}}{N} + \sum_{k \neq i}^{K} \sum_{m \neq j}^{K} \frac{r_k r_m \Psi_0}{\left(\Lambda_{kk}^2 + \Lambda_{ii}^2\right)\left(\Lambda_{mm}^2 + \Lambda_{jj}^2\right)} \right) \\ \text{where} \\ \Psi_0 = \Lambda_{kk}^2 \delta_{km} \sigma_{ij} - \Lambda_{kk}^2 \delta_{kj} \sigma_{im} - \Lambda_{ii}^2 \delta_{im} \sigma_{kj} + \Lambda_{ii}^2 \delta_{ij} \sigma_{km}\end{gathered} \tag{5.9}$$

Where

- Σ_{TRE} Is the covariance matrix of the expected value of TRE (or E_R).
- σ Is the covariance matrix of the expected value of FLE (or ILE).
- N Is the number of IREDs visible.
- Λ Are the eigenvalues of the SVD of the covariance matrix of the remaining visible IREDs. A function of the geometry of the remaining visible IREDs (X_{rb}).
- r Are the distances of the target IRED to the principal axes of the IRED configuration
- δ Is the Kronecker Delta.
- K is the dimension, in this case 3.

This determines the expected value of E_R rather than the actual value for an individual measurement as in equation 5.7. We cannot therefore calculate the individual ILEs using equation 5.9, however we can determine the expected value of ILE for a number of frames as in equation 5.10. $\langle |E_R| \rangle$ is equivalent to

$(\Sigma_{TRE})_{ij}$ from equation 5.9. $X_R[i]$ and $X_m[i]$ are known. Therefore $\langle |E_{ML}| \rangle$ can be determined, with equation 5.10 where N is the number of IREDs sampled.

$$\langle |E_{ML}| \rangle = \sum_{i=1}^{N} \frac{|X_R[i] - X_m[i]|}{N} - \langle |E_R| \rangle \tag{5.10}$$

However, $\langle |E_R| \rangle$ is a function of $\langle |E_{ML}| \rangle$ as well as the geometry of the remaining visible IREDs, (X_{rb}). Therefore we use an iterative process to determine $\langle |E_{ML}| \rangle$ as in equation 5.11. Here k is the iteration step.

$$\langle |E_{ML}| \rangle_{k+1} = \sum_{i=1}^{N} \frac{|X_R[i] - X_m[i]|}{N} - \langle |E_R| \rangle \ \text{ where } \ \langle |E_R| \rangle = \mathcal{F}(\langle |E_{ML}| \rangle_k, X_{rb}) \tag{5.11}$$

An isotropic ILE with an expected value of 0.2 mm was used as an estimate for $\langle |E_{ML}| \rangle_0$ and approximately 20 iterations were required before $\langle |E_{ML}| \rangle_{k+1} = \langle |E_{ML}| \rangle_k$ to machine accuracy, typically taking less than a minute. To simplify the calculation we assumed that $\langle |E_{ML}| \rangle_i$ was the same for all markers in the sample. This will not be strictly true as $\langle |E_{ML}| \rangle$ will depend on marker position and orientation. For greater accuracy a different estimate of $\langle |E_{ML}| \rangle$ should be used when calculating $\langle |E_R| \rangle$ for each marker. It should be noted that each marker point will still have a different value of $\langle |E_R| \rangle$ as X_{rb} varies.

Frames with less than 7 visible markers were not used. Below this threshold a few extremely high values of $\langle |E_R| \rangle$ could cause $\langle |E_{ML}| \rangle_{k+1}$ to become negative, halting the iteration process. Badly registered frames were detected and discarding by comparing the root mean square registration error with $3(1 - (2\langle E_{ML} \rangle^2)/N$, where N is the number of IREDs used for registration.

The next stage is to use the expected values of ILE to build a model that can be used for simulation of the endoscope calibration or tracking process. As the resulting errors did not follow any analytic statistical distribution we elected to use the measured errors to build a look up table that could be randomly sampled during the simulation process. Equation 5.12 defines how this was done, with N being the number of markers in the data set. Lower case is used to denote that this was done in each direction independently.

$$e_{ML}[i] = e_m[i] - \langle |e_R[i]| \rangle \times \frac{|e_m[i]|}{e_m[i]} \qquad \text{for } i = 1 \text{ to } N \tag{5.12}$$

The resulting values of e_{ML} should not be used as an actual measure of error for a given marker, as the actual value of e_R remains unknown. However, for Monte-Carlo simulation where a large number of repeats are used the errors will average out, so the expected value of e_{ML} will be correct.

A second look up table of marker visibility states was also created. In our application individual markers are not always visible. To enable accurate modelling the same marker visibility states as observed in practice must be used in the Monte-Carlo modelling. To enable this, each possible combination of visible markers was represented as a 14 bit binary number and the number of occurrences of this IRED combination in the sample recorded. This was then divided by the total number of samples, to give a probability of a particular state occurring.

We have now presented a method for calibration and tracking the endoscope together with a method to estimate the ILE to enable simulation of calibration and tracking. The remainder of this chapter is divided into sections detailing experiments to validate these methods. We first present the results of an actual endoscope calibration.

5.4 In Theatre Calibration of Endoscope

5.4.1 Aim

Here we perform a calibration of the *daVinci*^{TM} endoscope. In addition to determining the accuracy of the calibration we also seek to use the data gathered to estimate the ILE

5.4.2 Method

The calibration method used was as defined in Section 5.1.1. A calibration grid image as in Figure 5.4 was placed on a table in theatre. The endoscope was moved into various positions to view the grid. The intent was to spread the views as widely as possible, while still allowing sufficient IREDs to be tracked. The endoscope was held static at each view point and approximately 60 frames of video (approximately 8 seconds) captured. Both the video image and the IRED positions were averaged over the collected frames. The average IRED positions were used to determine estimates of $T_{O \Rightarrow ET}$ for the collected views. The averaged images and the estimates of $T_{O \Rightarrow ET}$ were used to estimate $T_{ET \Rightarrow EL}$, $T_{EL \Rightarrow ES}$, and $T_{Grid \Rightarrow Optotrak}$. The average back projection error was measured.

5.4.3 Results

Number and Location of Views

24 views were used. 3 of these were discarded as they were not adequately tracked. 5 more were not used as the automatic point extraction algorithm failed to determine the location of the grid points. This left 16 frames for calibration. Figure 5.5 shows the endoscope positions used. The endoscopes lens was between 156 and 52 mm from the grid, with an average distance of 102 mm. The angle from the horizontal plane was between 27 and -23 degrees, with a mean angle of 4 degrees. The angle from the vertical plane was between 26 and -22 degrees, with a mean angle of -4 degrees. The Optotrak Certus specifies a “characterised volume” which defines the volume in which the tracking accuracy of the machine has been checked prior to leaving the factory. Figure 5.6 shows the positions of the centroids of the IREDs for each frame used in the calibration process, along with the characterised volume of the Optotrak. All IREDs fall within the characterised volume.

Calibration Results

Equation 5.13 and Table 5.1 present the estimated endoscope parameters.

$$T_{ET \Rightarrow EL} = \begin{bmatrix} -0.9401 & -0.3401 & -0.0228 & 8.2053 \\ 0.3402 & -0.9404 & -0.0011 & 5.4197 \\ -0.0210 & -0.0088 & 0.9997 & 561.30 \\ 0.0000 & 0.0000 & 0.0000 & 1.0000 \end{bmatrix} \tag{5.13}$$

The average back projection error for the calibration was 73.63 pixels. Figure 5.7 shows three examples of the collected calibration images and the grid points projected onto the images using the estimated

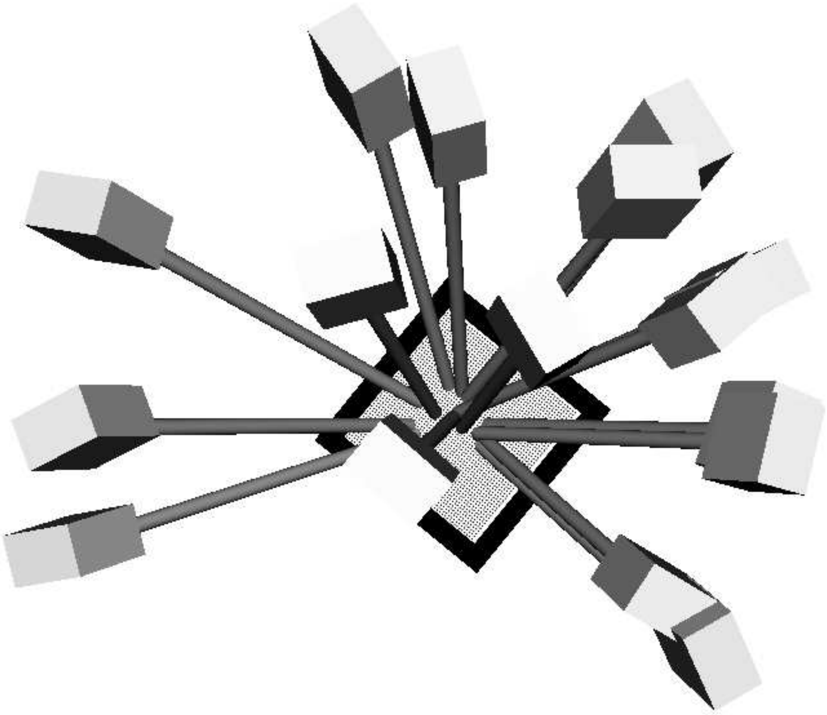

Figure 5.5: The 16 endoscope positions used for endoscope calibration.

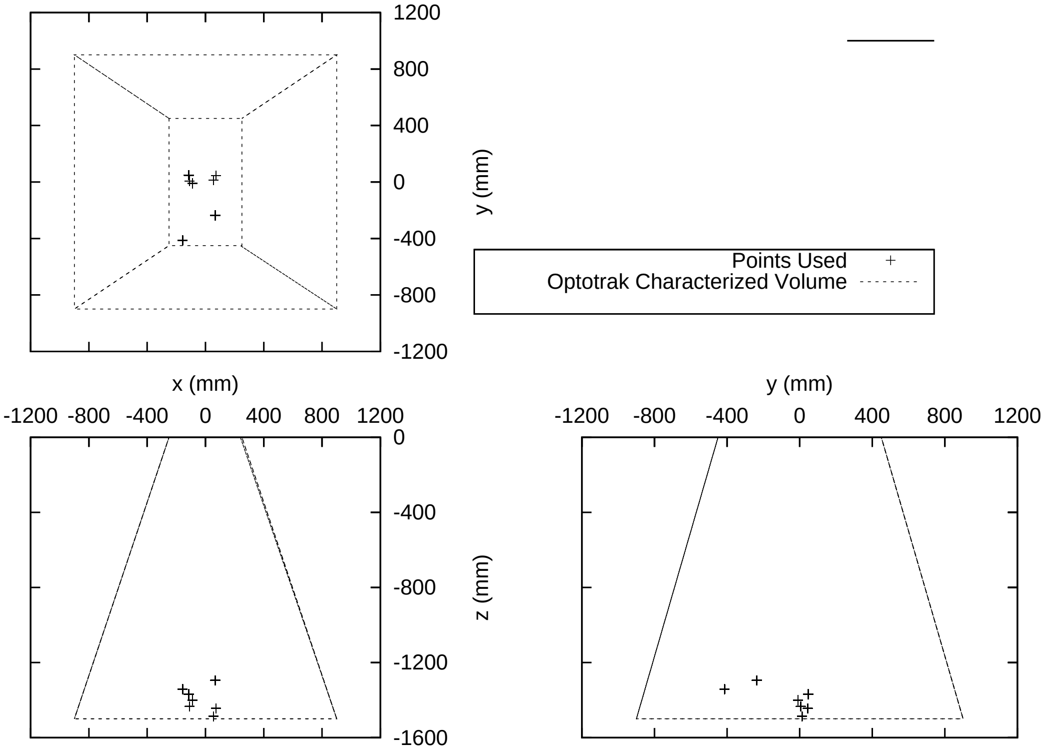

Figure 5.6: The centroids of the IREDs for the frames used in endoscope calibration. The measured IRED positions all fall within the characterised volume of the Optotrak.

$fc(1)$	$fc(2)$	$cc(1)$	$cc(2)$	$D_{rad}(1)$	$D_{rad}(2)$	$D_{tan}(1)$	$D_{tan}(2)$
839.36	908.55	428.60	257.94	-0.2743	0.6165	0.0038	0.0109

Table 5.1: The calculated intrinsic parameters for the endoscope.

calibration parameters. The back projection error achieved is very high when compared to the on screen errors due to registration calculated in Chapters 3 and 4. This indicates that for the image guidance system to function, further work is required here. The remainder of this chapter aims to determine the cause of the back projection error and hence methods to reduce it.

Calculation of IRED Visibility and IRED Localisation Error

Figure 5.8 shows the proportion of frames for which of the 14 IREDs were visible during the calibration process. Tracking of the IREDs on the top and bottom of the collar was only successful around 50 % of the time.

The method outlined in Section 5.3 was used to estimate the ILE occurring during the calibration process. After the rejection of frames with less than 7 visible IREDs, 5027 measurements of individual IRED positions were available. An isotropic error of 0.2 mm was used as a starting estimate for the ILE and an iterative process as per equation 5.11 used to find the expected values of the ILE. 24 iterations were required before the ILE converged, taking approximately 30 seconds. The expected values in each direction were $\langle|E_{ML}|\rangle_x = 0.053$, $\langle|E_{ML}|\rangle_y = 0.047$ and $\langle|E_{ML}|\rangle_z = 0.150$. Equation 5.12 was then used to create an estimate of e_{ml} for each frame. These estimates for each direction are shown in Figure 5.9. It is clear from Figure 5.9 that the errors are anisotropic, being particularly high in the z direction, and irregularly distributed. The irregular distribution is caused by the fact the sample is based on 16 discrete views.

IRED Localisation Error Versus Position and Angle

Wiles et al. (2004) note that the magnitude of the ILE is dependent on the position and angle of the IREDs relative to the Optotrak camera. They do not however define this dependence. Because of this dependence the ILE shown in Figure 5.9 can only be used to model endoscope views that are similar to those used during calibration. A more useful result would be to determine a relationship between the ILE and the position and/or angle of the IRED relative to the Optotrak. To this end the values of e_{ML} plotted in figure 5.9 were plotted against the position and angle of the IRED in Optotrak coordinates. Rather than plotting the individual directional errors the error magnitude ($\sqrt{e_{ml}X^2 + e_{ml}Y^2 + e_{ml}Z^2}$) was plotted. Figure 5.10 plots the error magnitude versus the X,Y,Z position and the angle from the Optotrak camera lens normal (the Z axis). Based on Figure 5.10 it was decided not to pursue this further. The limited sample size at each position/angle limits the reliability of the results.

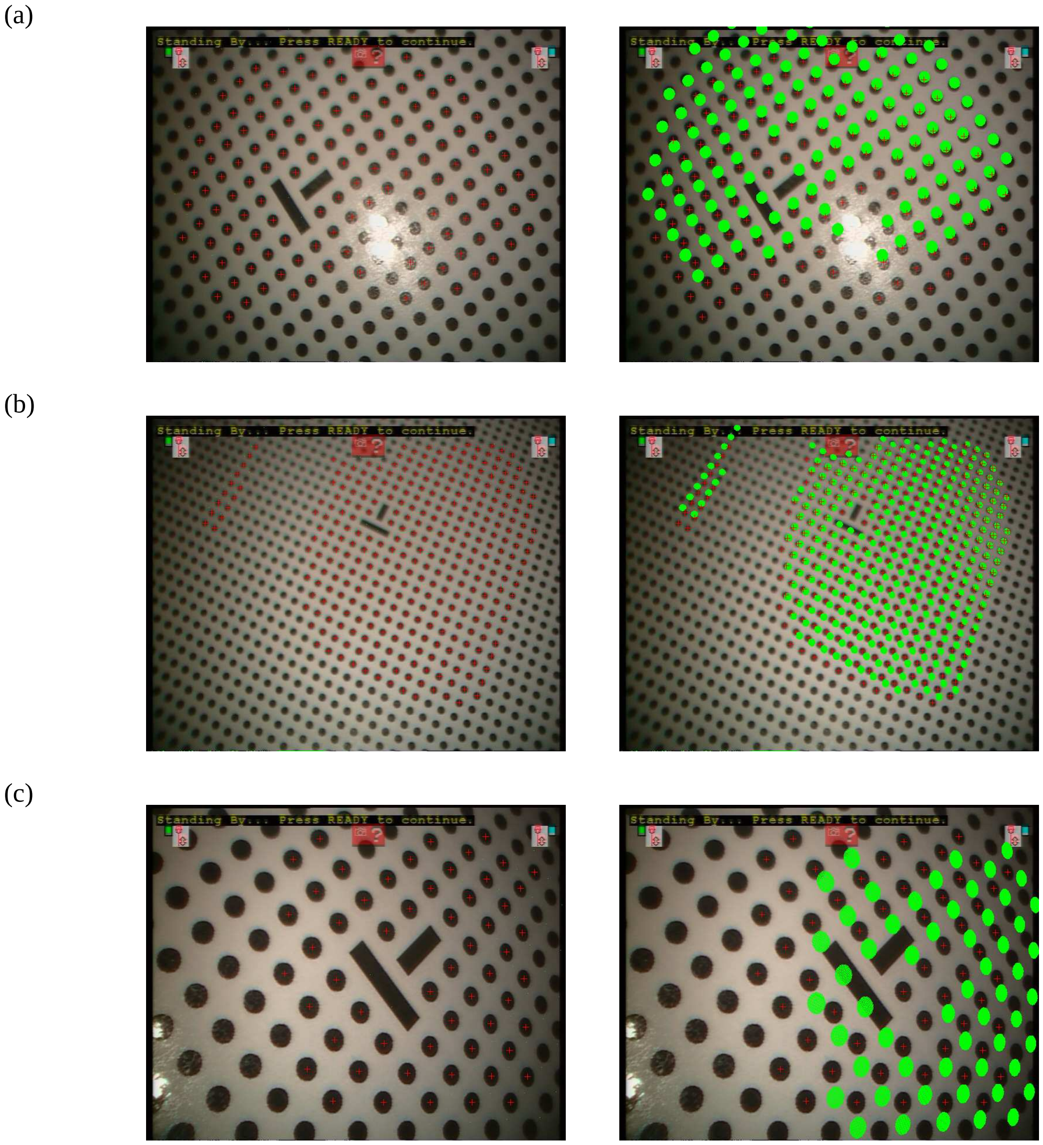

Figure 5.7: Three examples of the calibration images (image with extracted points on left) and the projected grid points (on right). Figure (a) shows the effect of a specular reflection on the automatic point extraction algorithm. Figure (b) shows a case where the automatic point extraction algorithm has failed to extract a large number of points. Figure (c) shows a case at the closest approach to the grid.

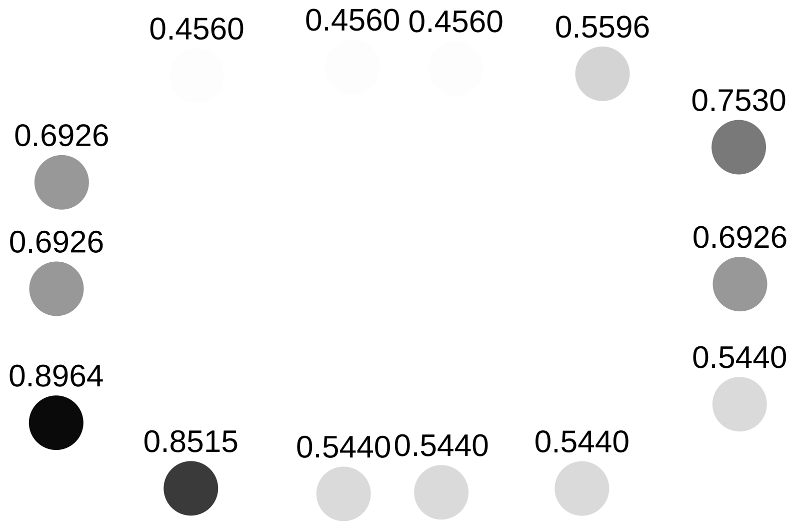

Figure 5.8: The Proportion of total frames that each IRED was visible during the calibration process.

5.5 Estimation of Error Due to Calibration

5.5.1 Aim

The preceding experiment determined the accuracy for a single calibration of the *daVinci*^{TM} endoscope. In practice the endoscope would require recalibration prior to each procedure. Therefore we are not particularly interested in the result of a single calibration. We want to repeat the calibration many times to get an expected value for the errors due to calibration. As the actual values of the intrinsic and extrinsic parameters for the preceding experiment are unknown, the actual errors due to calibration also remain unknown. We cannot calculate the error in estimating $T_{ET \Rightarrow EL}$ and $T_{EL \Rightarrow ES}$ from the average back projection error. We need to repeat the experiment using an endoscope for which these two transforms are known. The aim of this experiment is to measure the error due to calibration for repeated calibrations. The error can be measured as an on screen projection error in line with Section 1.6.2 by projecting a single point onto the screen using the values of $T_{ET \Rightarrow EL}$ and $T_{EL \Rightarrow ES}$ estimated by the calibration process.

5.5.2 Method

The calibration simulation method was described in Section 5.2. The simulated calibration was repeated 100 times. The average back projection errors were calculated for each calibration. To account for the fact that in Section 5.4 the IRED positions were averaged over a number of frames, the values of $X_m[i]$ used to build the look up tables for e_{ML} were averaged over 30 frames.[2] A normally distributed grid

---

[2] The effect of averaging the values of $X_m[i]$ was tested over averaging periods from 1 to 60 frames, and was found to be negligible. This is because the ILE is due not to random error for a static IRED, but rather the change in apparent IRED position as the IRED is imaged in different locations and angles. Because the effects of averaging are negligible we have not discussed them further in this thesis.

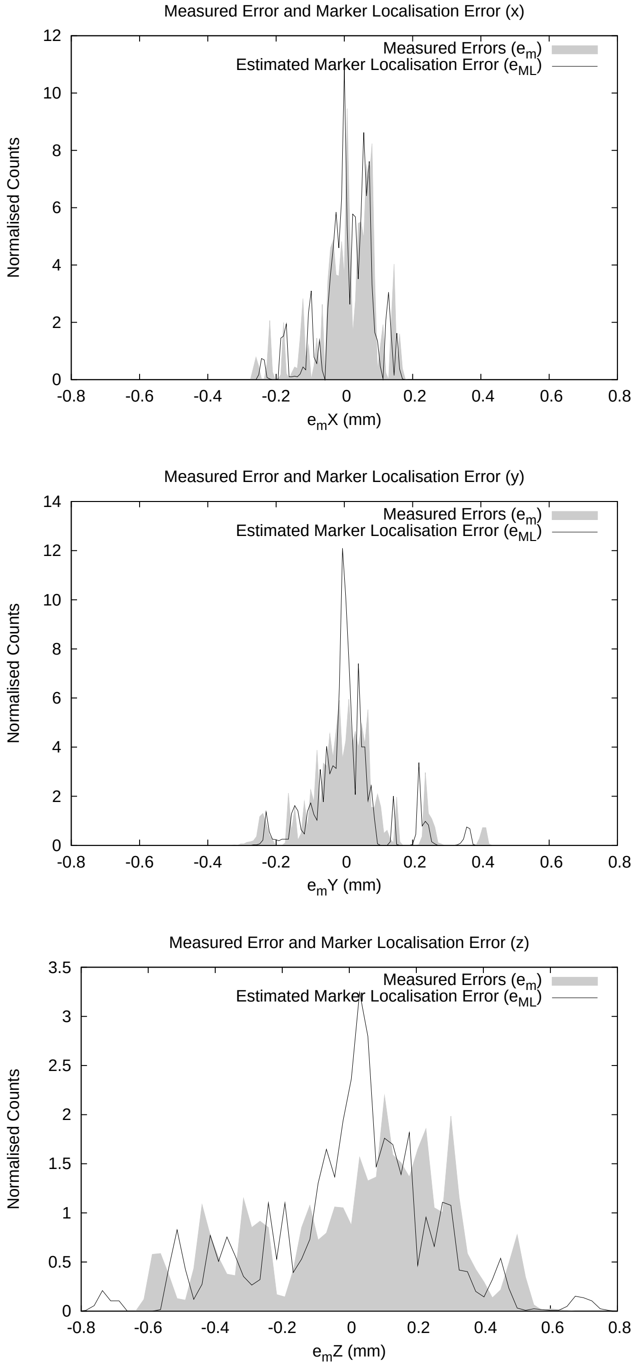

Figure 5.9: The measured marker errors (e_m) and the estimated marker localisation errors (e_{ml}) in each direction for the data collected during calibration. Equation 5.12 is used to estimate e_{ml} using e_m. Equation 5.12 slightly reduces the measured errors to account for the error in the registration process used to estimate the actual marker positions.

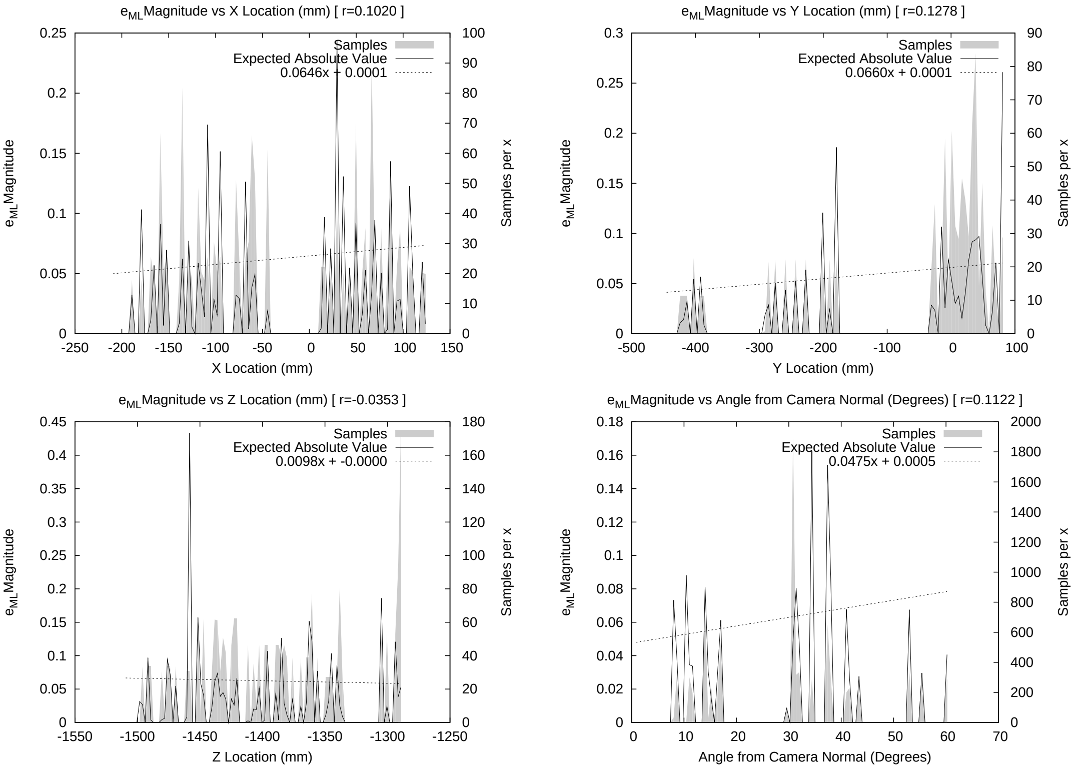

Figure 5.10: Examples of the IRED tracking error versus IRED position and angle. The irregular sampling and measurement noise in the data means that though correlations exist, they are of limited predictive value.

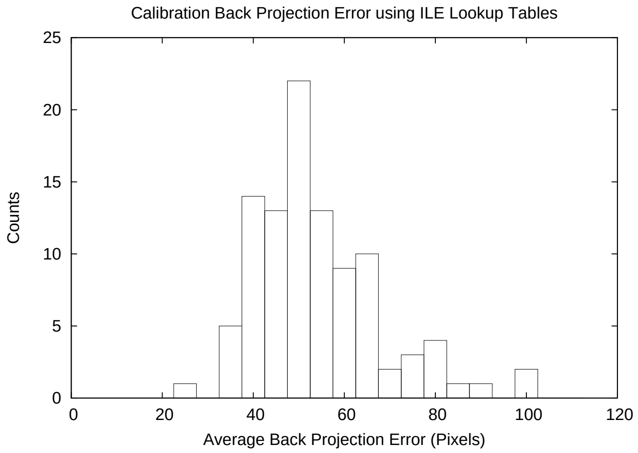

Figure 5.11: A histogram of the average back projection error for 100 simulated calibrations.

point extraction error with a standard deviation of 0.8 pixels was used.

5.5.3 Data

The values of e_{ML} determined in Figure 5.9 were used to perturb the IREDs. The intrinsic and extrinsic parameters and the position of the calibration grid found in Section 5.4 were used. The endoscope poses shown in Figure 5.5 were used.

5.5.4 Results

Figure 5.11 shows a histogram of the average back projection error for the 100 simulations. The average back projection error for the actual calibration (Section 5.4) was 73.63 pixels. This falls within the expected range shown in Figure 5.11, indicating that the simulation may be a plausible model of the actual calibration. Without further real calibrations it is not possible to disprove this hypothesis.

Calibration Error

It is possible to model the on screen error due to calibration as per Section 1.6.2. The nominal prostate apex point is here projected onto the screen using the 100 different estimates of $T_{ET \Rightarrow EL}$ and $T_{EL \Rightarrow ES}$ found through simulation. This gave an on screen error of 42.87 pixels, as shown in Figure 5.12. The error in estimating $T_{EL \Rightarrow ES}$ contributed only 2.89 pixels, while the error in estimating $T_{ET \Rightarrow EL}$ contributed 42.26 pixels.

Figure 5.13 is a scatter plot of the calibration back projection error versus the actual error due to the calibration. The calibration back projection error cannot be used as a predictor of the actual error due to calibration.

5.5.5 Comparison of Actual Versus Simulated Errors

The accuracy of the Monte-Carlo simulation of the calibration process can be validated by comparing the distribution of marker errors, in this case e_m. We make the assumption that if the distribution of e_m for the simulated data appears similar to the actual distribution of e_m then the Monte-Carlo model should be realistic. Figure 5.14 shows overlays of e_m in each of the x, y and z directions for the actual data in

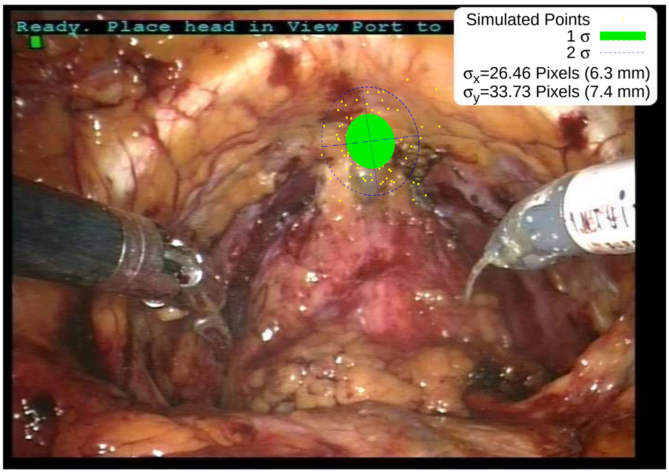
Ready. Place head in View Port to
Simulated Points
1 σ
2 σ
σ$_x$=26.46 Pixels (6.3 mm)
σ$_y$=33.73 Pixels (7.4 mm)

Figure 5.12: The on screen projection error due to the error in determining $T_{ET \Rightarrow EL}$ and $T_{EL \Rightarrow ES}$, based on 100 simulated calibrations.

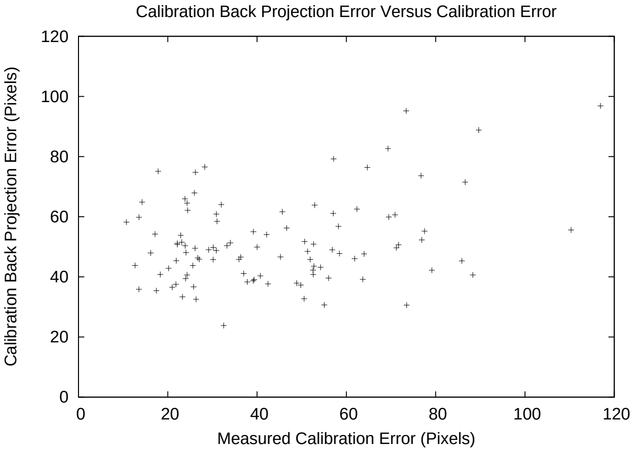

Figure 5.13: A scatter plot of the calibration back projection error versus the actual error due to the calibration.

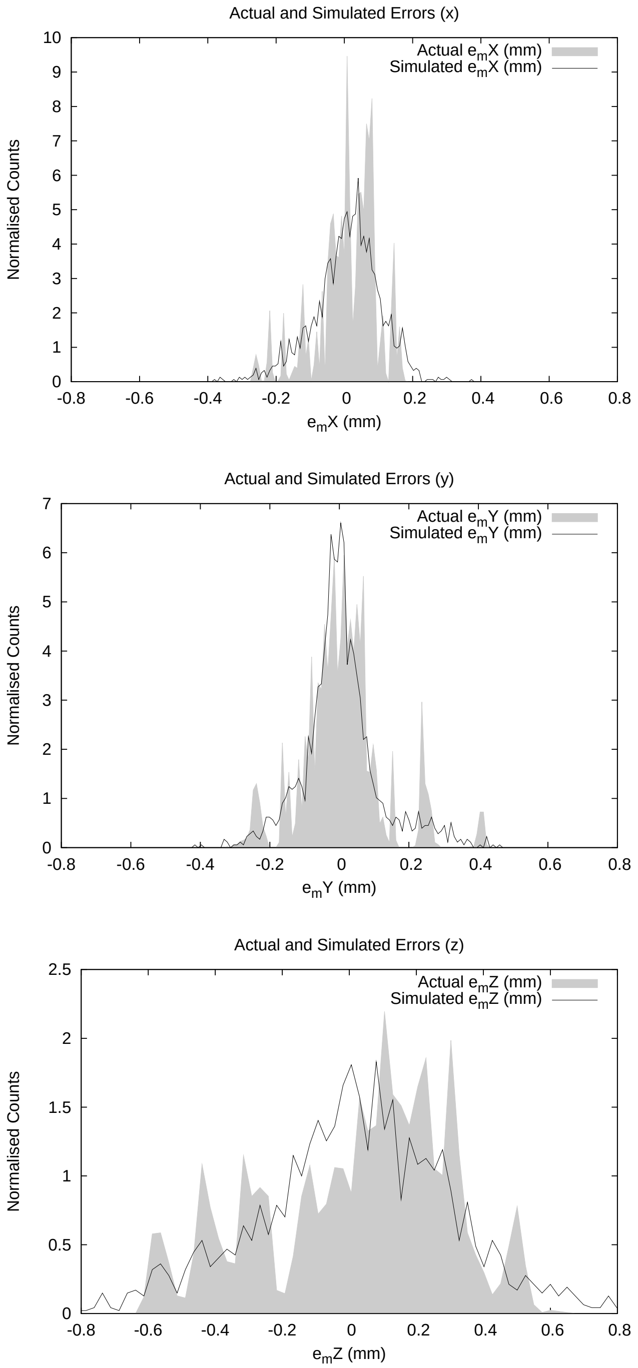

Figure 5.14: Plots of the measured marker error e_m in the x, y and z directions. The grey background is the distribution of errors measured in the actual data (from Section 5.4). The black line is the distribution of e_m for a simulated calibration.

Section 5.4 and the simulated data generated during a simulation of the calibration process. Quantitative comparison of the overlaid distributions is very difficult do to their irregular nature. Qualitatively the simulated error distribution appears to be a smoothed version of the actual errors.

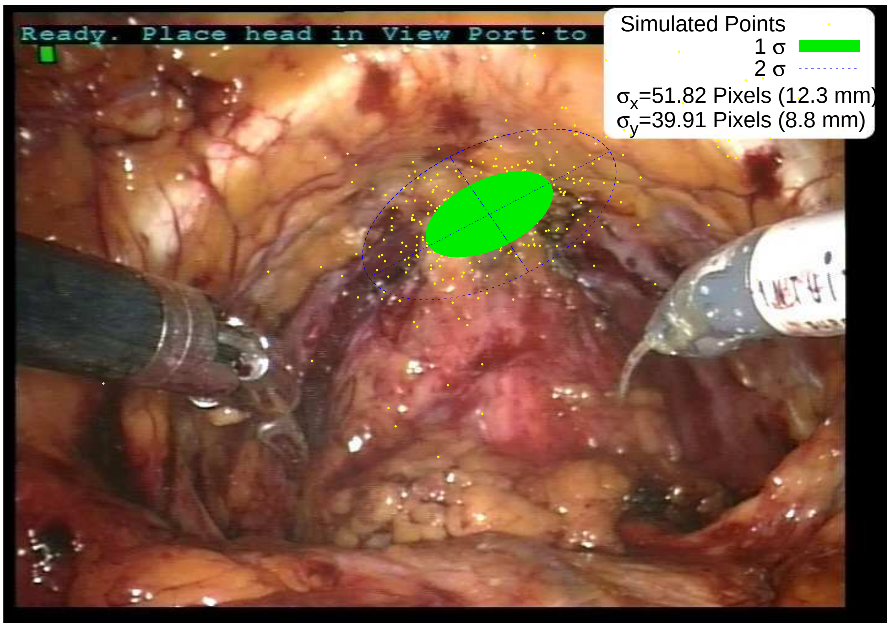

Figure 5.15: The on screen error due to tracking the endoscope. The errors are large and anisotropic.

5.6 Estimation of On Screen Error due to IRED Tracking Error

5.6.1 Aim

To use the ILEs determined in Section 5.4 to determine the tracking error of the system.

5.6.2 Method

The same ILEs used for the calibration can be used to perturb the IREDs for projection of on screen error using the simulation method outlined in Section 1.6.2 to determine the on screen error due to the ILE.

5.6.3 Data

The ILEs shown in Figure 5.9 are used to perturb the IREDs.

5.6.4 Results

The tracking errors and their distribution are shown overlaid on a typical surgical scene in Figure 5.15. The errors are very large in comparison to the on screen errors due to MRI segmentation determined in Chapter 3 and ultrasound to MRI registration determined in Chapter 4. It is clear that unless this tracking error is reduced the image guided surgery system will be of very limited use.

5.7 Effect of Errors on Calibration Accuracy

5.7.1 Aim

In Section 5.5 we defined three main sources of calibration error. These are the grid point extraction error in pixels, the error due to a misprinted calibration grid and the ILE. This experiment aims to determine the effect of these errors on the calibration accuracy.

5.7.2 Method

Grid Point Extraction Error

In Section 5.5 we assumed that the algorithm that finds the centre of grid points on screen has a normally distributed error with a standard deviation of 0.8 pixels. This was based on published results for a similar algorithm [Shahidi et al. (2002)] claiming a sub pixel accuracy. We have not measured the accuracy of the algorithm used here. Here we see what effect this error may have on calibration accuracy. The method in Section 5.5 was repeated with extraction errors of 0.0, 0.5, 1.0, 1.5 and 2.0 pixels standard deviation. Calibration simulation at each extraction error was repeated 20 times. The calibration error was measured using the projection method from Section 5.5. The average projection error was then plotted against grid point extraction error.

Grid Scaling Error

In Section 5.5 we assumed that the calibration grid was perfectly printed, so that the distance between grid centres was exactly 4.0 mm. However we have not used a perfectly printed grid, rather we have printed the grid using a commercial laser printer, so the actual grid centres may not be exactly 4.0mm apart. We can check for scaling errors by measuring the printed grid. The grid is 200 by 200 mm, and we can measure it easily to within 1mm. Therefore it is possible that the actual grid point centres may be between 3.98 mm and 4.02 mm. To test the effect of this error the calibration simulation was repeated with actual grid centre distances of 3.98, 3.99, 4.00, 4.01, and 4.02 mm. The simulation was repeated 20 times at each grid size and the results as projection error plotted against actual grid size.

IRED Tracking Error

To see what reducing the IRED tracking error will have on the calibration accuracy the calibration simulation from Section 5.5 was repeated with reduced IRED tracking errors. To preserve the anisotropic and irregular distribution of the ILE the same look up table as in Section 5.5 was used to perturb the IRED, however a scaling constant S was added to equation 5.12 to give equation 5.14.

$$e_{ML}[i] = S(e_m[i] - \langle |e_R[i]| \rangle \times \frac{|e_m[i]|}{e_m[i]}) \qquad \text{for } i = 1 \text{ to } N \tag{5.14}$$

S was varied from 0.0 to 1.0 in steps of 0.1. At each step the calibration simulation was repeated 20 times. The results are plotted as projection error due to calibration versus S.

5.7.3 Results

Grid Point Extraction Error

Figure 5.16 shows the effect of changing the grid point extraction error. Changing the grid point extraction error has a significant effect on calibration error. We used an off the shelf algorithm for grid point extraction, the results in figure 5.16 suggest that it would be worthwhile to further investigate the performance of this algorithm.

Grid Scaling Error

Figure 5.17 shows the effect of grid scaling errors over the range that is likely to occur in practice. The effect shown in Figure 5.17 is minimal, indicating that little is to be gained by more accurate printing of

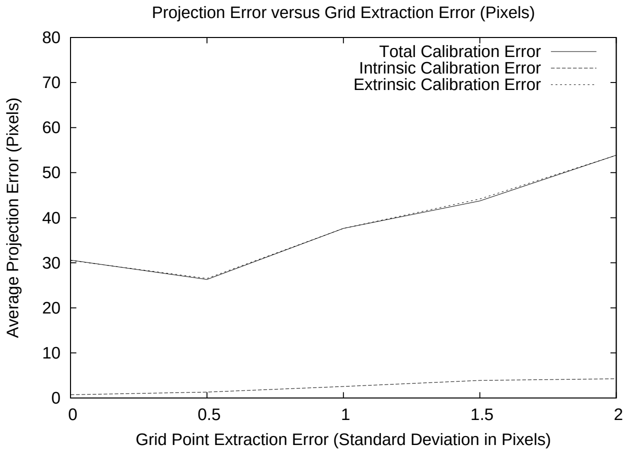

Figure 5.16: The projection error due to calibration versus the error in the automatic grid point extraction algorithm. The grid point extraction error has a significant effect on the accuracy of the calibration.

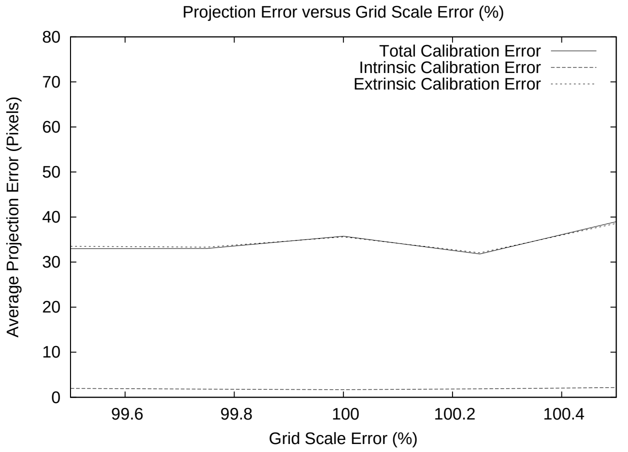

Figure 5.17: The projection error due to calibration versus the scaling error in printing the grid. Over the range likely to occur in practice the grid scaling error does not have a significant effect on calibration accuracy.

the calibration grid.

IRED Tracking Error

Figure 5.18 shows the effect of reducing the IRED tracking error. The calibration error reduces significantly as the tracking error reduces. Investigation of methods to improve the tracking during calibration should yield an improved calibration.

5.8 Improvement Of Tracking Error Using Motion Constraints

5.8.1 Aim

It is clear from the preceding experiments that the error due to tracking the endoscope is the most significant source of error in our proposed image guided surgical system. Here we investigate how repeating

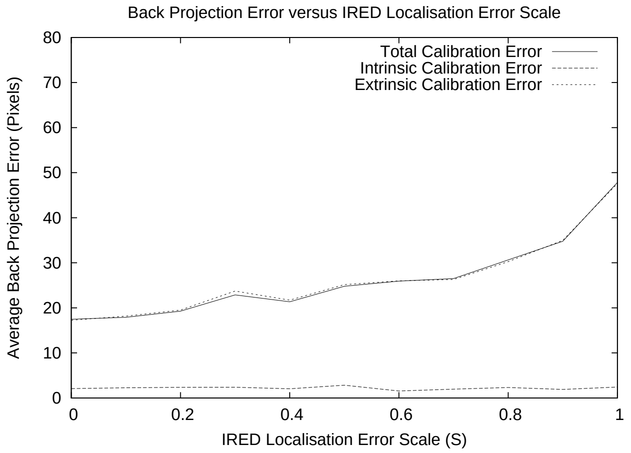

Figure 5.18: The projection error due to calibration versus the IRED tracking error. Reducing the IRED tracking error leads to a significant improvement in calibration accuracy.

the experiment in Section 5.6 building in prior knowledge about the allowable endoscope movement can significantly reduce the tracking error.

5.8.2 Method

The endoscope tracking method used in Section 5.6 registers two IRED sets under the assumption that one set is moving freely in space. In general, however, an endoscope does not move freely in space, it is constrained in its motion. We investigate two modes of constraint, the first is that the endoscope must pass through a fixed pivot point (the trocar) which is relevant to endoscopes in general. The second mode is rotation about a single axis of the *daVinci*^{TM} robot. Figure 5.19 shows the trocar point of the *daVinci*^{TM} robot and an axis of rotation of one of its joints.

Use of a Trocar Constraint

Here it is assumed that the endoscope is constrained to pivot about a known trocar point. This point can be added to the measured IRED positions in the SVD based registration used to track the endoscope. A trocar point is first defined in the coordinate system of the Optotrak. In practice it would be necessary to measure the location of the trocar, for example with a tracked pointer. To allow the endoscope to slide along the endoscope's axis through the trocar an iterative closest point on line algorithm is used. A starting estimate of the endoscope's axis is found using only the measured IRED positions. The point in this axis closest to the trocar point is then found. The registration is then repeated including this point, this gives a new estimate of the axis location, and hence a new point closest to the trocar. This is repeated until the position of the trocar point converges. As there will be an error either in localising the trocar in Optotrak coordinates or in the assumption that the trocar is fixed in space we added a normally distributed error to the trocar position. The value of this error is varied and the on screen tracking error in pixels is plotted against it.

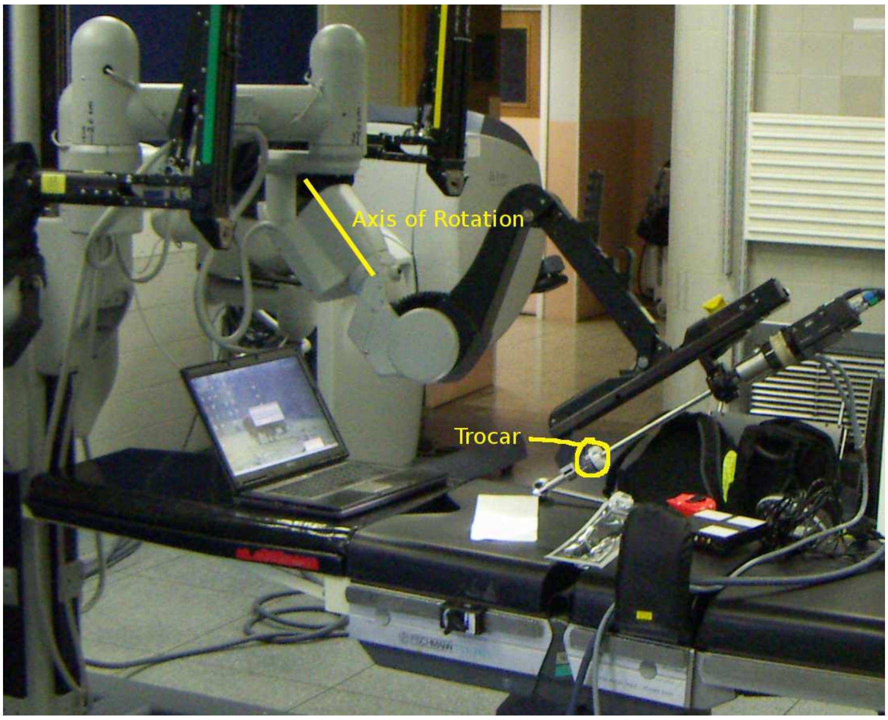

Figure 5.19: Two example motion constraints. In one case the endoscope is constrained to pass through a fixed trocar point (with sliding along the endoscope axis allowed). Although this does not appear to be the case for the *daVinci*^{TM} , it may be useful for other endoscopic systems. The second constraint allows motion around one axis of the *daVinci*^{TM} only. Such a constraint could be implemented using the *daVinci*^{TM} kinematic information.

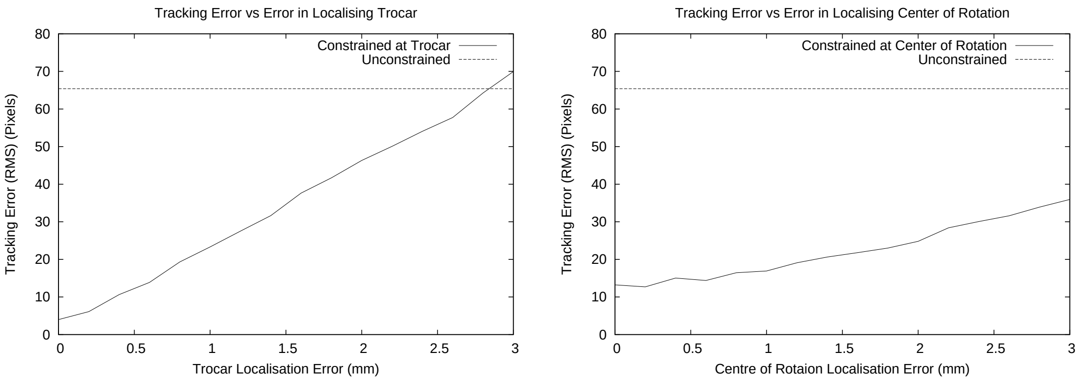

Figure 5.20: Using knowledge of how the endoscope motion is constrained can improve the tracking accuracy. However as the accuracy of this knowledge reduces the improvement in tracking accuracy is reduced.

Use of an Axis Of Rotation Constraint

An alternative to constraining the endoscope to pass through a trocar is to constrain the endoscope based on the known degrees of freedom of the *daVinci*^{TM} system. In practice the *daVinci*^{TM} will be moving through several axes simultaneously, here however we investigate the case in which the endoscope is moving through a single axes. If the centre of rotation of this axes is known (through a prior measurement and interrogation of the *daVinci*^{TM} application programming interface (API)) then this centre of rotation can be added as an additional point in the SVD registration method. Similarly to using a trocar constraint there will be an error in localising the centre of rotation. The tracking error is thus plotted against this localisation error.

5.8.3 Results

Figure 5.20 shows the tracking error (in pixels) versus constraint localisation error for a trocar constraint and a centre of rotation constraint. Applying the trocar constraint yields dramatic improvement in tracking accuracy if the trocar can be accurately localised. The improvement reduces rapidly as the error in localising the trocar reduces. The gains using a center of rotation constraint are less dramatic, but also less sensitive to an increase in the error in localising the centre of rotation.

5.9 Using Normally Distributed IRED Tracking Error

5.9.1 Aim

The calibration back projection error found in Section 5.4 and the tracking error estimated in Section 5.6 are significantly greater than would be expected if the data in Figure 5.3 was correct. The formula used to construct Figure 5.3 assumed isotropic normally distributed ILE. This experiment aims to test the effect of this assumption.

5.9.2 Methods

The calibration simulation from Section 5.5 was repeated twice. Rather than perturbing the IREDs with actual values of e_{ML} taken from a look up table, values of e_{ML} were randomly sampled from normal

	x	y	z
$\langle\lvert E_{ML}\rvert\rangle$	0.053	0.047	0.150
Anisotropic Errors	$\mathcal{N}(0, 0.067)$	$\mathcal{N}(0, 0.059)$	$\mathcal{N}(0, 0.199)$
Isotropic Errors	$\mathcal{N}(0, 0.120)$	$\mathcal{N}(0, 0.120)$	$\mathcal{N}(0, 0.120)$

Table 5.2: The IRED localisation error modelled as isotropic and anisotropic normal distributions $\mathcal{N}(\mu, \sigma)$ with mean μ and standard deviation σ

distributions. In the first case an anisotropic error was used and in the second an isotropic error was used. In Section 5.4 we estimated the expected values of e_{ML} in each direction, giving $\langle\lvert E_{ML}\rvert\rangle_x = 0.053$, $\langle\lvert E_{ML}\rvert\rangle_y = 0.047$ and $\langle\lvert E_{ML}\rvert\rangle_z = 0.150$. If E_{ML} were normally distributed with mean zero, equation 5.15 can be used to find the standard deviation.

$$\sigma_{E_{ML}} = \frac{\langle\lvert E_{ML}\rvert\rangle \times \sqrt{(\pi)}}{\sqrt{2}} \tag{5.15}$$

In the anisotropic case normal distributions with mean zero and standard deviations ($\sigma_{x,y,z}$) calculated using equation 5.15 were used to perturb the IREDs. In the isotropic case equation 5.16 was used to calculate an isotropic standard deviation that gives the same magnitude error when all three directions are summed.

$$\sigma = \sqrt{\frac{\sigma_x^2 + \sigma_y^2 + \sigma_z^2}{3}} \tag{5.16}$$

Table 5.2 defines the normal distributions that were sampled to perturb the IREDs in each case. In both cases the expected value of the total error is the same as when using the look tables from Section 5.4. In each case the calibration simulation was repeated 100 times and histograms of the average calibration back projection errors plotted, similarly to Figure 5.11.

The tracking error experiment from Section 5.6 was also repeated under the influence of anisotropic and isotropic ILE. As well as presenting the on screen error similarly to Figure 5.15, the TRE at the projected point was also calculated for the isotropic case. This should agree with Figure 5.3.

5.9.3 Results

Calibration Simulation

Figure 5.21 shows histograms of the back projection errors for 100 repeat calibrations for both isotropic and anisotropic normally distributed ILE. In both cases the back projection errors are significantly less than for calibration using non normally distributed data, see Figure 5.11. The calibration back projection error achieved in practice (73.63 pixels, see Section 5.4), is very unlikely to have been randomly drawn from either of the two distributions shown in Figure 5.21.

Tracking Error

Figure 5.22 shows the on screen error for the two models of normally distributed ILE. In both cases the on screen error is significantly less than was seen in Figure 5.15. Table 5.3 compares the on screen error

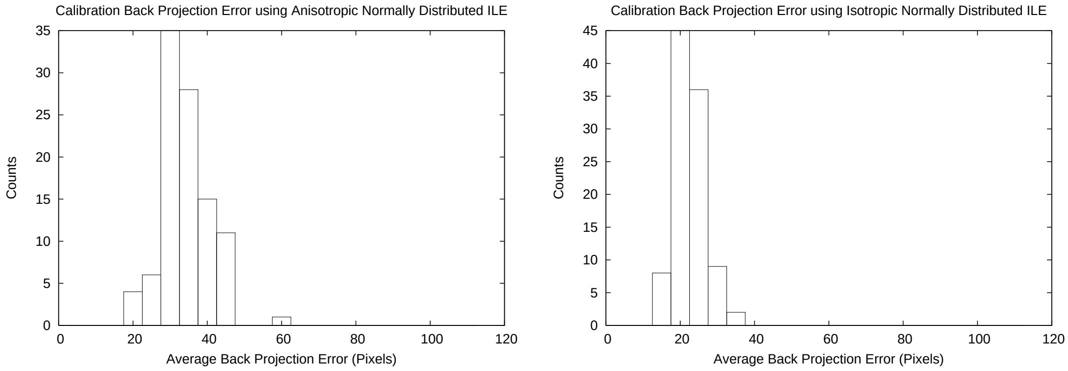

Figure 5.21: Histograms of the average back projection error for 100 simulated calibrations using normally distributed IRED localisation errors. The histogram on the left assumes anisotropic errors, the histogram on the right assumes isotropic errors.

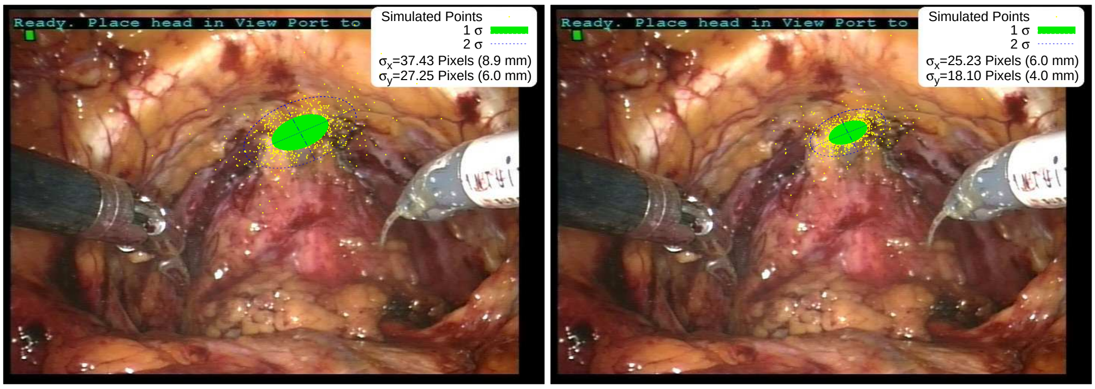
Ready. Place head in View Port to
Simulated Points
1 σ
2 σ
σx=37.43 Pixels (8.9 mm)
σy=27.25 Pixels (6.0 mm)
Ready. Place head in View Port to
Simulated Points
1 σ
2 σ
σx=25.23 Pixels (6.0 mm)
σy=18.10 Pixels (4.0 mm)

Figure 5.22: The on screen error due to tracking the endoscope, assuming normally distributed IRED localisation error. The plot on the left assumes anisotropic error, the plot on the right assumes isotropic tracking error.

	x	y	Magnitude
Anisotropic Non Normal (Section 5.6)	51.82	39.91	65.41
Anisotropic Normal	37.43	27.25	46.30
Isotropic Normal	25.23	18.10	31.06

Table 5.3: The on screen errors for 3 different models of IRED localisation error. The first model allows arbitrary errors based on look up tables generated from actual data. The second row is an isotropic normally distributed error and the third row is an isotropic normally distributed error.

from Section 5.6 with the errors assuming normally distributed error. Modelling the errors as normal distributions significantly underestimates the on screen error.

To check the method and the applicability of equation 5.1 the TRE at a point approximately corresponding to the prostate apex point projected in Figure 5.22 was measured for the isotropic case. This gave a TRE of 2.59 mm. The magnitude of the isotropic normally distributed error used is $\sqrt{3 \times 0.12^2} = 0.21$ mm. This gives an expected value of 0.16 mm, so the ratio of TRE to ILE is approximately 16. As the simulation uses multiple patterns of visible IREDs exact comparison of this number with Figure 5.3 is not possible, however it does fall within the range of expected values for the error ratio. This suggests that 5.1 gives a correct estimate of endoscope tracking error if the assumption of normally distributed isotropic IRED error is correct. However in our case the ILE is neither normally distributed nor isotropic.

5.10 Discussion

At the start of this Chapter, (see Figure 5.3) we presented a method that has been used in the image guided surgery literature [Shahidi et al. (2002)] to estimate the tracking accuracy of an endoscope tracked with optical markers. Using this method, assuming an isotropic, normally distributed ILE with an expected value of 0.2 mm gave an estimated tracking accuracy at a point 200mm from the endoscope tip of approximately 4mm. Using the projection method from Section 1.6.2 this would yield an on-screen error of the order of 18 pixels. However when we tried to calibrate the endoscope ourselves in Section 5.4 we found that the errors we measured were much higher than would be expected if this were the case.

Using a novel method to estimate the actual ILE rather than assuming normally distributed error we demonstrated in Section 5.6 that the actual tracking error is approximately 60 pixels. This is not due to an underestimation of the magnitude of the ILE, rather to an incorrect assumption about its distribution. In Section 5.9 we modelled the ILE with normal distributions, that had the same expected values as the measured ILE. The on-screen error for the normally distributed errors was significantly less than achieved in practice. The disparity between the measurements of on-screen tracking error is due the assumption of isotropic, normally distributed error, which is not appropriate in this case. Optical tracking systems do not provide isotropic error. The presence of anisotropic error has been addressed by Wiles et al. (2008). The problem of non-normally distributed error has not been addressed. The method we have presented allows the correct estimation of tracking error without requiring that the actual ILE fits any particular statistical distribution.

The tracking and calibration errors estimated (approximately 60 and 40 pixels respectively) overwhelm the registration errors determined in the preceding chapters. The calibration error could be improved significantly by improving the calibration process. In light of the result in Section 5.8 it would seem that constraining the endoscope to pass through a trocar during calibration could yield significant improvements. Based on the results in Section 5.7 it would also be worthwhile to investigate the automatic grid point extraction algorithm.

The tracking errors however are more problematic. At present we have no way of reducing the tracking error in theatre. In section 5.8 we introduced and tested via simulation a novel method of constrained tracking that appears to have the potential to reduce the tracking error to acceptable levels. Implementation of these methods is left as future work.

We made the decision at the start of the chapter to track the endoscope using the optical tracker. The more common alternative to this is to use the *daVinci*^{TM} kinematics to track the endoscope. There were two reasons for using the Optotrak to track the endoscope. The first was that as the preoperative data has already been transformed to the Optotrak's coordinate system in the preceding chapter, tracking the endoscope using the Optotrak avoids the need for a registration between the *daVinci*^{TM} kinematic data and the Optotrak. The second reason was that published results for tracking using the *daVinci*^{TM} kinematics were not very encouraging. Here we attempt to test the validity of the second reason by comparing our results with the best result from the literature for tracking using the *daVinci*^{TM} kinematics. The best published result for *daVinci*^{TM} tool tracking is 1.39 mm by Kwartowitz et al. (2009) using a mix of optical tracking and the *daVinci*^{TM} kinematics. This is the error at the tip of end effectors rather than the endoscope. We can compute an equivalent error for the simulated tracking data used in Section 5.6. A target point is defined approximately 360 mm from the tracking collar, which would correspond approximately to the tool tip for a *daVinci*^{TM} end effector. The estimated tracking transforms $T_{O \Rightarrow ET}$ are used to perturb this point and the position compared with that found using the known gold standard tracking transform. This gives a mean TRE at a point near the tool tip of 2.172mm. This is significantly higher than the figure determined by Kwartowitz et al. (2009). However if we assume a linear relationship between TRE and on-screen error, a TRE of 1.39 mm would correspond to an on-screen error of 38 pixels. Still very high in comparison to the other system errors. This suggests that the solution to our endoscope tracking problem does not lie in switching to using the *daVinci*^{TM} kinematics.

5.11 Error Summation

The preceding two chapters and this chapter have isolated and quantified the four main sources of errors in the proposed surgical guidance system. These are summarised in Table 5.4.

The task now is to combine them. It is assumed they are independent. The shape modelling and registration errors can be added to the Monte Carlo simulation by first perturbing the 3D point by a normally distributed random error of the appropriate standard deviation. Assuming that shape modelling and registration errors are independent and normally distributed, then the standard deviation of the combined

Error	Reference	X Err. (Pix.)	Y Err. (Pix.)	Magnitude (Pix.)
Shape Modelling Error	Figure 3.11	13.94	14.99	20.47
Registration Error	Figure 4.18	13.86	13.93	19.65
Endo. Calibration Error (Intrinsic)	Section 5.5	2.03	2.05	2.89
Endo.Calibration Error (Extrinsic)	Section 5.5	25.91	33.38	42.26
Endo. Calibration Error (Total)	Figure 5.12	26.46	33.73	42.87
Endo.Tracking Error	Figure 5.15	51.82	39.91	65.41
Total Error	Figure 5.23	59.73	50.42	78.17

Table 5.4: A summary of the system errors. The total system error is dominated by the endoscope tracking error.

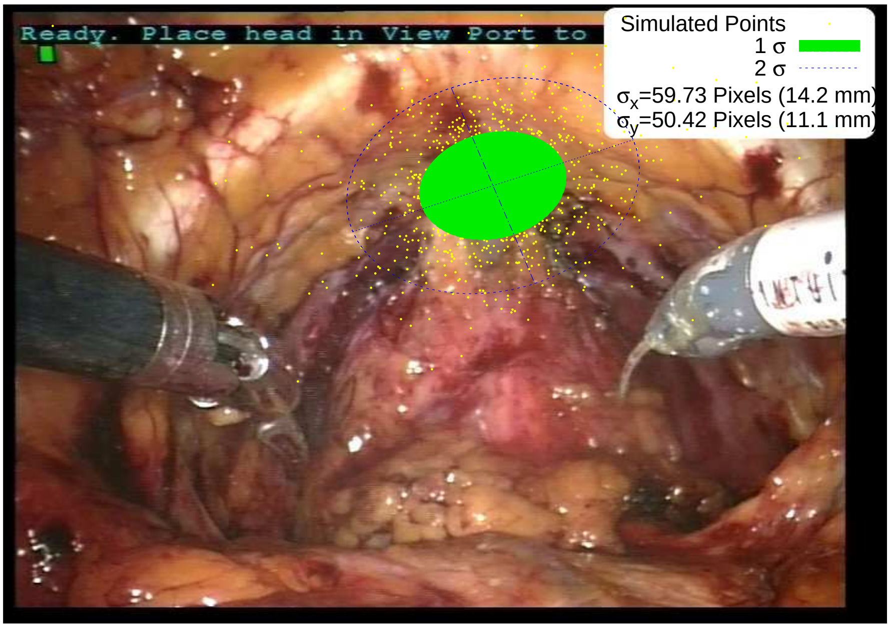
Ready. Place head in View Port to
Simulated Points
1 σ
2 σ
σ_x=59.73 Pixels (14.2 mm)
σ_y=50.42 Pixels (11.1 mm)

Figure 5.23: The total system error as an on-screen pixel error.

error will be $\sqrt{5.3^2 + 5.03^2} = 7.3mm$ in 3D. The resulting point can be projected on the screen under the influence of the tracking error and calibration errors as in Sections 5.5 and 5.6 give the total system error as shown in Figure 5.23. The total system error is dominated by the endoscope tracking error. Improvements in the other component errors will yield minimal improvements without first addressing the endoscope tracking error.

Chapter 6

Patient Trials

6.1 Introduction

The preceding three chapters have described our proposed image guidance system and determined its likely accuracy. At present the system has not been implemented in full for two reasons. The first is that the ultrasound to CT registration algorithm, see Chapter 4, takes too long to be used in-theatre. The second is that the endoscope cannot be tracked accurately enough to be clinically useful, see Chapter 5. Work on both these algorithms is ongoing, with good reason to expect they will result in a practical system. In parallel to this development we have implemented a basic image guidance system. This system aligns the MRI to the endoscope screen directly using the visible pubic bone surface. The system described in this chapter has been developed over the course of five RARP procedures. During development we have also gathered a significant amount of data (MRI, tracked ultrasound, and tracked endoscopic video) for later analysis.

Up to this point this thesis has primarily focused on understanding the accuracy of the image guidance system. By implementing the image guidance system described in this chapter we have been able to assess some of the other factors important to the functioning of the system. These primarily relate to how the guidance system fits into the surgical routine and how the clinical staff interact with the system. Being able to show the surgeon an overlay image during surgery, albeit one of limited accuracy, has enabled two important outcomes. Firstly it was possible to better understand the surgeon's expectations of the system. Secondly the surgeon was better able to understand the capabilities of an image guidance system. To enable the display of overlaid images we have built a minimal user interface which has allowed us to experiment with different ways of presenting the MRI images.

This chapter first details our experience of gathering data in-theatre, summarising our data acquisition methods and the data gathered. Following this is an analysis of the endoscope tracking accuracy during the procedure, using the methods developed in Chapter 5. We found that tracking the endoscope during surgery was more difficult than during the calibration process described in Chapter 5. The tracking accuracy achieved was significantly less than that shown in Figure 5.15. The last part of this chapter describes the direct alignment algorithm we developed for image to image registration and the projection algorithm used to present overlays to the surgeon. The results of this are encouraging.

Patient	Image Description	Size (Vox.)	Resolution (mm)	Slice Spacing (mm)	Acquisition Times (Days Before Surgery)
01	T2 Transverse	320:320:23	0.59:0.59:3.2	3.84	127
02	T2 Transverse	384:384:24	0.65:0.65:3	3.9	21
03	T2 Transverse	320:320:23	0.59:0.59:3.2	3.84	106
04	T2 Transverse	256:256:23	0.63:0.63:3	3.3	177
05	T2 Transverse	512:512:30	0.35:0.35:3.5	3.84	101

Table 6.1: Details of the T2 weighted MRI slices that were used for overlay images in theatre.

6.2 MRI Data

As currently implemented the system does not require additional MRI acquisitions. We have used existing MRI images for overlay. Table 6.1 summarises the MRI data used for overlay for each patient. Full details of the data gathered for each patient can be found in Appendix A. For in-theatre overlay we used the transverse T2 weighted images as these show the prostate in greatest detail.

All MRI acquisitions were supine scans and no attempt was made to control the patient's bladder or rectal filling. Furthermore there was a significant time gap between image acquisition and surgery. Therefore we would expect significant motion of the prostate between the image and the surgical position. For ongoing work it would be useful to acquire the MRI at a time closer to the surgery, with the bladder and colon empty. For registration using the pelvic bone it would also be useful to acquire a T1 weighted image of the entire pelvic bone. This is not however necessary for the early stage evaluations presented here.

6.3 Ultrasound Data Capture

Ultrasound data was collected for each of the 5 patients. Ultrasound images were captured using a frame grabber attached to the video output of the ultrasound machine. Ultrasound tracking data was written to a separate tracking file that recorded the IRED positions for each slice. As discussed in Chapter 4 two pieces of software were used for collecting the tracking data in theatre. The first version of the tracking code was used for Patients 1 and 2, while Patients 3 to 5 used the newer code. The earlier code did not record individual IRED positions, only the estimated tracking matrix $T_{UT \Rightarrow O}$ along with the RMS IRED registration error. The later code recorded the individual IRED positions allowing a more accurate estimation of the tracking error.

A bug in the earlier code caused it to crash after 100 slices had been collected. Therefore for Patients 1 and 2 we attempted to collect fewer ultrasound images but ensured that all images were of good quality. The ultrasound probe was first positioned carefully to get a good image of the patient's pelvic bone. It was then held static and a trigger pressed to capture a single image. This was a time consuming (approximately 15 seconds per image) procedure requiring two operators, one with the probe and another at the workstation. The new code did not suffer from this limitation, allowing images to be obtained continuously. For Patients 3 to 5 therefore we collected many more images, but many do not

Patient	(Old/New) Code	Number Of Slices
01	Old	43
02	Old	38
03	New	4338
04	New	10185
05	New	2064

Table 6.2: The number of Ultrasound slices collected for each of the five patients.

show the pelvic bone clearly, so sorting will be required before they are used for registration. For Patients 3 to 5 the following acquisition method was used. The tracking software was started, then the ultrasound probe was moved slowly over the patient's pelvis. At positions where the pelvic bone could be seen clearly the probe was momentarily held static. This procedure was quicker (taking about 5 minutes in total) than using the old code and could be completed by a single operator controlling the ultrasound probe.

The result of this is that for patients 1 and 2 less than 50 ultrasound images were collected, (much fewer than were found necessary for registration in Chapter 4). However the individual images are all usable. For patients 3 to 5 many more images were collected (over 2000), however the majority of these do not show the pelvic bone well. In order to use these for registration a sorting algorithm will be necessary. Images with poor tracking accuracy (calculated using the measured IRED positions) should first be discarded. It may also be necessary to discard images that do not show the pelvic bone well using some sort of image processing. However these images should in theory be removed when the images are converted to bone edge probability images. The resulting number of ultrasound images collected for each patient are shown in table 6.2.

6.3.1 Patient Positioning During Imaging

The original intent was to capture the ultrasound images of the patient immediately prior to surgery with the patient in the operative position, so that no patient motion occurred between imaging and surgery. In practice however this was not possible. The patient could not be imaged after the cutting of the surgical ports due to the presence of sterile drapes. The ultrasound images were therefore collected prior to cutting of the ports. The port cutting was done with the patient horizontal, the patient is then tilted to enable docking with the *daVinci*^{TM} . For these reasons we adopted the following procedure when gathering ultrasound data.

The patient was positioned on the operating table and their legs elevated to the operative position. The table was then tilted (1 degree of freedom motion), the angle of the table recorded, and the ultrasound images collected. Figure 6.1 shows the patient in the head down position with the probe positioned near the pelvic bone.The table was returned to the horizontal and the port cutting and further preparation completed. The table was then tilted to the same angle as measured previously and the *daVinci*^{TM} docked. This introduces the real possibility that the patient will move in between ultrasound imaging

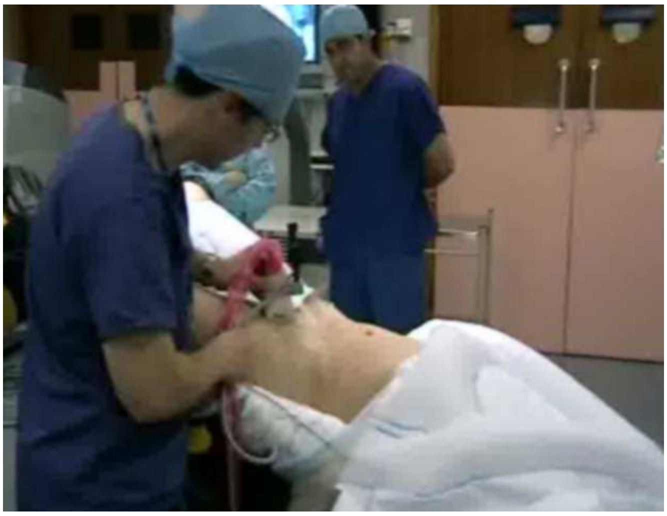

Figure 6.1: The patient is positioned head down in the operative position and ultrasound images of the pelvic bone collected using a tracked ultrasound probe.

and surgery. Further work is required to assess the effect of this and determine a solution. One possibility would be to place skin markers on the patient during surgery. We have not attempted this yet, however, for reasons that will be discussed in Section 6.3.2.

6.3.2 Processing of Ultrasound Data

The original intention for the in-theatre work was to use the ultrasound to bone registration algorithm (Chapters 3 and 4) to find $T_{M \Rightarrow O}$. The accuracy of this would then be assessed using the visible position of the pubic arch on the endoscopic images and $T_{O \Rightarrow ES}$ estimated as per Chapter 5. However, contrary to our expectations, it became apparent that the error in estimating $T_{O \Rightarrow ES}$ would be significantly larger than the error in $T_{M \Rightarrow O}$, see Chapter 5 for details. Therefore the approach of validating the registration method using the visible anatomy cannot be used. At present we therefore have no way of testing the accuracy of our estimate of $T_{M \Rightarrow O}$. With limited time available it was decided to put the ultrasound data aside for possible future analysis and concentrate instead on implementing a direct endoscope image to MRI registration algorithm, to be described in Section 6.6.

6.4 Acquisition of Video Data

The same software as was used for collecting the ultrasound images (for Patients 3 to 5) was used for collecting the endoscopic video data. Switching from ultrasound images to endoscopic images was done by connecting the frame grabber to the video output of one of the *daVinci*^{TM} assistant consoles. Video data was captured at approximately 7 frames per second. Each frame is matched to tracking data taken from the Optotrak. The video capture software was run for the entirety of the procedure. During capture of data for Patients 1, 3, and 4 a significant number of frames were lost due to overheating of the frame grabber. These problems did not occur for Patient 2, and a more reliable frame grabber was used for Patient 5, preventing re-occurrence. Table 6.3 lists the number of video frames captured for each

Patient	Number Of Endoscope Frames N	Elapsed Time T (h:mm)	Proportion Lost Data P_{lost}
01	33637	1:53	0.31
02	74315	2:40	0.00
03	61206	2:31	0.06
04	39764	1:50	0.16
05	78631	3:01	0.00

Table 6.3: The number of endoscopic video frames collected for each of the five patients.

procedure. The average frame rate (fr in frames per second) was calculated using the data for Patients 2 and 5, where the frame grabber functioned continuously over the elapsed time. In both cases fr was 7.2 frames per second. The proportion of frames lost due to the overheating frame grabber for the remaining patients was calculated using equation 6.1. Where N is the number of frames captured and T is the time elapsed from the first frame to the last frame. Values of P_{lost} for each patient are shown in Table 6.3.

$$P_{lost} = 1 - \frac{N}{T \times fr} \tag{6.1}$$

Similarly to Figure 5.6 we checked how the tracked positions of the endoscope collar during surgery were distributed in relation to the characterised volume of the Optotrak. Figure 6.2 shows the results. All tracked positions fall within the characterised volume of the Optotrak.

6.5 Estimation of Tracking Error for a Moving Endoscope

6.5.1 Aim

In Section 5.3 we presented a method for estimating the IRED localisation error for the tracked endoscope. This was based on static data captured in an uncluttered environment. A further complication is the presence of sterile drapes between the IREDs and the tracking camera. Figure 6.3 demonstrates some of these challenges. Many of these challenges could be overcome relatively easily. The surgical assistants could move to the side during image overlay and an effort could be made to better control the cable routing. Despite these challenges enough tracking data was collected to repeat the analyses from Section 5.6. Two factors, in addition to those present in Section 5.3, may lead to an increased tracking error for the surgical data. The first of these are motion of the endoscope, which could be avoided by only using the overlay when the endoscope is stationary. The second is the presence of optically transparent sterile surgical drapes. Here we attempt to quantify the effect of these additional factors on the IRED tracking error.

6.5.2 Data

We only used the tracking data from Patient 02 for the in-theatre data. Patient 02 was used as it provided a full record of the in-theatre tracking. Initial analysis of the data from the remaining patients indicates

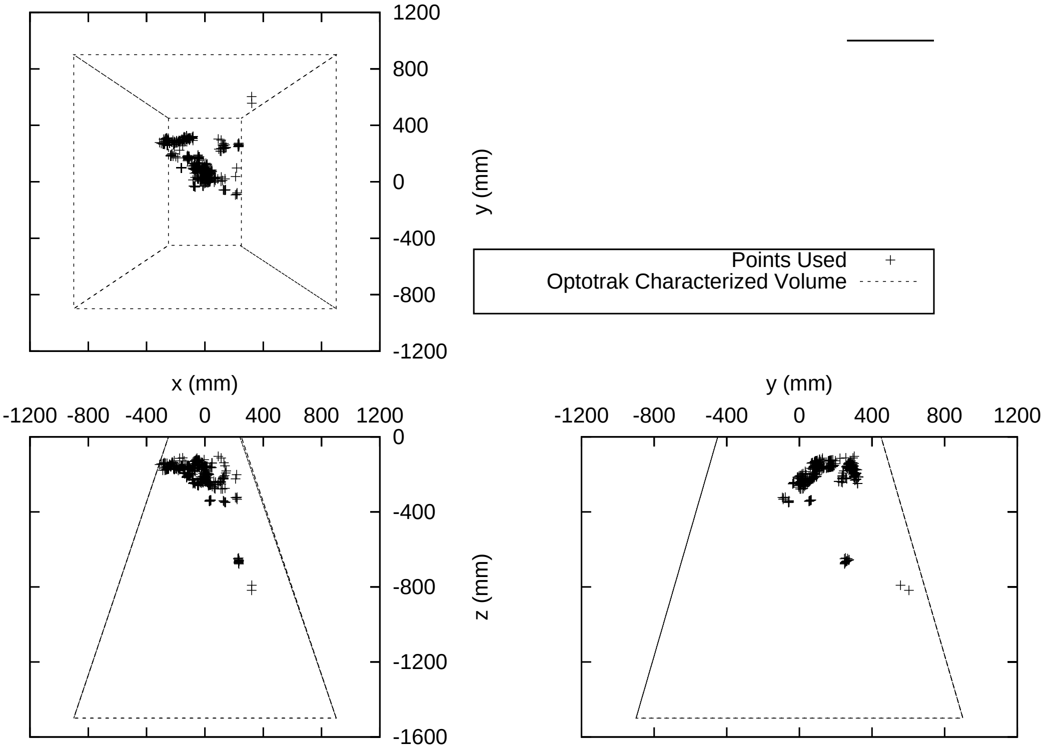

Figure 6.2: The positions of the tracking collar recorded during data acquisition on Patient 02. All tracked locations fall within the characterised volume of the Optotrak

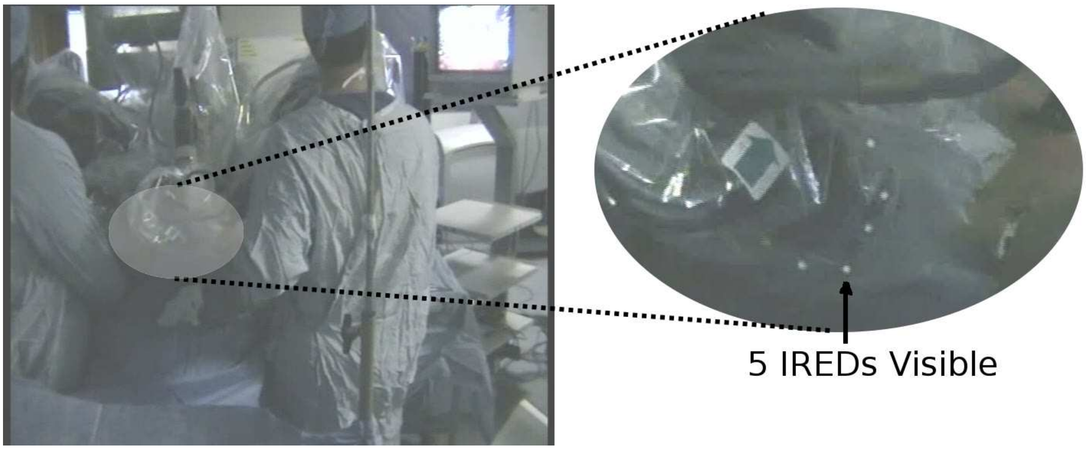

Figure 6.3: The view of the IREDs from the Optotrak camera system. The presence of the two surgical assistants, the anaesthetist, surgical drapes and numerous control and data cables makes endoscope tracking very difficult.

Direction	Patient 02	Static (Section 5.4)
X	0.204	0.053
Y	0.135	0.047
Z	0.445	0.150

Table 6.4: The expected values of the IRED localisation error based on the data collected from Patient 02. The static tracking values from Section 5.4 are also shown for comparison.

that the data from Patient 02 is representative of the data from the other patients. There is no reason to expect that the Patient 02 data is not representative of the other data sets. A full analysis of the remaining data has not yet been completed. In addition to this we recorded a separate sequence of tracking data with the endoscope moving but without the presence of surgical drapes. This was to test the effect of the surgical drapes on the result.

6.5.3 Method

The method was the same as employed in Section 5.3, summarised again here. The measured positions of a given IRED was compared with the IRED position estimated from the measured positions of the remaining IREDs on the tracking collar. An adjustment was made to account for the registration error. This gives a distribution of estimated IRED localisation errors that can be used as a look-up table for Monte-Carlo simulation of the tracking accuracy. To enable direct comparison of the tracking results with those in Section 5.6 we used the same IRED visibility states for this simulation.

6.5.4 Results

Estimated IRED Localisation Errors

Table 6.4 presents the expected values of the estimated IRED localisation errors. The results for the static tracking performed in Section 5.4 are also included for reference. It is clear that the errors seen in theatre are significantly (between 3 and 4 times) greater than than those observed during the endoscope calibration in Section 5.4. Figure 6.4 shows the distribution of the estimated IRED localisation errors.

6.5.5 Estimated Tracking Error

Using the same method as described in Section 5.6, the error distributions shown in Figure 6.4 were used to perturb the IREDs during a simulation of tracking error. Figure 6.5 shows the results as an on screen tracking error. The magnitude of the error standard deviations shown in Figure 6.5 is 164 pixels. By comparison, in Section 5.6 (for static endoscope tracking) the magnitude of the error standard deviations was found to be 65 pixels. Two factors could explain this increase. The first is the effect of endoscope motion. The second is the use of surgical drapes. To test the effect of endoscope motion without the presence of surgical drapes a further data set was collected of the tracking collar moving but without the presence of surgical drapes. The tracking error for this data was nearly identical to that shown in Figure 6.5. This suggests that the increase in IRED tracking error observed here is due exclusively to motion of the endoscope. Further work is required to determine the precise effects of the moving endoscope

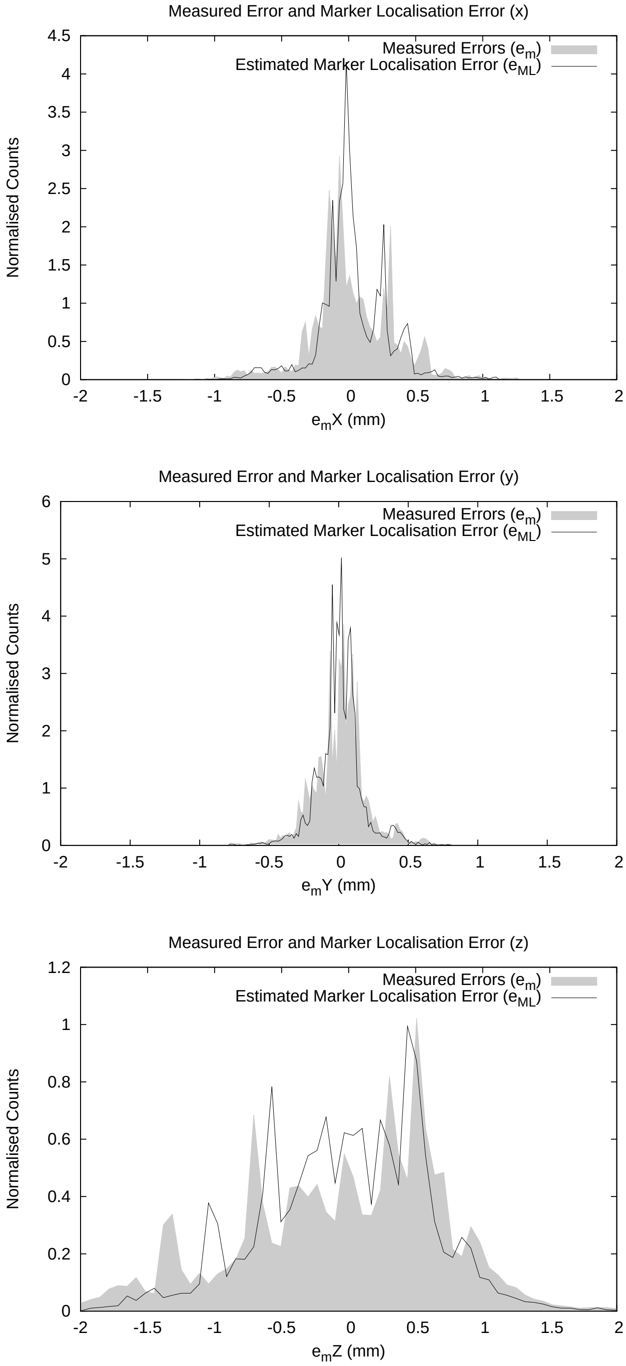

Figure 6.4: The measured marker errors (e_m) and the estimated marker localisation errors (e_{ml}) in each direction for the data collected during surgery on Patient 02. Equation 5.12 is used to estimate e_{ml} using e_m. Equation 5.12 slightly reduces the measured errors to account for the error in the registration process used to estimate the actual marker positions.

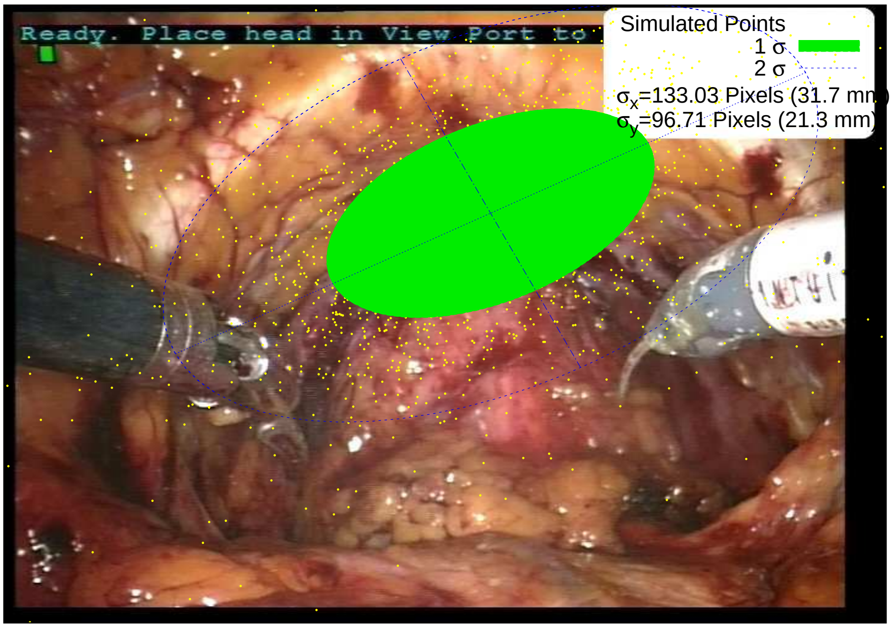
Ready. Place head in View Port to
Simulated Points
1 σ
2 σ
σx=133.03 Pixels (31.7 mm)
σy=96.71 Pixels (21.3 mm)

Figure 6.5: The on screen error due to tracking the endoscope during surgery on Patient 02. The errors are significantly larger than those seen for static endoscope tracking in Figure 5.15.

and the surgical drapes. The increase in tracking error for a moving endoscope agrees with the results of Barnes et al. (2007) who observed significant loss of accuracy when the Optotrak was used to track a moving target.

6.5.6 Discussion

In attempting to test whether the IRED localisation errors found in Section 5.4 are applicable in-theatre we have highlighted an important short coming of the Optotrak Certus system. The Optotrak system is significantly less accurate when tracking a dynamic target. [1] Despite this it may still prove practical track the moving endoscope by applying the methods proposed in Section 5.8. Work in this area is ongoing.

6.6 Direct Alignment of the MRI Image to the Endoscope Image

Due to the inability to validate the ultrasound to MRI registration in theatre we elected to concentrate solely on developing an algorithm to align the endoscope video image directly to the MRI data.. The need for this algorithm was realised after analysis of the results for Patients 01 and 02. Development commenced and deployment was attempted on Patients 03 to 05. We continued to collect ultrasound data for later analysis, though no attempt has yet been made to process this data.

Alignment is based on the fact that the inner surface of the bony pubic arch can be seen in both the endoscopic images and the MRI images. The first stage is to manually pick points on the inner surface of the MRI. Dedicated software was implemented in C++ using OpenGL. A screen shot of this is shown in Figure 6.6. Points are picked in transverse planes of the MRI, with 7 points picked over 6 slices, plus

---

[1] The apparent dynamic tracking error will be further increased by any synchronisation error between the video and tracking signals, discussed in Appendix C.

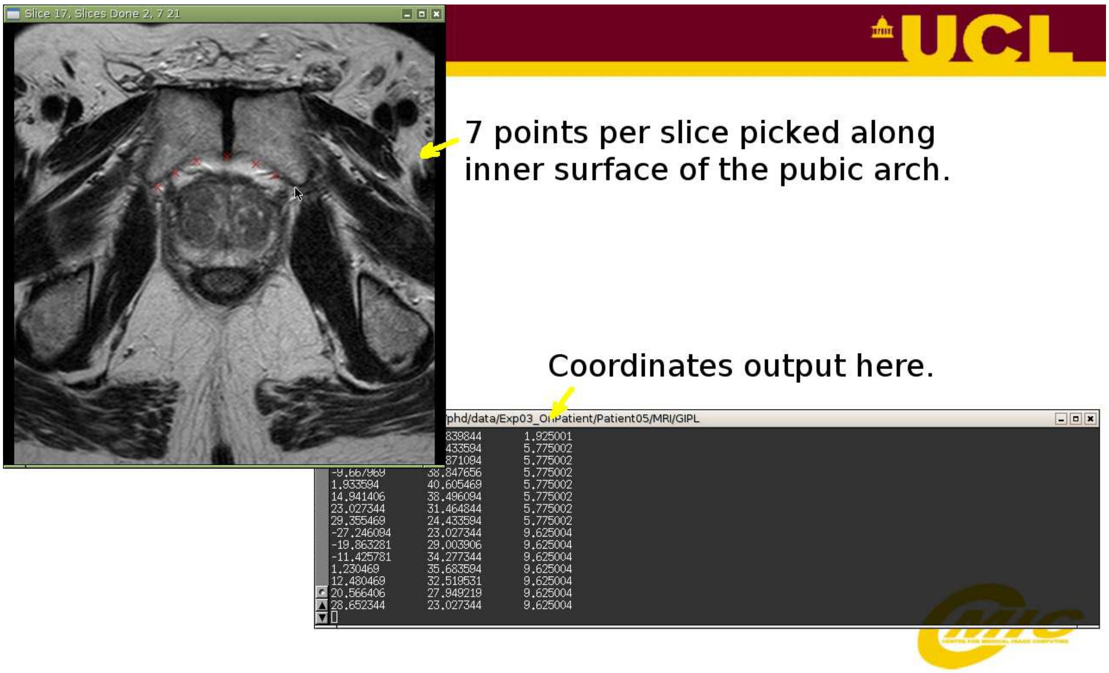

7 points per slice picked along inner surface of the pubic arch.

Coordinates output here.

```
/phd/data/Exp03_OnPatient/Patient05/MRI/GIPL
          839844      1.925001
          433594      5.775002
          871094      5.775002
-9.667969   38.847656   5.775002
1.933594    40.605469   5.775002
14.941406   38.496094   5.775002
23.027344   31.464844   5.775002
29.355469   24.433594   5.775002
-27.246094  23.027344   9.625004
-19.863281  29.003906   9.625004
-11.425781  34.277344   9.625004
1.230469    35.683594   9.625004
12.480469   32.519531   9.625004
20.566406   27.949219   9.625004
28.652344   23.027344   9.625004
```

Figure 6.6: 43 points on the inner surface of the pubic arch are manually picked from the preoperative MRI image. Using purpose built software this task can be completed in less than 5 minutes.

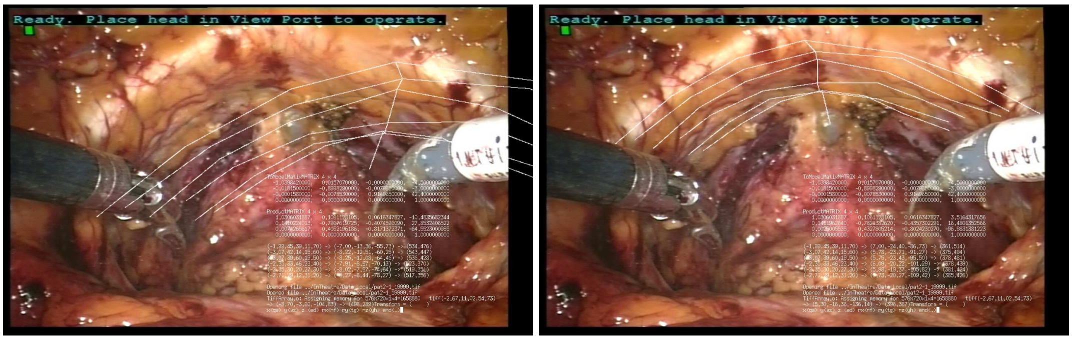
Ready. Place head in View Port to operate.
Ready. Place head in View Port to operate.

Figure 6.7: Manually selected points along the inner surface on of the pubic arch are projected onto the endoscopic image. The keyboard is used to increment the transform $T_{M \Rightarrow O}$ in each the six degrees of freedom. The image on the left shows the two images prior to alignment. The image on the right shows the two images after alignment.

an estimate of the prostate apex. Using our dedicated software this can be done in less than 5 minutes, at any point in time between MRI acquisition and surgery.

The points picked are used to form a wire-frame model of the inner surface of the pubic arch. This wire-frame model is projected onto the endoscope screen using initial estimates of the various transforms that make up $T_{M \Rightarrow ES}$. At this stage the distortion model of the endoscope is not used as it significantly slows the projection algorithm, especially if some points are off screen. This projection is shown in Figure 6.7. The estimate of $T_{M \Rightarrow ES}$ is interactively updated by varying the estimate $T_{O \Rightarrow ET}$.

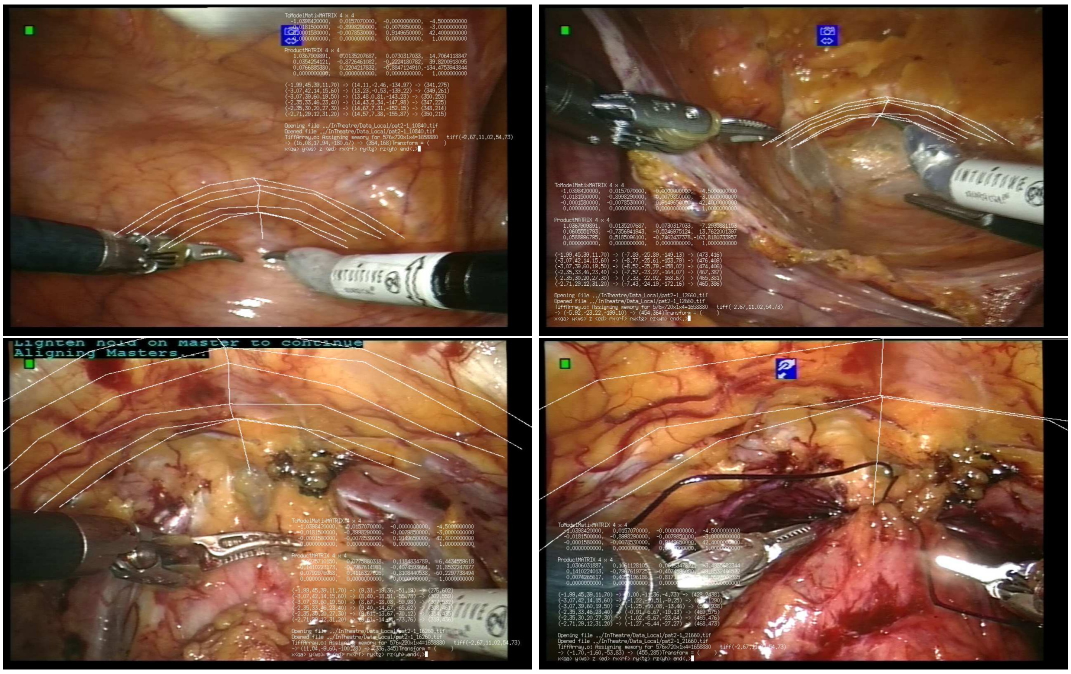

Figure 6.8: The direct alignment routine can be applied at any point in the surgery where the pubic arch is visible to some extent. The error in the alignment will depend on the extent to which the pubic arch can be discerned. We have attempted in each case here to perform the registration using our estimate of the position of the pubic arch. We have no method at present to validate the accuracy of the results.

$T_{O \Rightarrow ET}$ is incremented in each of its 6 degrees of freedom (3 translations and 3 rotations) using keyboard commands. Two increment sizes are implemented, $\pm 1(^\circ)(mm)$ and $\pm 10(^\circ)(mm)$, allowing course to fine registration. A skilled user can perform this registration from an arbitrary starting position in less than 5 minutes. With a good initial estimate of $T_{M \Rightarrow ES}$ registration can be performed in under 30 seconds. We have not implemented an intuitive user interface, so at present the alignment algorithm is not suitable for use by clinicians.

This rigid alignment algorithm was successfully used for Patients 4 and 5 to provide the surgeon with an image showing the patient's MRI overlaid on the surgical scene, during surgery. Due to the poor endoscope tracking accuracy the alignment routine must be repeated at each static location of the endoscope used for overlay. In figure 6.7 we performed the alignment on an image frame where the inner surface of the pubic arch is clearly visible. If necessary the alignment can be performed on frames where the pubic arch is not so clearly visible. This will lead to a less accurate registration, however when used for the overlay of MRI the effect still appears visually accurate.

6.7 Display of MRI Overlaid on Endoscope Screen

With $T_{M \Rightarrow ES}$ estimated we can project MRI data onto then endoscope screen. We developed software in C++ using OpenGL to perform this projection. This software was successfully used for Patient 05 to display the overlay in theatre. The overlay in this case was static with the back ground image being a static frame captured from the endoscope. Since then the software has been developed to allow the endoscope image and the estimate of $T_{M\Rightarrow ES}$ to be updated at up to 8 frames per second, ie. it can keep pace with the images being captured by the frame grabber. Updating of $T_{M\Rightarrow ES}$ is done by updating $T_{O\Rightarrow ET}$ based on the tracking data from the Optotrak. If $T_{n,\mathrm{manual}}$ is the estimate of $T_{O\Rightarrow ET}$ for frame n as estimated manually using the method detailed in Section 6.6, and $T_{n,\mathrm{Optotrak}}$ is the corresponding estimate of $T_{O\Rightarrow ET}$ using the Optotrak then $T_{O\Rightarrow ET}[i]$ can be estimated for an arbitrary frame i by equation 6.2.

$$T_{O\Rightarrow ET}[i] = T_{i,\mathrm{Optotrak}} \times T^{-1}_{n,\mathrm{Optotrak}} \times T_{n,\mathrm{Manual}} \tag{6.2}$$

The revised estimate of $T_{O\Rightarrow ET}[i]$ can be used to calculate a new estimate of $T_{M\Rightarrow ES}$ and this is used to project the MRI data. If an estimate of $T_{i,\mathrm{Optotrak}}$ is unavailable (due to insufficient IREDs being visible, then the last known version of T_{Optotrak} is used instead. When this occurs an on-screen marker changes from green to red to alert the surgeon that the endoscope is not being tracked. We have not yet used this live update in theatre. Shortly we will present some example images that provide a qualitative assessment of the tracking accuracy. Prior to this we will finish description of the algorithm.

6.7.1 Projection Model

The MRI image is first loaded into memory. The image is then cropped to show a particular region of interest, for example the prostate. Each remaining voxel in the MRI image is first transformed to Optotrak coordinates using $T_{M\Rightarrow O}$. Projection to on screen pixel coordinates is then performed using $T_{O\Rightarrow ES}[i]$.

MRI Interpolation

Unless the MRI image is very far from the screen this will result in large spaces between the projected pixels, as shown at the left in Figure 6.9. To prevent this the MRI image is first interpolated. The amount of interpolation can be controlled interactively, depending on the distance if the MRI image from the endoscope lens. More interpolation is necessary to get a realistic image when the MRI is close to the lens, however there is a direct trade off in projection speed, so we aim to keep interpolation to the minimum required. To the right of Figure 6.9 we show the effect of dividing each voxel in the MRI into 64 (4 times interpolation).

The quickest interpolation method is nearest neighbour interpolation. The result of this is shown at the left of Figure 6.10. The resulting image is somewhat blocky. This can be reduced by using tri-linear interpolation. The overlay on the right of Figure 6.10 shows the same overlay using tri-linear interpolation. The image is effectively blurred, giving a more attractive overlay but with a slight trade off in projection speed.

Overlay Opacity

The interface allows the opacity of the overlay image to be incremented in steps of 10% from 0 to 100%. Figure 6.11 shows an example of changing the slice opacity.

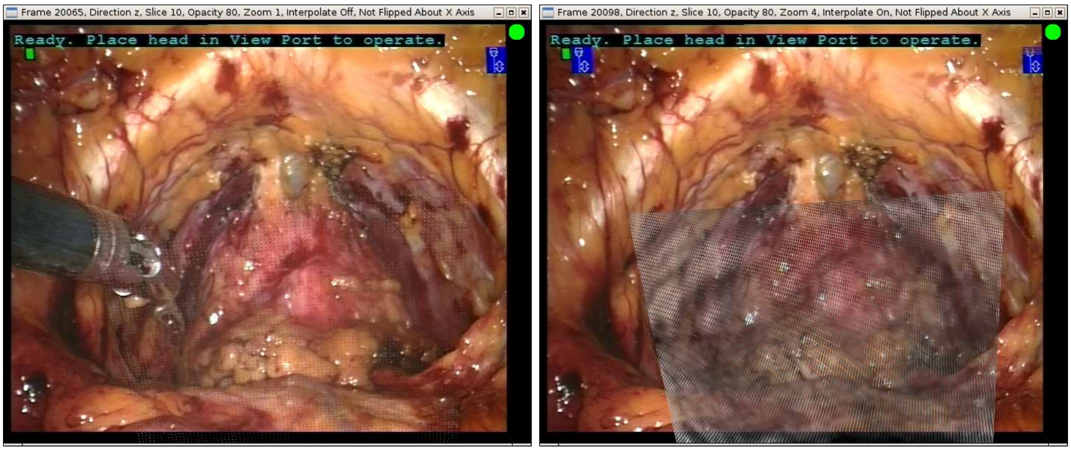

Frame 20065, Direction z, Slice 10, Opacity 80, Zoom 1, Interpolate Off, Not Flipped About X Axis

Ready. Place head in View Port to operate.

Frame 20098, Direction z, Slice 10, Opacity 80, Zoom 4, Interpolate On, Not Flipped About X Axis

Ready. Place head in View Port to operate.

Figure 6.9: In order to show the MRI on screen as a coherent set of points it in necessary to interpolate the MRI image. On the left the MRI is projected without interpolation resulting in a sparse set of pixels. On the right the MRI image has been interpolated, placing 3 additional data points between each voxel vertex. The projected image is starting to fill in. The interesting fill pattern is an artefact of the distortion model used.

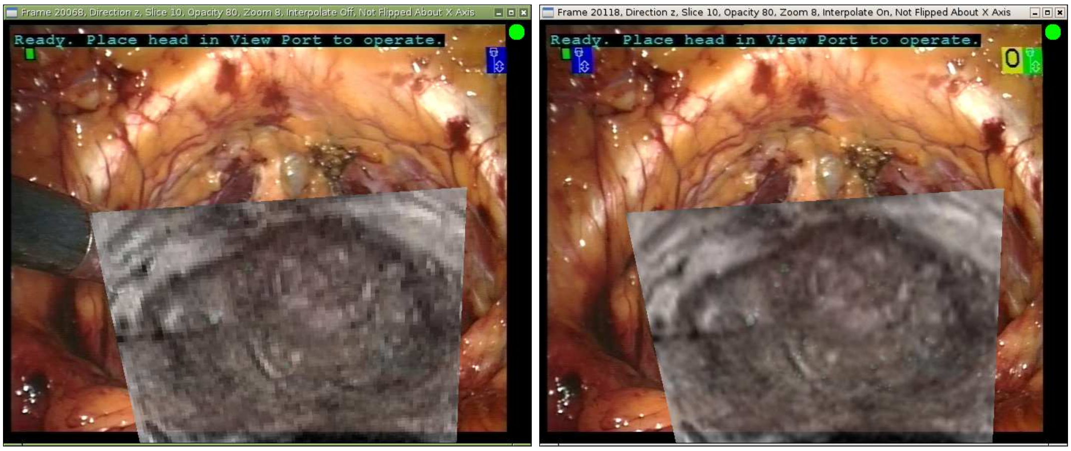

Frame 20068, Direction z, Slice 10, Opacity 80, Zoom 8, Interpolate Off, Not Flipped About X Axis

Ready. Place head in View Port to operate.

Frame 20118, Direction z, Slice 10, Opacity 80, Zoom 8, Interpolate On, Not Flipped About X Axis

Ready. Place head in View Port to operate.

Figure 6.10: The effect of using nearest neighbour interpolation (on the left) or tri-linear interpolation (on the right).

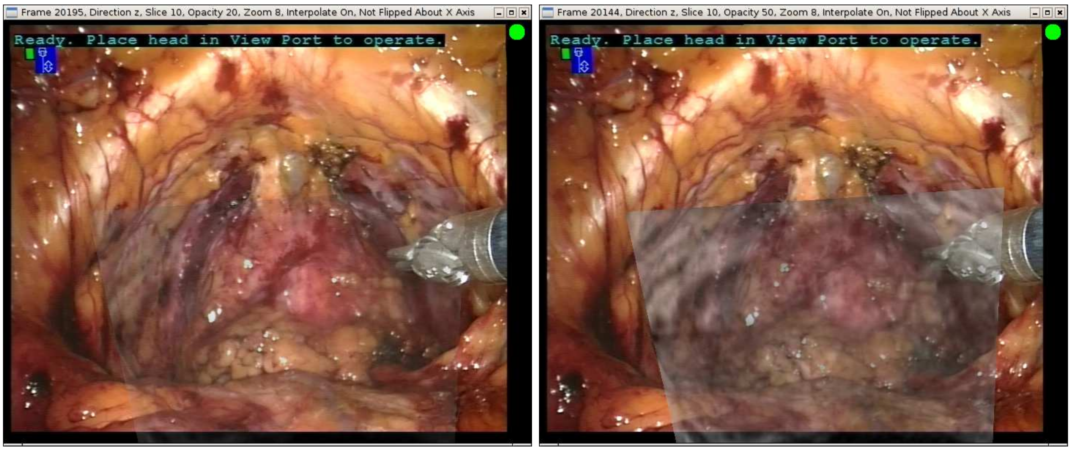

Frame 20195, Direction z, Slice 10, Opacity 20, Zoom 8, Interpolate On, Not Flipped About X Axis

Ready. Place head in View Port to operate.

Frame 20144, Direction z, Slice 10, Opacity 50, Zoom 8, Interpolate On, Not Flipped About X Axis

Ready. Place head in View Port to operate.

Figure 6.11: The opacity of the overlay can be varied interactively between 0 and 100%. Here the same overlay is shown at 20% opacity (at left) and 50% opacity (at right).

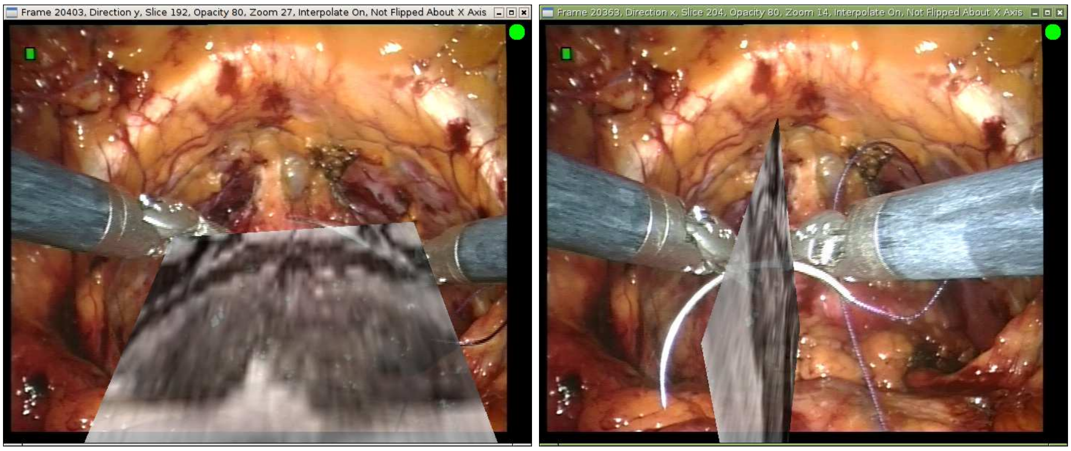

Frame 20403, Direction y, Slice 192, Opacity 80, Zoom 27, Interpolate On, Not Flipped About X Axis

Frame 20363, Direction x, Slice 204, Opacity 80, Zoom 14, Interpolate On, Not Flipped About X Axis

Figure 6.12: It is also possible to overlay Coronal (at left) or Sagittal (at right) slices. In this case these are re-sampled from a stack of transverse MRI slices.

Slice Orientation

The orientation of the projection slice can also be changed interactively. We have generally used transverse slices as these provide the most intuitive overlay, however it is possible to also show sagittal and coronal slices, see Figure 6.12.

Moving Through the Slice Stack

Difference slices in the MRI volume can be selected interactively. Figure 6.13 shows projection of two different transverse slices from the same MRI volume.

6.7.2 Motion Tracking

Figure 6.14 shows the movement of the overlay image as the endoscope undergoes motion. The direct alignment method described in Section 6.6 was first used to align the MRI with the endoscopic data with the endoscope static in the pose shown in image 1 of Figure 6.14. The endoscope remained static in this

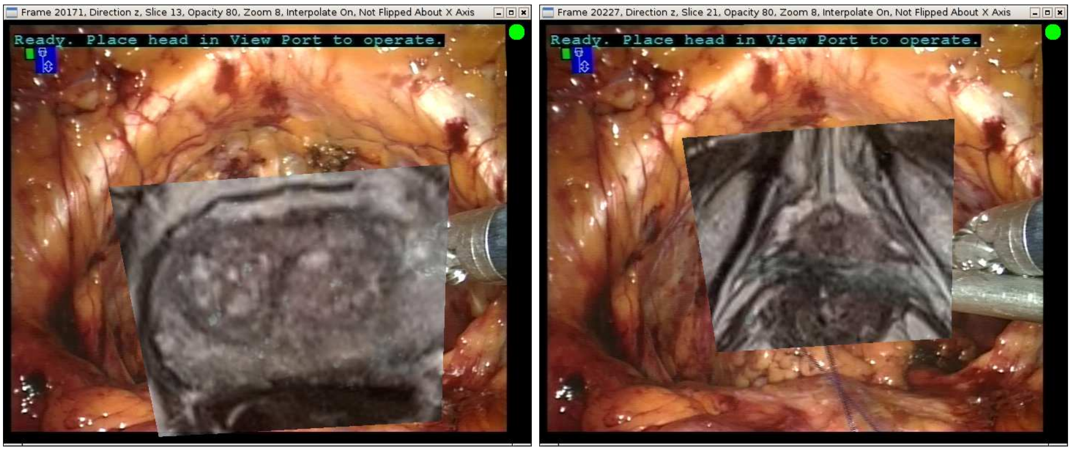

Frame 20171, Direction z, Slice 13, Opacity 80, Zoom 8, Interpolate On, Not Flipped About X Axis

Ready. Place head in View Port to operate.

Frame 20227, Direction z, Slice 21, Opacity 80, Zoom 8, Interpolate On, Not Flipped About X Axis

Ready. Place head in View Port to operate.

Figure 6.13: Different slices from the MRI can be selected interactively. The image at right shows a slice 8 slices further away (inferior) to the image at the left

position for a prolonged period (about a minute) during which time the tracking data remained reasonably steady, with the overlay moving less than 5 pixels over the image. This represents random jitter in the IRED position measurement. In image 2 the overlay image begins to move before the endoscope video image, due to the frame grabber lag (see Appendix C). During motion, images 2 to 4, the position of the overlay image moves about very unpredictably, in line with the tracking error shown in Figure 6.5. Once endoscope motion has stopped, images 5 and 6, the overlay position settles into a new steady position, which is slightly wrong (based on the visible position of the prostate apex) compared to the position in image 1. This is in line with the static tracking error seen in Figure 5.15.

6.8 Clinical Experience

The aim of this chapter is not to validate the work of the preceding chapters in theatre. This cannot be done as no way of validating the results against a gold standard value is available. The aim of this chapter was to get some form of image guidance up and running in theatre, so that the clinicians could interact with it. To that end we have had some success. We present here the clinical outcomes of the 5 patients involved in the trial. The system as implemented does not involve any invasive procedure in addition to the prostatectomy, for example the use of fiducial markers. We have also taken care to ensure that the surgeon does not use the overlaid images to make clinical decisions yet. Therefore we would not expect any negative outcomes. However the implementation to date has also not been sufficient to show any positive outcome.

The reports of the clinicians involved were positive. The impact of being able to refer to the preoperative MRI in an intuitive way during surgery was seen as positive. No formal evaluation, for example using a questionnaire, of surgeons' experience of the system has been done yet. Significantly more work is required to demonstrate any real positive impact of the system.

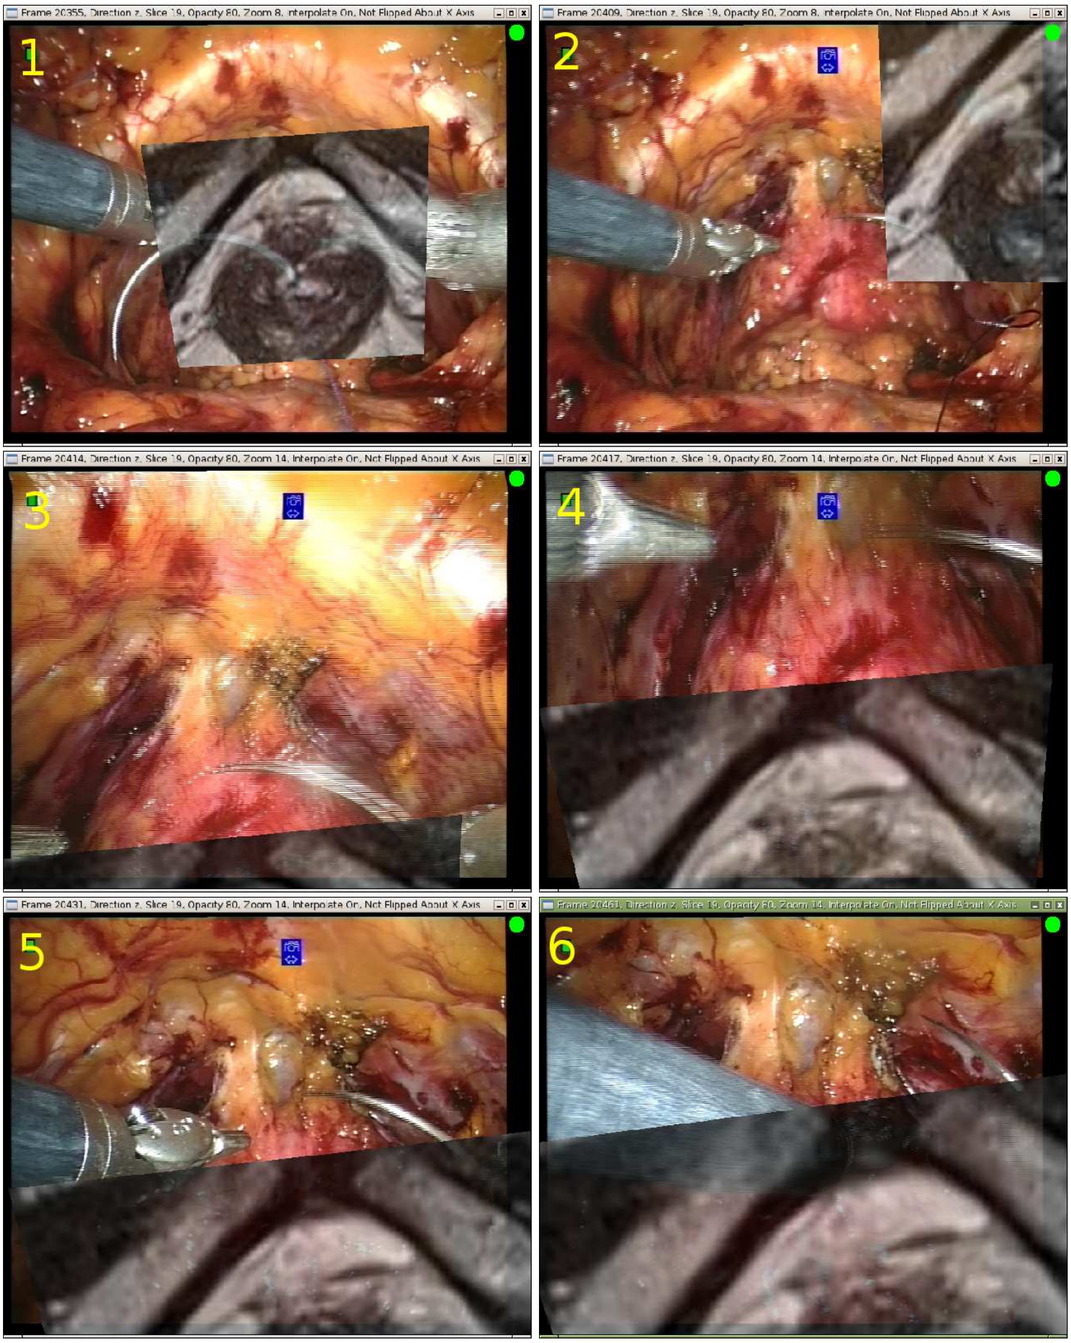
Frame 20355, Direction z, Slice 19, Opacity 80, Zoom 8, Interpolate On, Not Flipped About X Axis
1
Frame 20409, Direction z, Slice 19, Opacity 80, Zoom 8, Interpolate On, Not Flipped About X Axis
2
Frame 20414, Direction z, Slice 19, Opacity 80, Zoom 14, Interpolate On, Not Flipped About X Axis
3
Frame 20417, Direction z, Slice 19, Opacity 80, Zoom 14, Interpolate On, Not Flipped About X Axis
4
Frame 20431, Direction z, Slice 19, Opacity 80, Zoom 14, Interpolate On, Not Flipped About X Axis
5
Frame 20461, Direction z, Slice 19, Opacity 80, Zoom 14, Interpolate On, Not Flipped About X Axis
6

Figure 6.14: A series of overlay images (non consecutive) shows tracking of the endoscope during motion. $T_{O \Rightarrow ET}$ is first estimated using the method from Section 6.6 on image 1. Subsequent overlays use estimates of $T_{O \Rightarrow ET}$ calculated using equation 6.2. The green dot in the top right corner of each image indicates that the Optotrak succeeding tracking each of these frames. When insufficient IREDs are visible this is replaced with a red dot. Frames 3 and 4 show significant blurring of the video image due to endoscope motion. The striations are an artefact of the interlacing used by the frame grabber.

Age	Preop PSA	Stage	Grade	Margins	Post-Op PSA	8 Week Continence
67	9.8	pT2a	3+3	negative	< 0.03	Good, no pads
58	6.1	pT3b	3+4	negative	< 0.03	Good, no pads
68	11.2	pT2c	3+4(+5)	negative	< 0.03	1 pad/day
58	7.2	pT2c	3+4	negative	< 0.03	1 pad/day for security
57	3.8	pT3a	3+4	apex +ve	< 0.03	Good no pads

Table 6.5: Clinical outcomes for the five patients. There is no reason to believe that the system as implemented would have affected clinical outcomes.

6.9 Conclusions

The implementation of the image guidance system shown in this chapter is quite different to what was envisioned at the start of this thesis. This change was necessitated by the high tracking errors found in Chapter 5 and restated in Section 6.5. None the less we implemented an image guidance system in theatre which was found to be useful by the clinical staff. The system as shown can be implemented in theatre without the need for Optotrak system or the ultrasound probe. Direct registration and overlay can be done on a laptop attached the *daVinci*^{TM} assistant console. Some work on improving the user interfaces of the alignment and overlay software is required, but apart from that the system works for the limited case of static overlay.

Chapter 7

Conclusions and Further Work

7.1 Summary of Results

At the start of this thesis we proposed a method for incorporating preoperative images into the intra-operative endoscope image. At the core of this system is the transformation $T_{M \Rightarrow ES}$ which transforms individual voxels of the MRI image to on-screen coordinates. The method we proposed to find $T_{M \Rightarrow ES}$ was to break it into to two fundamentally separate transformations, firstly we transform the MRI image to the coordinate system of the Optotrak tracking system, using $T_{M \Rightarrow O}$, then we project the MRI image onto the screen using $T_{O \Rightarrow ES}$. This methodology for image guided surgery is well established in the literature, Birkfellner et al. (2002), Shahidi et al. (2002), and King et al. (1999) all present similar systems. In these cases transforms equivalent to $T_{M \Rightarrow O}$ were determined using fiducial markers, whereas we present a method (in Chapters 3 and 4) that allows determination of $T_{M \Rightarrow O}$ with out the need for fiducial markers. We invested considerable effort into developing this method and quantifying its accuracy, with significant successes.

In Chapter 3 we presented a novel algorithm for fitting inter patient CT pelvic bone to an MRI scan. As it is relatively easy to segment bone from CT we can treat this algorithm as a bone segmentation from MRI algorithm. On its own this algorithm compares well to the state of the art in literature, with RMS boundary errors only 0.6 mm greater than that achieved by Schmid and Magnenat-Thalmann (2008) (see table 3.11). It should be noted that Schmid and Magnenat-Thalmann (2008) used MRI images better suited to bone imaging and optimised their method to achieve a low RMS boundary error. In comparison we used MRI more suited to imaging the soft tissue and optimised our method to achieve a low registration error at the prostate.

In Chapter 4 we presented a method for registering ultrasound slices to the pseudo CT image generated in Chapter 3. By storing the ultrasound image as a set of non zero points rather than a set of slices we reduced the computational time required for registration in comparison with Penney et al. (2006), who used a similar method. We also optimised our method for use on the pseudo CT, which presents a slightly inaccurate CT image of the patient as opposed to the actual CT used by Penney et al. (2006). We showed that due to the use of the pseudo CT the distribution of ultrasound slices used for registration had a strong influence on the registration accuracy. We also suggested a method for checking the convergence of the registration algorithm in-theatre, by repeating the registration from randomly perturbed starting estimates. The method appeared robust, however at present its run-time is excessive (300 minutes). Finally we computed the overall error in estimating $T_{M \Rightarrow O}$. We found this error to be 7.3 mm RMS, corresponding to an on-screen error of roughly 30 pixels. To determine the effect of this error we now needed to calculate $T_{O \Rightarrow ES}$ and take the system into the theatre.

The decision to concentrate on estimating $T_{M \Rightarrow O}$ was based on the assumption that the determination of $T_{O \Rightarrow ES}$ would be trivial and the resulting errors would be small in comparison to the error in estimating $T_{M \Rightarrow O}$. We based this assumption partly on the results of Birkfellner et al. (2002), Shahidi et al. (2002), and King et al. (1999), but more so on a model of the endoscope tracking using the method of Fitzpatrick et al. (1998). For this model we assumed that the error in localising each IRED would be isotropic, normally distributed and homogeneous. We also were optimistic in our estimate of how many IREDs would be visible on the tracking collar at a given point in time. We didn't have a good estimate of the IRED tracking error so we looked at best and worst case possibilities. In the best case we assumed a IRED tracking error of 0.02 mm from Barnes et al. (2007), which resulted in negligible tracking error. In the worst case we assumed a IRED tracking error of 0.2 mm, from NDI (1992), which gave significant errors, but these were still less than the error in estimating $T_{M \Rightarrow O}$. However, when we started to test the system in-theatre it became clear, primarily from our attempts to calibrate the endoscope, that the tracking error achieved in practice was significantly greater than predicted.

In Chapter 5 we analysed the calibration and tracking errors. We first presented a novel method to estimate the IRED localisation error using the IRED location estimated by rigid body registration of the remaining IREDs. Using this we modelled the calibration and tracking of the endoscope. We came to the following conclusions. The assumption of isotropic, homogeneous, normally distributed error does not hold for optical tracking using the Optotrak Certus. This agrees with Wiles et al. (2004). We found that the expected value of the magnitude error was approximately 0.2 mm, agreeing with NDI (1992). However the error was anisotropic, non-homogeneous, and non-normally distributed. By using the estimated marker errors as a look-up table for Monte-Carlo simulation we quantified the effect of our erroneous initial assumptions, see Table 5.3. The actual tracking errors (65 pixels) are more than double that found assuming normally distributed, isotropic, homogeneous error (31 pixels [1]). We are not the first to show the effect of error anisotropy (see Wiles et al. (2008)). To our knowledge, however, we are the first to model the effect of non-normally distributed and non-homogeneous IRED localisation error. The effect is quite dramatic, and may explain the trend towards image guided surgery systems using direct alignment of the visible anatomy with the endoscopic image, rather than trying to track the endoscope. For a recent example see Su et al. (2009). In our case the very high endoscope tracking errors meant that it would not be possible to use the in-theatre data to assess the accuracy in finding $T_{M \Rightarrow O}$, as we had initially hoped.

The preceding work demonstrated that the image guidance system we had initially envisioned would not function until endoscope tracking errors ware addressed. In the meantime, however, we wanted to implement image guidance in-theatre to get an appreciation of the surgeon's reaction and interaction with it along with an understanding of an acceptable alignment error. We presented this in Chapter 6. The tracking errors were significantly larger as we did not filter the data to include only static data. The errors in tracking a moving endoscope were very large (around 170 pixels, approx three times the static error). This agrees with the results of Barnes et al. (2007). With the impracticality of in-theatre endoscope tracking established we limited our overlay system to showing overlays of the MRI based on a manual image to MRI registration using a wireframe image of the inner surface of the pubic arch. We implemented an overlay projection algorithm and user interface using OpenGL. It has not been possible to make any quantitative analysis of this method. However the feedback from the surgeon has been very encouraging, with the availability of MRI images during surgery aiding decision making and communication within the surgical team.

---

[1] In this case the error assuming normally distributed, isotropic, homogeneous error is also higher than we expected and actually roughly the same as the registration error from Chapter 4. This is due to the difficultly we had in tracking the IREDs. We found in practice that it was common to be tracking the endoscope with 5 or less IREDs, whereas we had assumed that at least 10 IREDs would be visible.

7.2 Future Work

There are several areas where future work could be directed, we will discuss these separately.

7.2.1 Shape Model Fitting

The results presented in Chapter 3 are quite promising. However the validation of them has been very limited. The availability of a larger data set of MRI pelvic data, ideally with matching CT data (or expert bone segmentations) would enable the algorithm to be properly validated. It is very unlikely that such a data set will be available for work on radical prostatectomy. There are two alternative applications for the algorithm we developed in Chapter 3. The first is orthopaedics, as per Schmid and Magnenat-Thalmann (2008). The second is radiotherapy planning from MRI as per Boettger et al. (2008). Seeking an opportunity to apply the algorithm in one of these areas may provide some useful results, both to validate the method and as a useful method for each application.

7.2.2 Registration Using Ultrasound

We did not complete a practical ultrasound registration algorithm. For the method to work in-theatre a substantially faster and more robust optimiser should be implemented. Two obvious choices are the widely used Levenberg-Marquardt algorithm [Levenberg (1944)] and the differential evolution algorithm we used in Chapter 3 [Price et al. (2005)]. With this done it may be worthwhile to further analyse the ultrasound data we collected in-theatre.

7.2.3 Endoscope Tracking

In Chapter 5 we proposed two methods to drastically improve the endoscope tracking accuracy. One was to assume (or physically enforce) a sliding pivot constraint at the trocar. The second was to constrain the motion using known motion axes taken from the *daVinci*^{TM} API. The promising results presented in Figure 5.20 were based on simulations and need to be confirmed with experimental results.

Throughout Chapter 5 we used a point base registration algorithm (SVD) that itself assumes isotropic normally distributed error. It would be worthwhile to repeat the various experiments in Chapter

5 using a point registration algorithm that does not require these assumptions, see Chu and Trendafilo (1998) or Batchelor and Fitzpatrick (2000).

Alternatively, investigation of hybrid tracking systems could be worthwhile. The optical tracking data could be supplemented with data from magnetic tracking, as per Birkfellner et al. (1998), or the *daVinci*^{TM} kinematics [Kwartowitz et al. (2009)].

Our modelling of the IRED tracking errors was done by sampling measured errors, rather than building a predictive model of IRED tracking errors. If a predictive model of IRED tracking errors was developed, it should enable some of the systematic tracking errors to be adjusted for. Such an adjustment could significantly improve the endoscope tracking accuracy. A further useful experiment would be to compare our method with the model of optical tracking systems presented by Sielhorst et al. (2007).

7.2.4 Clinical Implementation

In Chapter 6 we presented a limited image guidance system, that only allowed overlay for certain static positions. Nonetheless the feedback from clinicians was very encouraging. It would be very worthwhile to continue implementing this system on a larger cohort of patients. As it does not require endoscope tracking or ultrasound acquisition this could be done with minimal additional work. If some effort was investing in implementing an intuitive user interface this system could be used by the clinicians themselves.

7.3 Conclusions

At the start of this thesis we set two goals. Firstly to develop and test in theatre an image guided surgery system. Secondly, to determine the accuracy of the system and to identify the key contributors to the system error. Section 7.1 detailed how these goals were met. We can now combine these results to answer the question of whether the system is clinically useful, and if so what steps to take to improve the system.

Repeating Table 5.4 here, this time with the pixel errors converted to errors in mm, we can put the system performance into context. Table 7.1 shows these results. The results shown in Table 7.1 raise the question of what an error of 18.0 mm means clinically.

7.3.1 Defining Clinical Utility

In this section we start with the assumption that as the system becomes more accurate it becomes more useful to the surgeon. For the purposes of visualisation we can define a measure of clinical utility. If the system has no effect on the clinical outcome the clinical utility is zero. If the system has no error the clinical utility is one. A real system will fall somewhere between these two points. [2] By estimating clinical utility values for errors greater than zero, we can fit an approximate curve to show how the system error effects clinical utility. Figure 7.1 shows a plot of clinical utility versus system error in mm.

The first intermediate point defined is the current system accuracy. As discussed in Chapter 6 the

---

[2] A negative clinical utility could occur if the system has a negative effect on clinical outcome. This may be the case if the system has a high error but the error has been underestimated. The obvious need to avoid this situation is the reason this thesis has placed such high importance on measuring the system error.

Error	Reference	X Err. (mm)	Y Err. (mm)	Magnitude (mm)
Shape Modelling Error	Figure 3.11	3.3	3.3	4.7
Registration Error	Figure 4.18	3.3	3.1	4.5
Endo. Calibration Error (Intrinsic)	Section 5.5	0.5	0.5	0.7
Endo.Calibration Error (Extrinsic)	Section 5.5	6.2	7.3	9.6
Endo. Calibration Error (Total)	Figure 5.12	6.3	7.4	9.7
Endo.Tracking Error	Figure 5.15	12.3	8.8	15.2
Total Error	Figure 5.23	14.2	11.1	18.0

Table 7.1: A summary of the system errors. The total system error is dominated by the endoscope tracking error. The errors in this table are the same as in Table 5.4 but have been converted to mm by back projecting them onto a plane 200 mm from the endoscope lens.

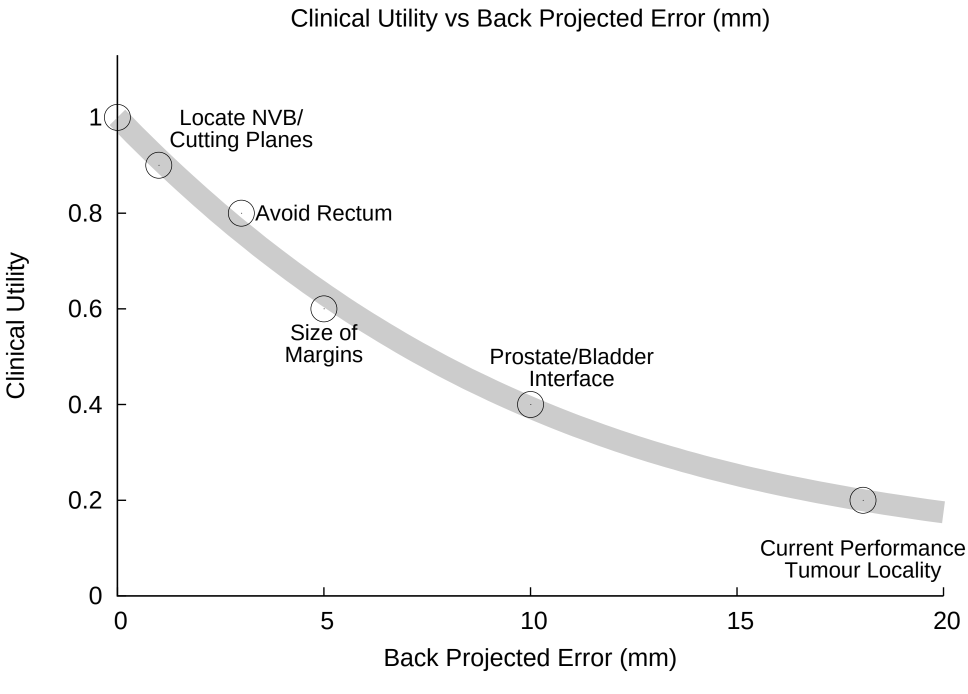

Figure 7.1: Clinical utility versus system error in mm. Clinical utility is an imprecise figure based on the assumption that at zero error the system is as good as possible, so clinical utility is one. If the system has no effect on clinical outcome then clinical utility is zero. Intermediate points have been estimated based on estimates of what the system could be used for at different accuracies.

system is useful at this accuracy. The surgeon can refer to MRI to see the general location (which side of the prostate) the tumours are located. A clinical utility of 0.2 was assigned to this point. At an error of 10 mm the system may be useful for showing the shape of the prostate at the interface with the bladder. At this point the surgeon is often unsure of the size and shape of the medial lobe of the prostate. Being able to show this anatomy would increase the clinical utility. At an error of 5 mm the system may be useful for defining the size of the margins around the prostate capsule. At errors of less than 5 mm the system could be used to show fine anatomy, like the nuerovascular bundles and the cutting planes around the rectum.

Clearly the values of clinical utility and hence the chart shown in Figure 7.1 are somewhat arbitrary at present. The values will also vary from patient to patient. The only way to accurately populate Figure 7.1 is to perform a significant number of procedures using guidance systems of known accuracies. At present Figure 7.1 merely provides a roadmap by which future improvement can be measured. The next stage is to create a similar plot detailing the feasibility of achieving a given accuracy.

7.3.2 Defining And Improving Accuracy

Table 7.1 gives the current system accuracy for the various components. Strategies for reducing errors were discussed in Section 7.2. To visualise this process we can assign a feasibility value to various error levels and plot them. Similarly to the preceding discussion on clinical utility, limits on feasibility can be set at the current known levels. A system with a total error of 18.0 is entirely feasible, hence a feasibility of 1.0. Similarly a system with an error of zero is impossible, hence a feasibility of zero. Intermediate points can be plotted with good accuracy based on the work presented in this thesis. Figure 7.2 shows a plot of feasibility versus system error.

Some improvements are relatively easy to implement. Accuracy could be significantly improved by calibrating the endoscope in a fixture, so that its position could be measured more accurately. Whilst various approaches could be used to improve the accuracy of the proposed ultrasound to MRI registration process, it may be more feasible to utilise fiducial markers. Past this point the feasibility of proposed improvement drops. It seems quite likely that the endoscope tracking accuracy could be improved by allowing for systematic tracking errors and/or constraining the allowable motion of the endoscope. Beyond this any improvement would rely on methods we have not yet thought of.

7.3.3 Combining Feasibility and Utility

Combining Figures 7.2 and Figures 7.1 we can gain an understanding of where the current system stands, and what should be possible with further development. At present the system is useful, but only in a general way. There is a clear roadmap to reduce the system error to around 5mm, which would make the system significantly more useful. At this level the system may be used to help define surgical margins. We have not specified a method to improve the system beyond this accuracy here. However it is likely that a method would be found at some point in the future. The only way to properly understand the clinical utility of such changes is to implement the system on a significant number of patients.

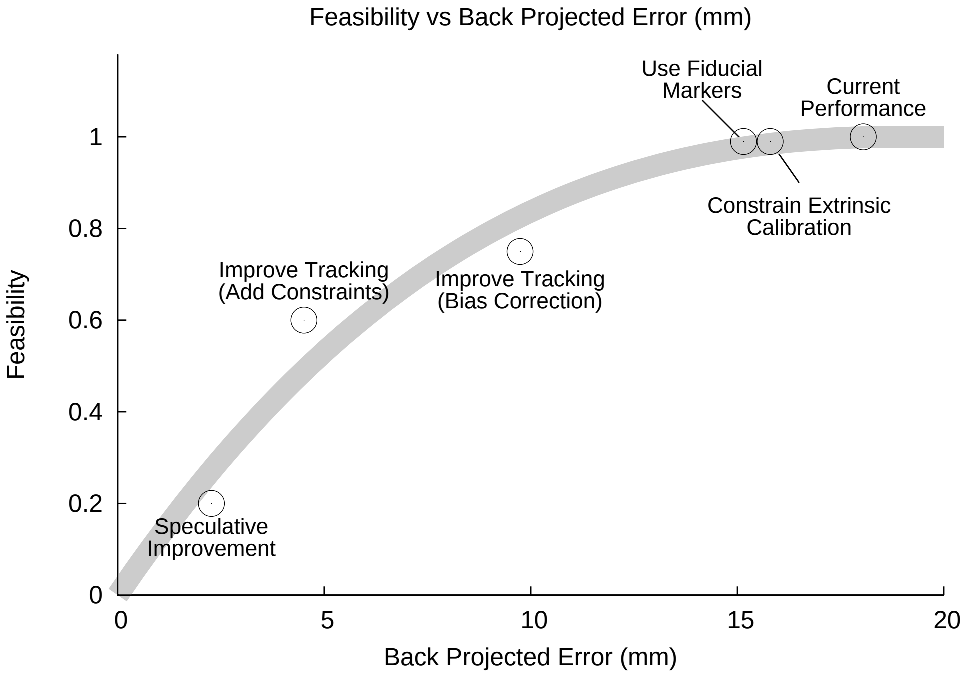

Figure 7.2: The feasibility of a given system versus its accuracy. Limits can be set at a feasibility of 1 for the current system accuracy of 18.0 mm, and zero for a system with zero error. Various strategies for reducing the total system error can then be applied. Some are more easily applied than others, so have been assigned different feasibility scores. The effect on accuracy is based on the values shown in Table 7.1.

Glossary

API application programming interface. 159, 185

ASM active shape model. 50

BEP bone edge probability. 104

CT computed tomography. 22, 24, 25, 29, 44, 47, 48, 50, 51, 56, 59, 61, 63, 69, 70, 75, 78, 80, 82, 84, 86, 90, 92, 94, 96, 97, 99, 101, 114, 115, 117, 119–121, 123–132, 165, 183, 185, 241

DOF degree of freedom. 69, 92, 241

FEM finite element model. 66

FLE fiducial localisation error. 109, 136, 137

FRE fiducial registration error. 108

GPU graphics processing unit. 66

ILE IRED localisation error. 109, 110, 141–144, 146, 148, 154, 155, 159, 160, 162, 163

IRE IRED registration error. 108, 109, 111

IRED infra red emitting diodes. 60, 62, 65, 102, 103, 105, 108–112, 129, 135–146, 148, 151, 154–157, 159, 160, 162, 166, 167, 169, 171, 173, 176, 179, 180, 184, 186

MRI magnetic resonance imaging. 21–26, 28–30, 33, 35, 43–48, 50–53, 56, 57, 59, 61–65, 67, 69, 70, 72, 76, 78, 80, 82, 84, 86, 88, 90, 92, 93, 96, 97, 99, 101, 104, 114, 130, 132, 154, 165, 166, 168, 173–179, 183, 185, 188

NMI normalised mutual information. 47, 55, 70, 73, 75, 82, 88–92, 96, 97, 245

PCA principal component analysis. 54

RARP robot assisted radical prostatectomy. 21, 22, 25, 28, 33, 35, 42, 43, 46, 59, 62, 65, 70, 132, 165

RMS root mean square. 51, 65, 87, 94, 96, 99, 109, 116, 120, 127, 130, 131, 138, 166, 183, 184, 249, 250

SDM statistical deformation model. 52, 54, 55, 70, 71

SSM statistical shape model. 29, 52

SVD singular value decomposition. 71, 135, 157, 159, 185

TRE target registration error. 30, 60, 64, 80–82, 84–92, 94, 96, 99, 109, 110, 112, 120–125, 127, 130–132, 136, 137, 142, 160, 162, 163, 249, 250

TRUS trans-rectal ultrasound. 47, 61, 63

Bibliography

Adhami, L. and Coste-Maniére, E. (2003), Optimal planning for minimally invasive surgical robots, *IEEE Transactions on Robotics and Automation* **19**(5), 854–863.

Atug, F., Castle, E. P., Woods, M., Davis, R. and Thomas, R. (2006), Robotics in urologic surgery: an evolving new technology., *International Journal of Urology* **13**(7), 857–863.

Aylward, S., Jomier, J., Guyon, J.-P. and Weeks, S. (2002), Intra-operative 3D ultrasound augmentation, *in Proceedings of IEEE International Symposium on Biomedical Imaging*, pp.421–424.

Baldock, R. and Graham, J. (eds.) (2000), *Model-Based Methods in Analysis of Biomedical Images*, Oxford University Press.

Barnes, P., Baldock, C., Meikle, S. and Fulton, R. (2007), Benchmarking of a Motion Sensing System for Medical Imaging and Radiotherapy, *in IEEE Nuclear Science Symposium Conference Record*, Vol. 6, pp.4513–4520.

Barratt, D. C., Chan, C. S., Edwards, P. J., Penney, G. P., Slomczykowski, M., Carter, T. J. and Hawkes, D. J. (2008), Instantiation and registration of statistical shape models of the femur and pelvis using 3D ultrasound imaging, *Medical Image Analysis* **12**(3), 358–374.

Barratt, D., Penney, G., Chan, C., Slomczykowski, M., Carter, T., Edwards, P. and Hawkes, D. (2006), Self-calibrating 3D-ultrasound-based bone registration for minimally invasive orthopedic surgery, *IEEE Transactions on Medical Imaging* **25**(3), 312–323.

Batchelor, P. G. and Fitzpatrick, J. M. (2000), A Study of the Anisotropically Weighted Procrustes Problem, *in MMBIA '00: Proceedings of the IEEE Workshop on Mathematical Methods in Biomedical Image Analysis*, IEEE Computer Society, Washington, DC, USA, pp.212–218.

Bharatha, A., Hirose, M., Hata, N., Warfield, S. K., Ferrant, M., Zou, K. H., Suarez-Santana, E., Ruiz-Alzola, J., D'Amico, A., Cormack, R. A., Kikinis, R., Jolesz, F. A. and Tempany, C. M. (2001), Evaluation of three-dimensional finite element-based deformable registration of pre- and intraoperative prostate imaging., *Medical Physics* **28**(12), 2551–2560.

Birkfellner, W., Figl, M., Huber, K., Watzinger, F., Wanschitz, F., Hummel, J., Hanel, R., Greimel, W., Homolka, P., Ewers, R. and Bergmann, H. (2002), A head-mounted operating binocular for augmented reality visualization in medicine - design and initial evaluation, *IEEE Transactions on Medical Imaging* **21**(8), 991–997.

Birkfellner, W., Franz, W., Wanschitz, F., Ewers, R. and Bergmann, H. (1998), Calibration of Tracking Systems in a Surgical Environment, *IEEE Transactions on Medical Imaging* **17**, 737–742.

Boettger, T., Nyholm, T., Karlsson, M., Nunna, C. and Celi, J. C. (2008), Radiation therapy planning and simulation with magnetic resonance images, *in* M. I. Miga and K. R. Cleary (eds.), *Medical Imaging 2008: Visualization, Image-guided Procedures, and Modeling*, Vol. 6918, SPIE, pp.69181C–1 – 69181C–11.

Braeckman, J., Autier, P., Garber, C., Marichal, M. P., Soviany, C., Nir, R., Nir, D., Michielsen, D., Bleiberg, H., Egevad, L. and Emberton, M. (2007), Computer-aided ultrasonography (HistoScanning): a novel technology for locating and characterising prostate cancer, *British Journal of Urology* **101**, 293–298.

Brailsford, J. F. (1946), Roentgen's discovery of X rays, their application to medicine and surgery, *British Journal of Radiology* **19**(227), 453–461.

Breslow, N., Chan, C. W., Dhom, G., Drury, R. A., Franks, L. M., Gellei, B., Lee, Y. S., Lundberg, S., Sparke, B., Sternby, N. H. and Tulinius, H. (1977), Latent carcinoma of prostate at autopsy in seven areas., *International Journal of Cancer* **120**(5), 680–688.

Bro-Nielsen, M. (1996), Surgery simulation using fast finite elements, *in Proceedings of the 4th International Conference on Visualization in Biomedical Computing*, Springer-Verlag, London, UK, pp.529–534.

Carballido-Gamio, J., Belongie, S. and Majumdar, S. (2004), Normalized cuts in 3-D for spinal MRI segmentation, *IEEE Transactions on Medical Imaging* **23**(1), 36–44.

Cash, D. M., Miga, M. I., Sinha, T. K., Galloway, R. L. and Chapman, W. C. (2005), Compensating for intraoperative soft-tissue deformations using incomplete surface data and finite elements, *IEEE Transactions on Medical Imaging* **24**(11), 1479–1491.

Challacombe, B. J., Khan, M. S., Murphy, D. and Dasgupta, P. (2006), The history of robotics in urology., *World Journal of Urology* **24**(2), 120–127.

Chan, C. S. K., Barratt, D. C., Edwards, P. J., Penney, G. P., Slomczykowski, M., Carter, T. J. and Hawkes, D. J. (2004), Cadaver validation of the use of ultrasound for 3D model instantiation of bony anatomy in image guided orthopaedic surgery, *in Medical Image Computing and Computer Assisted Intervention*, Vol. 3217 of *LNCS*, Springer, pp.397–404.

Chen, D., Mayer, E., Vale, J., Anstee, A., Mei, L., Darzi, A., Yang, G.-Z. and Edwards, P. (2008), A 3D Stereo System to Assist Surgical Treatment of Prostate Cancer, *in AMIARCS (Satellite WorkShop of Medical Imaging Computing and Computer Aided Interventions)*.

Chu, M. T. and Trendafilo, N. T. (1998), On a differential equation approach to the weighted orthogonal procrustes problem, *Statistics and Computing* **8**, 125–133.

Cootes, T. F. and Taylor, C. J. (1995), Combining Point Distribution Models with Shape Models Based on Finite Element Analysis, *Image Vis. Comput.,* **13**, 403–409.

Cootes, T. F. and Taylor, C. J. (2004), Anatomical statistical models and their role in feature extraction, *British Journal of Radiology* **77**(2), S133–S139(1).

Cootes, T. F., Taylor, C. J., Cooper, D. H. and Graham, J. (1994), Active Shape Models - Their Training and Application, *Computer Vision and Image Understanding* **61**(1), 38–59.

Crum, W. R., Tanner, C. and Hawkes, D. J. (2005), Anisotropic multi-scale fluid registration: evaluation in magnetic resonance breast imaging, *Physics in Medicine and Biology* **50**, 5153–5174.

Dasgupta, P., Challacombe, B., Murphy, D. and Khan, M. S. (2006), Coming full circle in robotic urology., *British Journal of Urology International* **98**(1), 4–5.

Dasgupta, P. and Kirby, R. S. (2009), Outcomes of robotic assisted radical prostatectomy., *International Journal of Urology* **16**(3), 244–248.

Davies, B. L., Hibberd, R. D., Ng, W. S., Timoney, A. G. and Wickham, J. E. A. (1992), A Robotics Assistant For Prostate Surgery, *in Proceedings of the Annual International Conference of the IEEE Engineering in Medicine and Biology Society*, Vol. 14, pp.1052–1053.

Davies, R. H., Twining, C. J., Cootes, T. F., Waterton, J. C. and Taylor, C. J. (2002), 3D Statistical Shape Models Using Direct Optimisation of Description Length, *in European Conference on Computer Vision*, *European Conference on Computer Vision LNCS 2352* **2352/2002**, pp.1–17.

Dey, D., Gobbi, D. G., Slomka, P. J., Surry, K. J. M. and Peters, T. M. (2002), Automatic fusion of freehand endoscopic brain images to three-dimensional surfaces: creating stereoscopic panoramas, *IEEE Transactions on Medical Imaging* **21**(1), 23–30.

Dice, L. R. (1945), Measures of the Amount of Ecologic Association Between Species, *Ecology* **26**, 297–302.

Edwards, P. J., King, A. P., Maurer, C. R., de Cunha, D. A., Hawkes, D. J., Hill, D. L. G., Gaston, R. P., Fenlon, M. R., Jusczyzck, A., Strong, A. J., Chandler, C. L. and Gleeson, M. J. (2000), Design and Evaluation of a System for Microscope-Assisted Guided Interventions (MAGI), *IEEE Transactions on Medical Imaging* **19**(11), 1082–1093.

Fei, B., Duerk, J. L., Boll, D. T., Lewin, J. S. and Wilson, D. L. (2003a), Slice-to-volume registration and its potential application to interventional MRI-guided radio-frequency thermal ablation of prostate cancer, *IEEE Transactions on Medical Imaging* **22**(4), 515–525.

Fei, B., Kemper, C. and Wilson, D. L. (2003b), A comparative study of warping and rigid body registration for the prostate and pelvic MR volumes., *Computerized Medical Imaging and Graphics* **27**(4), 267–281.

Fitzpatrick, J. M. (2009), Fiducial registration error and target registration error are uncorrelated, *in* M. I. Miga and K. H. Wong (eds.), *Medical Imaging 2009: Visualization, Image-Guided Procedures, and Modeling, Proceedings of the SPIE,*, Vol. 7261, pp.102–112.

Fitzpatrick, J. M., Hill, D. L. G. and Maurer Jr, C. R. (2000), *Handbook of Medical Imaging*, Vol. II, SPIE Press, Chapter Image Registration.

Fitzpatrick, J. M., West, J. B. and Maurer Jr, C. R. (1998), Predicting Error in Rigid-Body, Point-based Registration, *IEEE Transactions on Medical Imaging* **17**(5), 694–702.

Freedman, D., Radke, R. J., Zhang, T., Jeong, Y., Lovelock, D. M. and Chen, G. T. Y. (2005), Model-based segmentation of medical imagery by matching distributions, *IEEE Transactions on Medical Imaging* **Volume 24**(3), 281–292.

Fried, M. P., Kleefield, J., Gopal, H., Reardon, E., Ho, B. T. and Kuhn, F. A. (1997), Image-guided endoscopic surgery: results of accuracy and performance in a multicenter clinical study using an electromagnetic tracking system., *Laryngoscope* **107**(5), 594–601.

Fripp, J., Crozier, S., Warfield, S. K. and Ourselin, S. (2007), Automatic segmentation of the bone and extraction of the bone-cartilage interface from magnetic resonance images of the knee, *Physics in Medicine and biology* **52**, 1617–1631.

Gleason, D. F. (1977), *Urologic Pathology: The Prostate*, Lea and Febiger, Philadelphia, Chapter The Veteran's Administration Cooperative Urologic Research Group: histologic grading and clinical staging of prostatic carcinoma, pp.171–198.

Goksel, O., Salcudean, S. E., DiMaio, S. P., Rohling, R. and Morris, J. (2005), 3D needle-tissue interaction simulation for prostate brachytherapy, *in* J. Duncan and G. Gerig (eds.), *Medical Image Computing and Computer Aided Interventions*, Vol. 3749 of *LNCS*, Springer, pp.827–834.

Guthart, G. and Salisbury, J.K., J. (2000), The Intuitive telesurgery system: overview and application, *in Proceedings of the IEEE International Conference on Robotics and Automation.*, Vol. 1, pp.618–621.

Hawkes, D. J., Barratt, D., Blackall, J. M., Chan, C., Edwards, P. J., Rhode, K., Penney, G. P., McClelland, J. and Hill, D. L. G. (2005), Tissue deformation and shape models in image-guided interventions: a discussion paper., *Medical Image Analysis* **9**(2), 163–175.

Heikkila, J. and Silven, O. (1997), A Four-step Camera Calibration Procedure with Implicit Image Correction, *in Proceedings of the 1997 Conference on Computer Vision and Pattern Recognition*, IEEE Computer Society, Washington, DC, USA, p.1106.

Hoad, C. L. and Martel, A. L. (2002), Segmentation of MR images for computer-assisted surgery of the lumbar spine, *Physics in Medicine & Biology* **47**(19), 3503–3517.

Hoogeman, M. S., van Herk, M., de Bois, J. and Lebesque, J. V. (2004), Strategies to reduce the systematic error due to tumor and rectum motion in radiotherapy of prostate cancer, *Radiotherapy and Oncology* **74**, 177–185.

Housden, R. J., Gee, A. H., Treece, G. M. and Prager, R. W. (2006), Sensorless reconstruction of freehand 3D ultrasound data, *in* R. Larsen, M. Nielsen and J. Sporring (eds.), *Medical Image Computing and Computer Assisted Interventions*, Vol. 4191 of *LNCS*, Springer, pp.356–363.

Hricak, H. (2005), MR imaging and MR spectroscopic imaging in the pre-treatment evaluation of prostate cancer, *British Journal Radiology* **78**(Special Issue 2), S103–111.

Hu, M., Penney, G., Edwards, P., Figl, M. and Hawkes, D. (2007), 3D Reconstruction of Internal Organ Surfaces for Minimal Invasive Surgery, *in* N. Ayache, S. Ourselin and A. Maeder (eds.), *Medical Image Computing and Computer Aided Interventions*, Vol. 4791 of *LNCS*, Springer, pp.68–77.

Hu, M., Wen, F., Yuan, B. and Tang, X. (2002), A novel method of three-dimensional reconstruction based on the double algebra, *in Proceedings. 2002 IEEE Region 10 Conference on Computers, Communications, Control and Power Engineering*, Vol. 1, pp.549–552.

Hu, Y., Morgan, D., Ahmed, H. U., Pends, D., Sahu, M., Allen, C., Emberton, M., Hawkes, D. and Barratt, D. (2008), A Statistical Motion Model Based on Biomechanical Simulations for Data Fusion during Image-Guided Prostate Interventions, *in MICCAI*.

Kapur, T., Beardsley, P., Gibson, S., Grimson, W. and Wells, W. (1998), Model-based segmentation of clinical knee MRI, *in IEEE Internationall Workshop on Model-Based 3D Image Analysis*, pp.97–106.

King, A. P., Edwards, P. J., Pike, M. R., Hill, D. L. G. and Hawkes, D. J. (1999), An Analysis of Calibration and Registration Errors in an Augmented Reality System for Microscope-Assisted Guided Interventions, *in Medical Image Understanding and Analysis*.

Kirby, R. S., Eeles, R. A., Kote-Jarai, Z., Guy, M., Easton, D. and Fitzpatrick, J. M. (2010), Screening for prostate cancer: the way ahead., *British Journal of Urology International* **105**(3), 295–297.

Koikkalainen, J., Tlli, T., Lauerma, K., Mattila, K. A. E., Lilja, M. and Ltjnen, J. (2008), Methods of Artificial Enlargement of the Training Set for Statistical Shape Models, *IEEE Transactions on Medical Imaging* **27**(11).

Kuhnapfel, U., Cakmak, H. K. and Maas, H. (1999), 3D modeling for endoscopic surgery, *in Proceedings of the IEEE Symposium on Simulation*, pp.22–32.

Kwartowitz, D. M., Miga, M. I., Herrel, S. D. and Galloway, R. L. (2009), Towards image guided robotic surgery: multi-arm tracking through hybrid localization, *International Journal of Computer assisted radiology* **4**, 281–286.

Lalonde, N.-M., Dansereau, J., Aissaoui, R., Pauget, P. and Cinquin, P. (2003), Differences between pelvic skin and bone landmark identification in different seated positions on spinal-cord injured subjects., *IEEE Transactions on Biomedical Engineering* **50**(8), 958–966.

Lamecker, H., Seebass, M., Hege, H. C. and Deuflhard, P. (2004), A 3D statistical shape model of the pelvic bone for segmentation, *in* J. M. Fitzpatrick and M. Sonka (eds.), *Proceedings of the SPIE Medical Imaging 2004: Image Processing.*, Vol. 5370, pp.1341–1351.

Lee, Z., Sodee, D. B., Resnick, M. and Maclennan, G. T. (2005), Multimodal and three-dimensional imaging of prostate cancer., *Computerized Medical Imaging and Graphics* **29**(6), 477–486.

Leroy, A., Mozier, P., Payan, Y. and Troccaz, J. (2004), Rigid registration of freehand 3D ultrasound and CT-scan kidney images, *in Medical Image Computing and Computer Assisted Intervention*, Vol. 3216 of *LNCS*, Springer, pp.837 – 844.

Leven, J., Burschka, D., Kumar, R., Zhang, G., Blumenkranz, S., Dai, X., Awad, M., Hager, G. D., Marohn, M., Choti, M., Hasser, C. and Taylor, R. H. (2005), DaVinci Canvas: A telerobotic surgical system with integrated, robot-assisted, laparoscopic ultrasound capability, *in Medical Image Computing and Computer Assisted Intervention*, Vol. 3749 of *LNCS*, Springer, pp.811 – 818.

Levenberg, K. (1944), A Method for the Solution of Certain Non-Linear Problems in Least Squares, *The Quarterly of Applied Mathematics* **2**, 164–168.

Lin, H. C., Shafran, I., Murphy, T. E., Okamura, A. M., Yuh, D. D. and Hager, G. D. (2005), Automatic detection and segmentation of robot-assisted surgical motions, *in Medical Image Computing and Computer Assisted Intervention*, Vol. 3749 of *LNCS*, Springer, pp.802–810.

Lorigo, L. M., Faugeras, O., Grimson, W. E. L., Keriven, R. and Kikinis, R. (1998), Segmentation of bone in clinical knee MRI using texture-based geodesic active contours, *in* W. M. Wells, A. C. F. Colchester and S. L. Delp (eds.), *Medical Image Computing and Computer Assisted Intervention*, Vol. 1496 of *LNCS*, Springer, pp.1195–1204.

Maes, F., Collignon, A., Vandermeulen, D., Marchal, G. and Suetens, P. (1997), Multimodality image registration by maximization of mutual information., *IEEE Transactions on Medical Imaging* **16**(2), 187–198.

Maurer, C. R., Gaston, R. P., Hill, D. L. G., Gleeson, M. J., Taylor, M. G., Fenlon, M. R., Edwards, P. J. and Hawkes, D. J. (1999), AcouStick: A tracked A-mode ultrasonography system for registration in image-guided surgery, *in Medical Image Computing and Computer Assisted Intervention*, Vol. 1679, pp.953–963.

Mavrich, H. V. (2006), Da Vinci advanced robotic laparoscopic surgery: origin and current clinical application in urology, and comparison with open and laparoscopic surgery, *Actas Urol Esp* **30**(1), 1–12.

McLaughlin, R. A., Hipwell, J. H., Hawkes, D. J., Noble, J. A., Byrne, J. V. and Cox, T. C. S. (2005), A comparison of a similarity-based and a feature-based 2D-3D registration method for neurointerventional use., *IEEE Transactions on Medical Imaging* **24**(8), 1058–1066.

Mei, L., Figl, M., Ruekert, D., Darzi, A. and Edwards, P. (2008), Sample Sufficiency and Number of Modes to Retain in Statistcal Shape Modelling, *in Medical Image Computing and Computer Assisted Intervention*, Vol. 5241 of *LNCS*, Springer, pp.425–433.

Mohamed, S. S. and Salama, M. M. A. (2008), Prostate Cancer Spectral Multifeature Analysis Using TRUS Images, *IEEE Transactions on Medical Imaging* **27**(4), 548–556.

Mori, K., Deguchi, D., Kitasaka, T., Suenaga, Y., Takabatake, H., Mori, M., Natori, H. and Maurer, C. R. (2006), Bronchoscope Tracking Based on Image Registration Using Multiple Initial Starting Points Estimated by Motion Prediction, *in Medical Image Computing and Computer Assisted Intervention*, Vol. 4191 of *LNCS*, Springer, pp.645–652.

Mountney, P., Stoyanov, D., Davison, A. and Yang, G.-Z. (2006), Simultaneous stereoscope localization and soft-tissue mapping for minimal invasive surgery, *in Medical Image Computing and Computer Assisted Intervention*, Vol. 4190 of *LNCS*, Springer, pp.347–354.

Mourgues, F. and Coste-Maniére, E. (2002), Flexible Calibration of Actuated Stereoscopic Endoscope for Overlay in Robot Assisted Surgery, *in Medical Image Computing and Computer Aided Interventions*, Vol. 2488 of *LNCS*, Springer, pp.25–34.

NDI (1992), *Optotrak Certus Specifications*, NDI, Ontario, 103 Randall Drive. Waterloo, Ontario, Canada, N2V 1C5.

Office of National Statistics (2008), Cancer Registration Statistics England 2008, Technical Report, Office of National Statistics,, http://www.statistics.gov.uk/.

Paul, P., Fleig, O. and Jannin, P. (2005), Augmented virtuality based on stereoscopic reconstruction in multimodal image-guided neurosurgery: methods and performance evaluation, *IEEE Transactions on Medical Imaging* **24**(11), 1500–1511.

Penney, G. P., Barratt, D. C., Chan, C. S. K., Slomczykowski, M., Carter, T. J., Edwards, P. J. and Hawkes, D. J. (2006), Cadaver validation of intensity-based ultrasound to CT registration, *Medical Image Analysis* **10**, 385–395.

Penney, G. P., Blackall, J. M., Hamady, M. S., Sabharwal, T., Adam, A. and Hawkes, D. J. (2004), Registration of freehand 3D ultrasound and magnetic resonance liver images., *Medical Image Analysis* **8**(1), 81–91.

Peters, T. and Cleary, K. (eds.) (2008), *Image-Guided Interventions*, Springer, ISBN: 978-0-387-73856-7.

Porter, B. C., Rubens, D. J., Strang, J. G., Smith, J., Totterman, S. and Parker, K. J. (2001), Three-dimensional registration and fusion of ultrasound and MRI using major vessels as fiducial markers, *IEEE Transactions on Medical Imaging* **20**(4), 354–359.

Prager, R. W., Rohling, R. N., Gee, A. H. and Berman, L. (1998), Rapid calibration for 3D freehand ultrasound., *Ultrasound in Medicine and Biology* **24**(6), 855–869.

Price, K., Storn, R. and Lampinen, J. (2005), *Differential Evolution - A Practical Approach to Global Optimization*, Springer, ISBN: 3-540-20950-6.

Rasch, C., Barillot, I., Remeijer, P., Touw, A., van Herk, M. and Lebesque, J. V. (1999), Definition of the prostate in CT and MRI: a multi-observer study., *International Journal of Radiation Oncology Biol. Phys.* **43**(1), 57–66.

Rueckert, D., Frangi, A. F. and Schnabel, J. A. (2001), Automatic Construction of 3D Statistical Deformation Models Using Non-rigid Registration, *in Medical Image Computing and Computer Assisted Intervention*, Vol. 2208 of *LNCS*, Springer, pp.77–84.

Rueckert, D., Sonoda, L. I., Hayes, C., Hill, D. L. G. and Leach, M. O. (1999), Nonrigid registration using free-form deformations: Application to breast MR images, *IEEE Transactions on Medical Imaging* **18**, 712–721.

Schmid, J. and Magnenat-Thalmann, N. (2008), MRI Bone Segmentation Using Deformable Models and Shape Priors, *in Medical Imaging Computing and Computer Aided Interventions*, Vol. 5241 of *LNCS*, Springer, pp.119–126.

Shahidi, R., Bax, M. R., Maurer, C. R., Johnson, J. A., Wilkinson, E. P., Wang, B., West, J. B., Citardi, M. J., Manwaring, K. H. and Khadem, R. (2002), Implementation, calibration and accuracy testing of an image-enhanced endoscopy system., *IEEE Transactions on Medical Imaging* **21**(12), 1524–1535.

Shao, W., Wu, R., Ling, K. V., Thng, C. H., Ho, H. S. S., Cheng, C. W. S. and Ng, W. S. (2006), Evaluation on similarity measures of a surface-to-image registration technique for ultrasound images, *in* R. Larsen, M. Nielsen and J. Sporring (eds.), *Medical Image Computing and Computer Assisted Intervention*, Vol. 4191 of *LNCS*, Springer, pp.742–749.

Sielhorst, T., Baur, M., Wenisch, O., Klinker, G. and Navab, N. (2007), Online Estimation of the Target Registration Error for n-Ocular Optical Tracking Systems, *in Medical Image Computing and Computer Aided Interventions*, Vol. 4792 of *LNCS*, Springer, pp.652–659.

Skrinjar, O., Nabavi, A. and Duncan, J. (2001), A stereo-guided biomechanical model for volumetric deformation analysis, *in IEEE Workshop on Mathematical Methods in Biomedical Image Analysis*, pp.95–102.

Skrinjar, O., Tagare, H. and Duncan, J. (2000), Surface growing from stereo images, *in Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition.*, Vol. 2, pp.571–576.

Soler, L., Nicolau, S., Schmid, J., Koehl, C., Marescaux, J., Pennec, X. and Ayache, N. (2004), Virtual Reality and Augmented Reality in Digestive Surgery, *in Proceedings of the third IEEE and ACM international Symposium on Mixed and Augmented Reality*, pp.278–279.

State, A., Chen, D. T., Tector, C., Brandt, A., Chen, H., Ohbuchi, R., Bajura, M. and Fuchs, H. (1994), Observing a volume rendered fetus within a pregnant patient, *in IEEE Conference on Visualization*, Vol. CP41, pp.364–368.

Stefansic, J. D., Herline, A. J., Shyr, Y., Chapman, W. C., Fitzpatrick, J. M., Dawant, B. M. and Galloway, R. L. (2000), Registration of physical space to laparoscopic image space for use in minimally invasive hepatic surgery, *IEEE Transactions on Medical Imaging* **19**(10), 1012–1023.

Stoyanov, D., Darzi, A. and Yang, G.-Z. (2005), Laparoscope self-calibration for robotic assisted minimally invasive surgery, *in Medical Image Computing and Computer Assisted Intervention*, Vol. 3750 of *LNCS*, Springer, pp.114–121.

Strzelecki, M. (2004), Texture boundary detection using network of synchronised oscillators, *Electronics Letters* **40**(8), 466–467.

Studholme, C., Hill, D. L. G. and Hawkes, D. J. (1999), An Overlap Invariant Entropy Measure of 3D Medical Image Alignment, *Pattern Recognition* **32**(1), 71–86.

Su, L., Vagvolgyi, B. P., Agarwal, R., Reiley, C. E., Taylor, R. H. and Hager, G. D. (2009), Augmented Reality During Robot-assisted Laparoscopic Partial Nephrectomy: Toward Real-Time 3D-CT to Stereoscopic Video Registration, *Urology* **73**, 896–900.

Szpala, S., Wierzbicki, M., Guiraudon, G. and Peters, T. (2005), Real-Time Fusion of Endoscopic Views With Dynamic 3-D Cardiac Images: A Phantom Study, *IEE Transactions on Medical Imaging* **24**(9), 1207–1215.

Takenaka, A., Leung, R. A., Fujisawa, M. and Tewari, A. K. (2006), Anatomy of autonomic nerve component in the male pelvis: the new concept from a perspective for robotic nerve sparing radical prostatectomy., *World Journal of Urology* **24**(2), 136–143.

Tan, G. Y., , Goel, R. K., Kaouk, J. H. and Tewari, A. K. (2009), Technological Advances in Robotic-Assisted Laparoscopic Surgery, *Urologic Clinics of North America* **36**(2), 237–249.

Taylor, L. S., Rubens, D. J., Porter, B. C., Wu, Z., Baggs, R. B., di Sant'Agnese, P. A., Nadasdy, G., Pasternack, D., Messing, E. M., Nigwekar, P. and Parker, K. J. (2005), Prostate cancer: three-dimensional sonoelastography for in vitro detection., *Radiology* **237**(3), 981–985.

Tewari, A., Rao, S., Martinez-Salamanca, J. I., Leung, R., Ramanathan, R., Mandhani, A., Vaughan, E. D., Menon, M., Horninger, W., Tu, J. and Bartsch, G. (2008), Cancer control and the preservation of neurovascular tissue: how to meet competing goals during robotic radical prostatectomy., *British Journal of Urology International* **101**(8), 1013–1018.

Thompson, J. C., Wood, J. and Feuer, D. (2007), Prostate cancer: palliative care and pain relief., *British Medical Bulletin* **83**, 341–354.

Tonet, O., Focacci, F., Piccigallo, M., Cavallo, F., Uematsu, M., Megali, G. and Dario, P. (2006), Comparison of control modes of a hand-held robot for laparoscopic surgery, *in Medical Image Computing and Computer Assisted Intervention*, Vol. 4190 of *LNCS*, Springer, pp.429–436.

Tsai, A., Wells, W., Tempany, C., Grimson, E. and Willsky, A. (2004), Mutual information in coupled multi-shape model for medical image segmentation., *Medical Image Analysis* **8**(4), 429–445.

Tsai, R. Y. and Lenz, R. K. (June 1989), A New Technique for Fully Autonomous and Efficient 3D Robotics Hand/Eye Calibration, *IEEE Transactions on Robotics and Automation,* **5**(3), 345–358.

Tsaitgaist (2008), http://commons.wikimedia.org/wiki/File:Male_anatomy_en.svg, Electronic.

Ukimura, O. and Gill, I. S. (2009), Image-Fusion, Augmented Reality, and Predictive Surgical Navigation, *Urologic Clinics of North America* **36**(2), 115–123.

Ukimura, O., Magi-Galluzzi, C. and Gill, I. S. (2006), Real-time transrectal ultrasound guidance during laparoscopic radical prostatectomy: impact on surgical margins., *Journal of Urology* **175**(4), 1304–1310.

van de Kraats, E. B., Penney, G. P., van Walsum, T. and Niessen, W. J. (2005), Multispectral MR to X-ray registration of vertebral bodies by generating CT-like data., *in* J. S. Duncan and G. Gerig (eds.), *Medical Image Computing and Computer Assisted Intervention*, Vol. 3750 of *LNCS*, Springer, pp.911–918.

Varkarakis, I. M., Rais-Bahrami, S., Kavoussi, L. R. and Stoianovici, D. (2005), Robotic surgery and telesurgery in urology., *Urology* **65**(5), 840–846.

Venugopal, N., McCurdy, B., Hnatov, A. and Dubey, A. (2005), A feasibility study to investigate the use of thin-plate splines to account for prostate deformation., *Physics in Medicine & Biology* **50**(12), 2871–2885.

Visentini-Scarzanella, M., Mylonas, G. P., Stoyanov, D. and Yang, G. Z. (2009), i-BRUSH: A gaze-contigent virtual paintbruch for dense 3d reconstruction in robotic assisted surgery, *in Medical Image Computing and Computer Aided Surgery*, Vol. 5761 of *LNCS*, Springer, pp.353–360.

Warfield, S. K., Ferrant, M., Gallez, X., Nabavi, A., Jolesz, F. A. and Kikinis, R. (2000a), Real-Time Biomechanical Simulation of Volumetric Brain Deformation for Image Guided Neurosurgery, *in ACM/IEEE Conference on Supercomputing*, pp.23–23.

Warfield, S. K., Nabavi, A., Butz, T., Tuncali, K., Silverman, S. G., Black, P., Jolesz, F. A. and Kikinis, R. (2000b), Intraoperative segmentation and nonrigid registration for image guided therapy, *in Medical Image Computing and Computer Assisted Intervention*, Vol. 1935 of *LNCS*, Springer-Verlag, London, UK, pp.176–185.

Wengert, C., Reeff, M., Cattin, P. C. and Székely, G. (2006), Fully Automatic Endoscope Calibration for Intraoperative Use, *in Bildverarbeitung fr die Medizin*, Springer, pp.419–423.

West, J. B. and Maurer Jr, C. R. (2004), Designing Optically Tracked Instruments for Image-Guided Surgery, *IEEE Transaction on Medical Imaging* **23**(5), 533–545.

Wiles, D. A., Likholyot, A., Frantz, D. D. and Peters, T. M. (2008), A Statistical Model for Point-Based Target Registration Error With Anisotropic Fiducial Localizer Error, *IEEE Transactions on Medical Imaging* **27**(3), 378–390.

Wiles, D. A., Thompson, D. G. and Frantz, D. D. (2004), Accuracy assement and interpretation for optiical tracking systems, *in SPIE Medical Imaging, Visualisation, Image-Guided Procedures, Display*, Vol. 5367, pp.421–432.

Yao, J. and Taylor, R. H. (2003), Non-rigid registration and correspondence finding in medical image analysis using multiple-layer flexible mesh template matching., *International Journal of Pattern Recognition and Artificial Intelligence* **17**(7), 1145–1165.

Zhan, Y., Zhou, X. S. and Krishnan, A. (2008), An information-theoretic view of the scheduling problem in whole-body CAD, *in SPIE Medical Imaging*, Vol. 6915, pp.691515.1–691515.10.

Zhang, Z. (1999), Flexible camera calibration by viewing a plane from unknown orientations, *in The Proceedings of the Seventh IEEE International Conference on Computer Vision*, Vol. 1, pp.666–673.

Publications and Presentations

Thompson, S., Penney, G., Buie, D., Hawkes, D. and Dasgupta, P. (2008), Use of a CT statistical deformation model for multi-modal pelvic bone segmentation., *in Proceedings of the SPIE, Medical Imaging: Image Processing.* **6914**

Thompson, S., Penney, G., Dasgupta, P., and Hawkes, D. (2008), Image Guidance in Robot Assisted Radical Prostatectomy., *Presented at European Association of Urology 23rd Annual Congress.*

Thompson, S., Penney, G., Dasgupta, P., and Hawkes, D. (2008), Bone segmentation using a CT based statistical shape model for image guidance in robot assisted prostatectomy *Presented at British Association of Urological Surgeons Annual Meeting.*

Thompson, S., Penney, Carter, T., G., Dasgupta, P., and Hawkes, D. (2009), Accuracy Analysis of an Image Guided Robotic Urology Surgery System, *Presented at Satellite Workshop on Geometric Accuracy in Image Guided Intervention, MICCAI September 2009.*

Thompson, S., Penney, G., Dasgupta, P., and Hawkes, D. (2009), Image Guidance in Robot Assisted Prostatectomy *Presented at World Congress on Endourology.*

Appendix A

Details of Data used in Study

A.1 Shape Model Patient A

Adult male pelvis

Name	SM-A
Description	Adult Male Pelvis
Modality	CT
Size (Voxels)	512x512x100
Voxel Dimensions (mm)	0.783x0.783x3.200
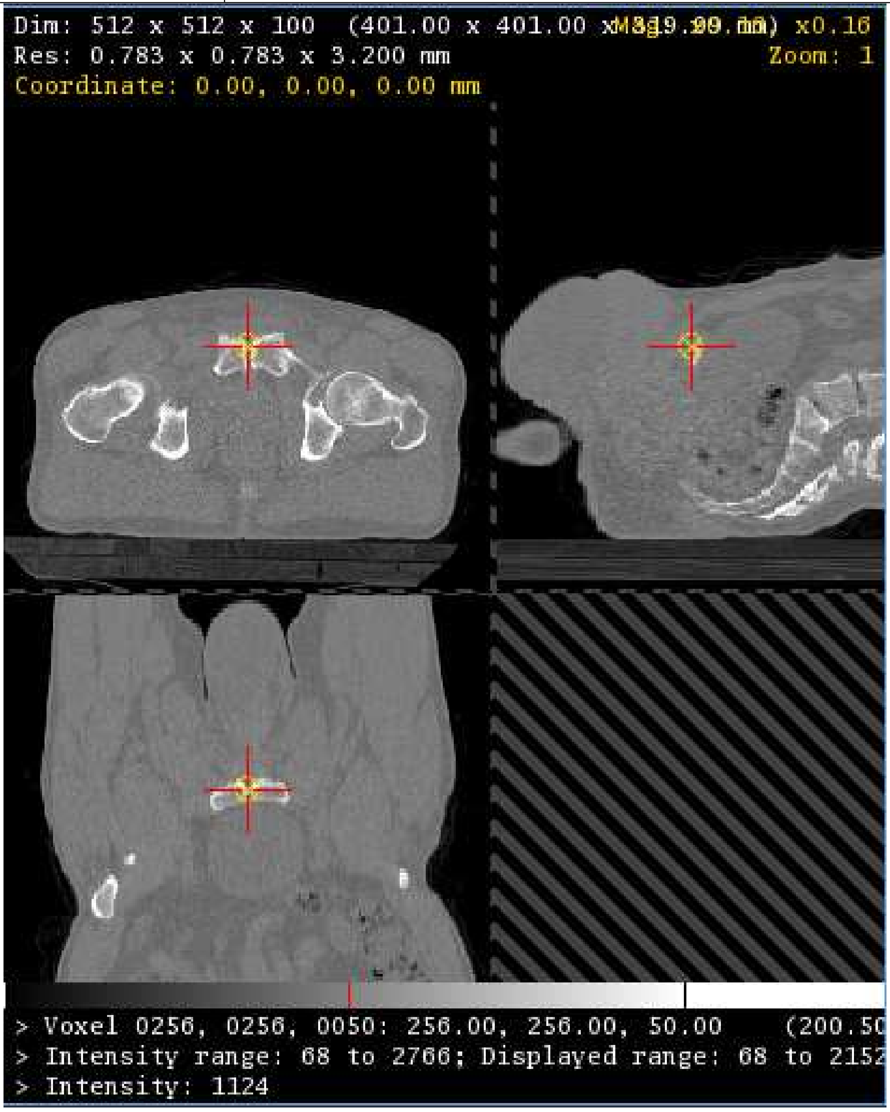	

```
Dim: 512 x 512 x 100  (401.00 x 401.00 x 319.99 mm) x0.16
Res: 0.783 x 0.783 x 3.200 mm                      Zoom: 1
Coordinate: 0.00, 0.00, 0.00 mm
```

```
> Voxel 0256, 0256, 0050: 256.00, 256.00, 50.00    (200.50
> Intensity range: 68 to 2766; Displayed range: 68 to 2152
> Intensity: 1124
```

A.2 Shape Model Patient B

Adult male pelvis

Name	SM-B
Description	Adult Male Pelvis
Modality	CT
Size (Voxels)	512x512x109
Voxel Dimensions (mm)	0.865x0.865x3.200
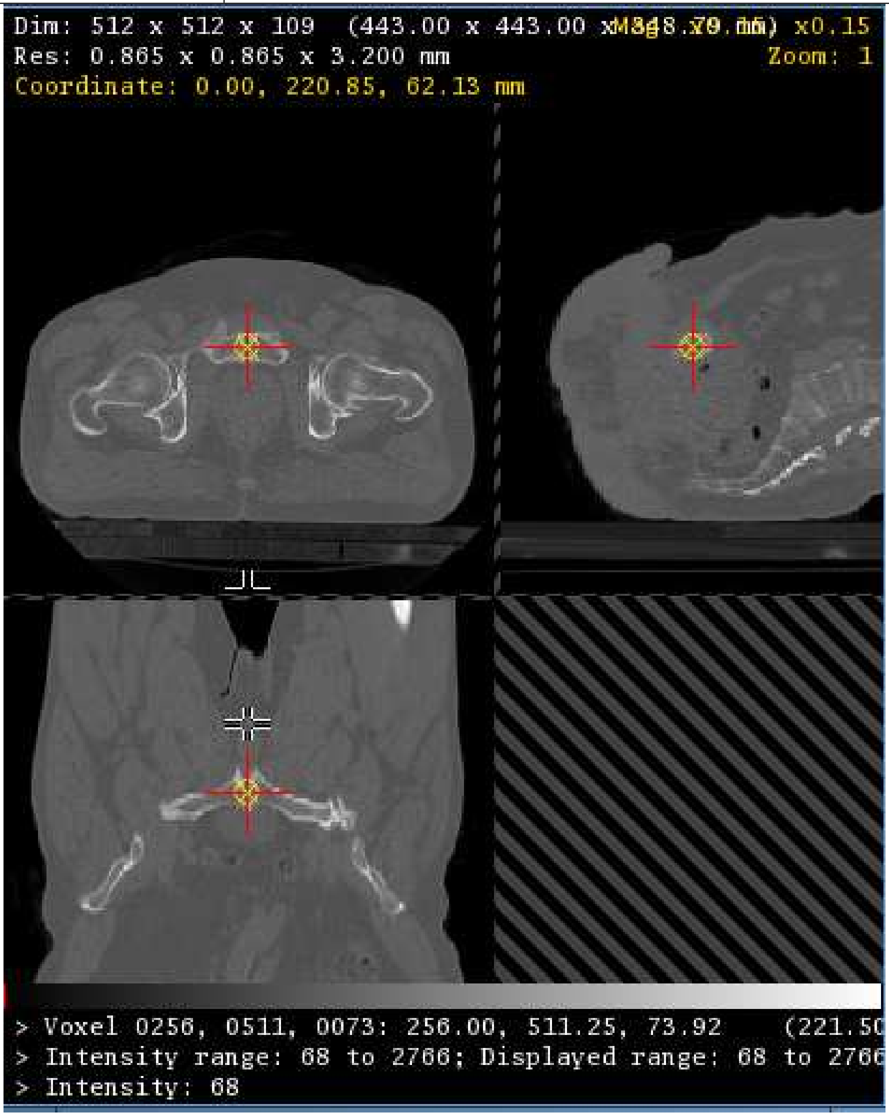	

```
Dim: 512 x 512 x 109  (443.00 x 443.00 xMag8.x0.mm) x0.15
Res: 0.865 x 0.865 x 3.200 mm                  Zoom: 1
Coordinate: 0.00, 220.85, 62.13 mm
```

```
> Voxel 0256, 0511, 0073: 256.00, 511.25, 73.92    (221.50
> Intensity range: 68 to 2766; Displayed range: 68 to 2766
> Intensity: 68
```

A.3 Shape Model Patient C

Adult male pelvis

Name	SM-C
Description	Adult Male Pelvis
Modality	CT
Size (Voxels)	512x512x100
Voxel Dimensions (mm)	0.762x0.762x3.200
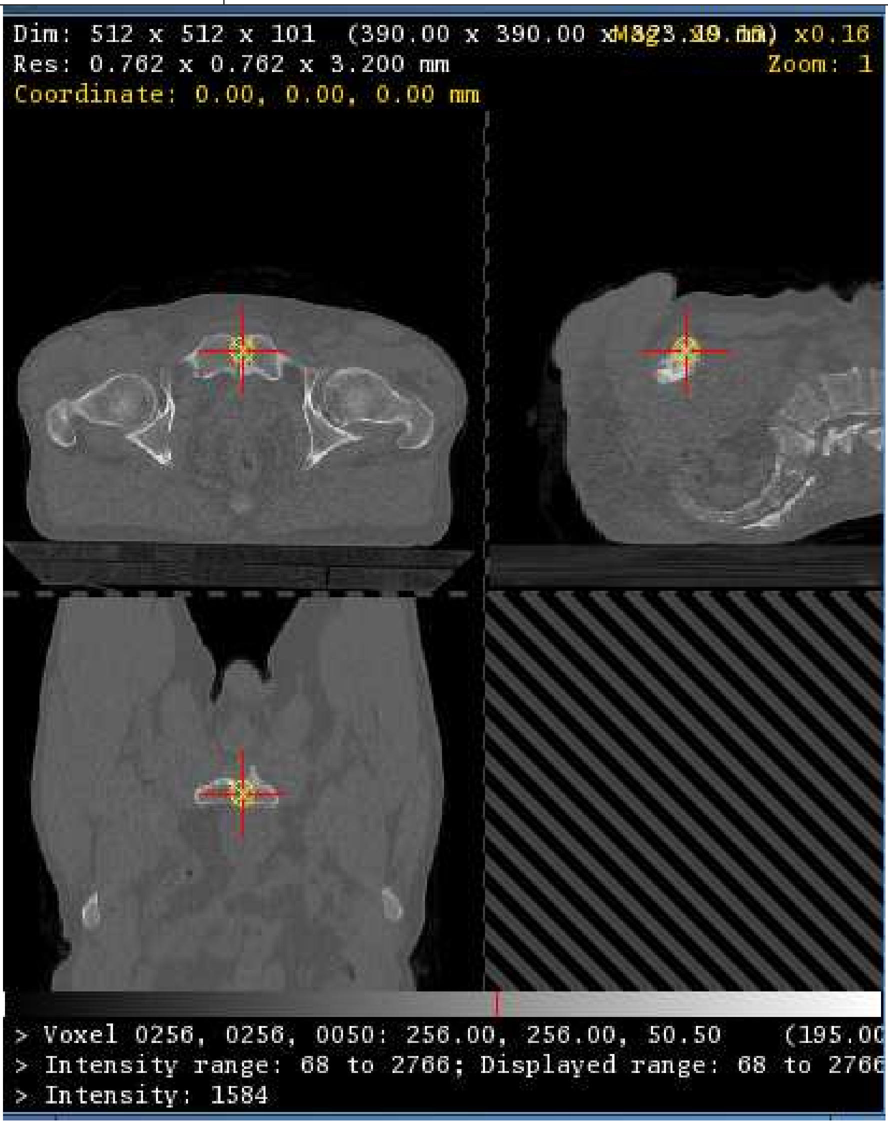	

```
Dim: 512 x 512 x 101  (390.00 x 390.00 xMag323.20.mm) x0.16
Res: 0.762 x 0.762 x 3.200 mm                        Zoom: 1
Coordinate: 0.00, 0.00, 0.00 mm
```

```
> Voxel 0256, 0256, 0050: 256.00, 256.00, 50.50    (195.00
> Intensity range: 68 to 2766; Displayed range: 68 to 2766
> Intensity: 1584
```

A.4 Shape Model Patient D

Adult male pelvis

Name	SM-D
Description	Adult Male Pelvis
Modality	CT
Size (Voxels)	512x512x103
Voxel Dimensions (mm)	0.779x0.779x3.200
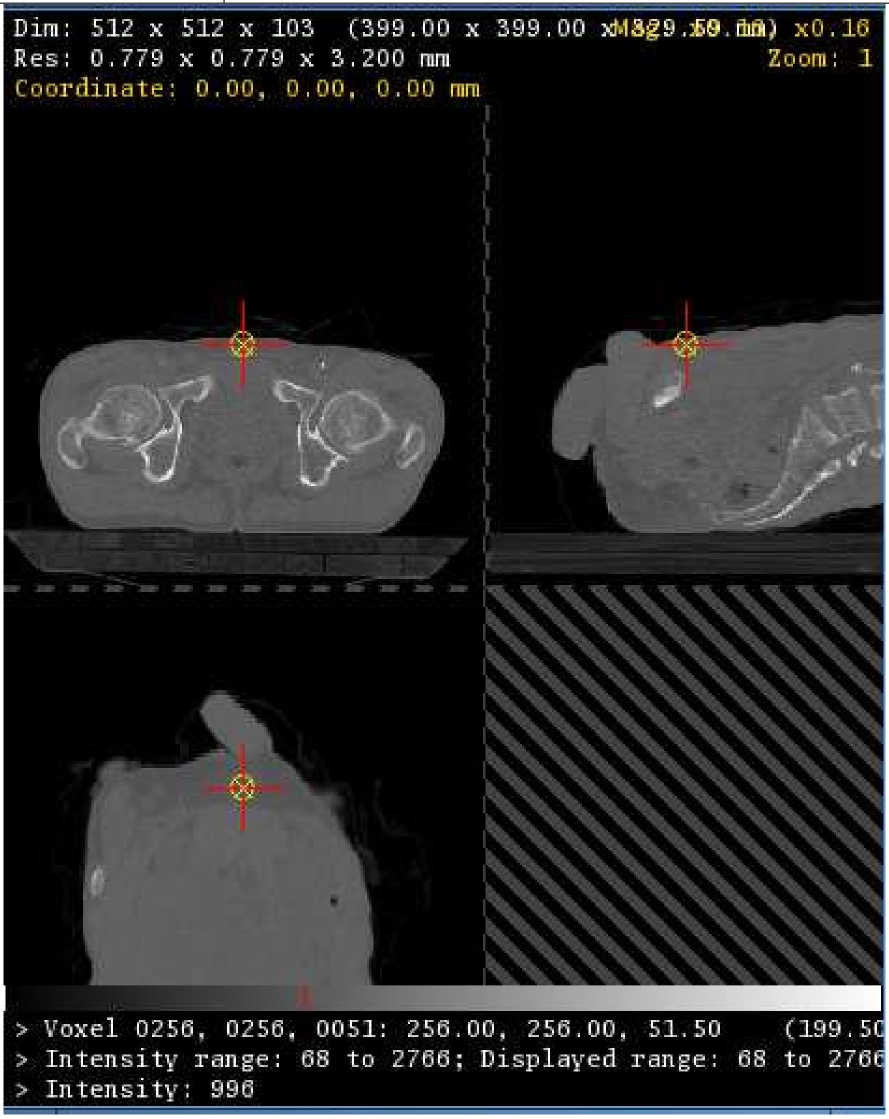	

```
Dim: 512 x 512 x 103  (399.00 x 399.00 x329.60 mm)  Mag: x0.16
Res: 0.779 x 0.779 x 3.200 mm                        Zoom: 1
Coordinate: 0.00, 0.00, 0.00 mm
```

```
> Voxel 0256, 0256, 0051: 256.00, 256.00, 51.50     (199.50
> Intensity range: 68 to 2766; Displayed range: 68 to 2766
> Intensity: 996
```

A.5 Shape Model Patient E

Adult male pelvis

Name	SM-E
Description	Adult Male Pelvis
Modality	CT
Size (Voxels)	512x512x173
Voxel Dimensions (mm)	0.738x0.738x2.000
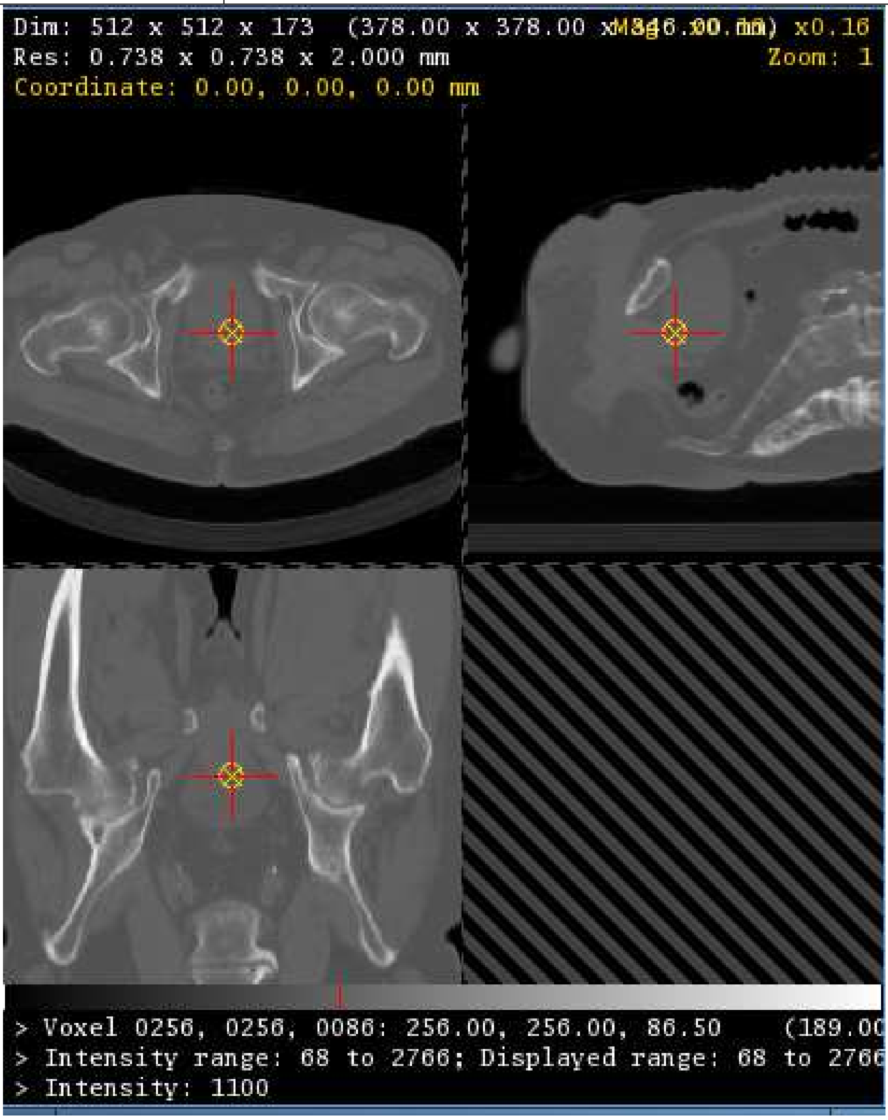	

```
Dim:  512 x 512 x 173  (378.00 x 378.00 x 346.00 mm) Mag: x0.16
Res:  0.738 x 0.738 x 2.000 mm                        Zoom: 1
Coordinate: 0.00, 0.00, 0.00 mm
```

```
> Voxel 0256, 0256, 0086: 256.00, 256.00, 86.50    (189.00
> Intensity range: 68 to 2766; Displayed range: 68 to 2766
> Intensity: 1100
```

A.6 Shape Model Patient F

Adult male pelvis

Name	SM-F
Description	Adult Male Pelvis
Modality	CT
Size (Voxels)	512x512x148
Voxel Dimensions (mm)	0.879x0.879x2.000
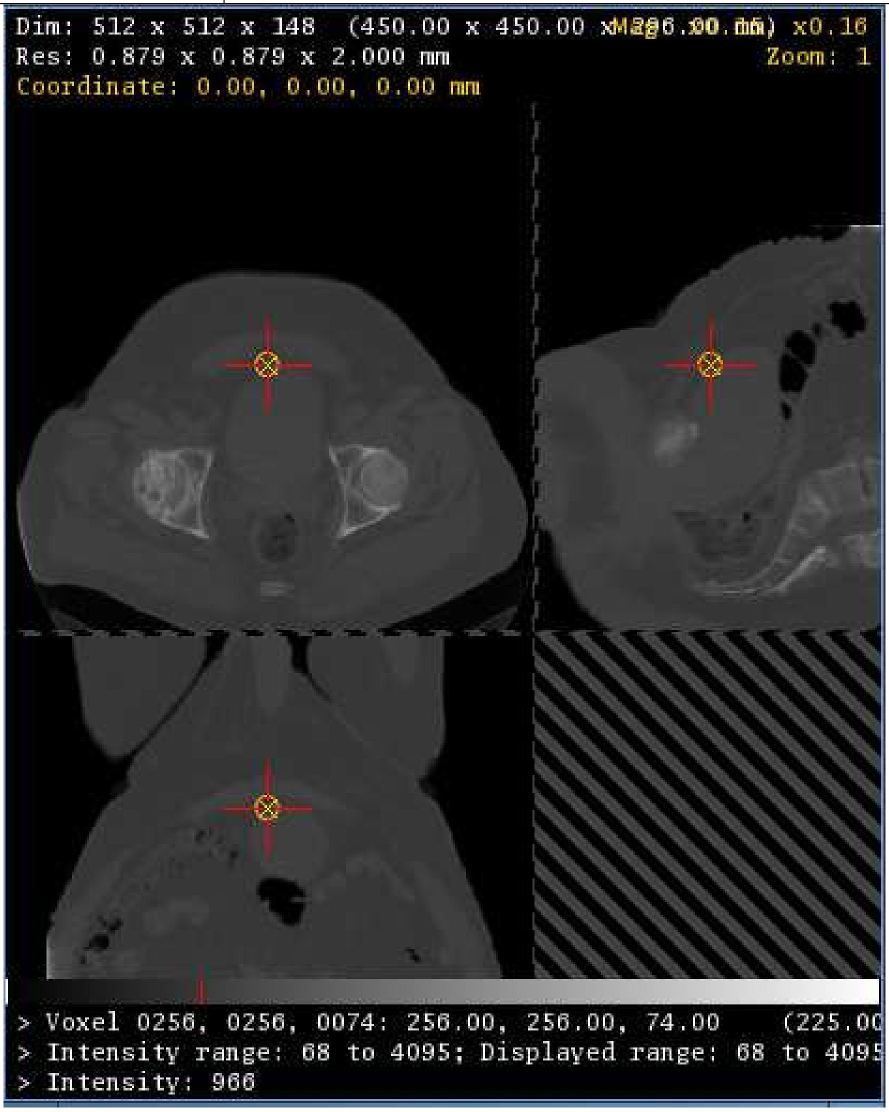	

```
Dim: 512 x 512 x 148  (450.00 x 450.00 x 296.00 mm)  Mag: x0.16
Res: 0.879 x 0.879 x 2.000 mm                          Zoom: 1
Coordinate: 0.00, 0.00, 0.00 mm
```

```
> Voxel 0256, 0256, 0074: 256.00, 256.00, 74.00    (225.0
> Intensity range: 68 to 4095; Displayed range: 68 to 409
> Intensity: 966
```

A.7 Shape Model Patient G

Adult male pelvis

Name	SM-G
Description	Adult Male Pelvis
Modality	CT
Size (Voxels)	512x512x158
Voxel Dimensions (mm)	0.732x0.732x2.000
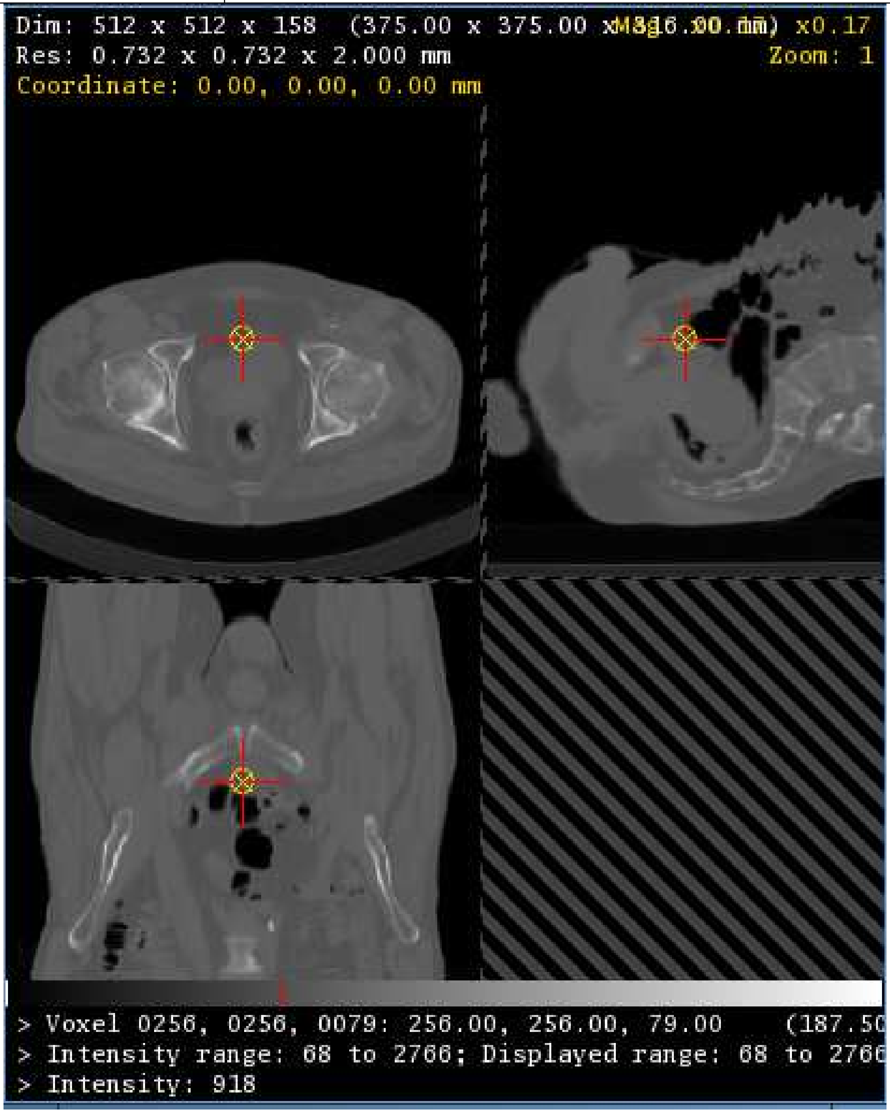	

```
Dim: 512 x 512 x 158  (375.00 x 375.00 x316.00 mm) Mag: x0.17
Res: 0.732 x 0.732 x 2.000 mm                          Zoom: 1
Coordinate: 0.00, 0.00, 0.00 mm
```

```
> Voxel 0256, 0256, 0079: 256.00, 256.00, 79.00    (187.5
> Intensity range: 68 to 2766; Displayed range: 68 to 2766
> Intensity: 918
```

A.8 Shape Model Patient H

Adult male pelvis

Name	SM-H
Description	Adult Male Pelvis
Modality	CT
Size (Voxels)	512x512x138
Voxel Dimensions (mm)	0.701x0.701x2.000
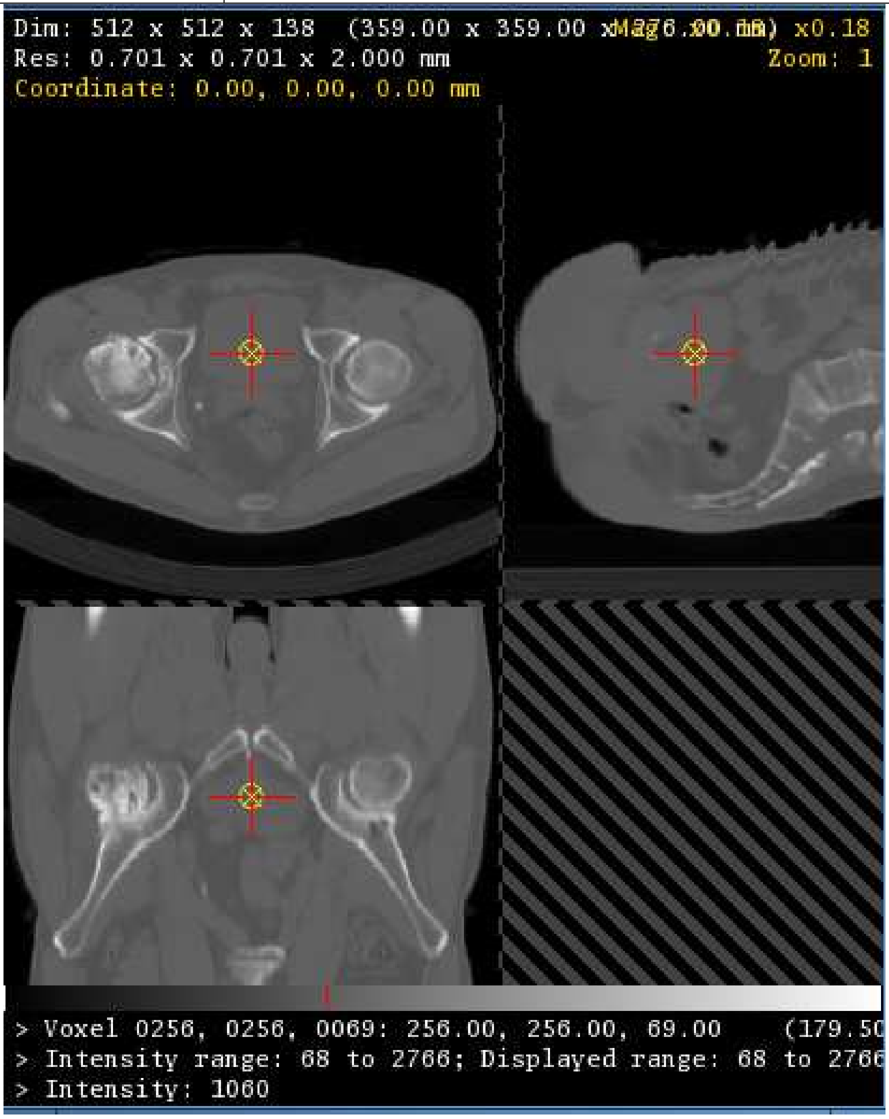 Dim: 512 x 512 x 138 (359.00 x 359.00 x 276.00 mm) Mag: x0.18 Res: 0.701 x 0.701 x 2.000 mm Zoom: 1 Coordinate: 0.00, 0.00, 0.00 mm > Voxel 0256, 0256, 0069: 256.00, 256.00, 69.00 (179.5 > Intensity range: 68 to 2766; Displayed range: 68 to 276 > Intensity: 1060	

A.9 Shape Model Patient I

Adult male pelvis

Name	SM-I
Description	Adult Male Pelvis
Modality	CT
Size (Voxels)	512x512x105
Voxel Dimensions (mm)	0.789x0.789x3.000
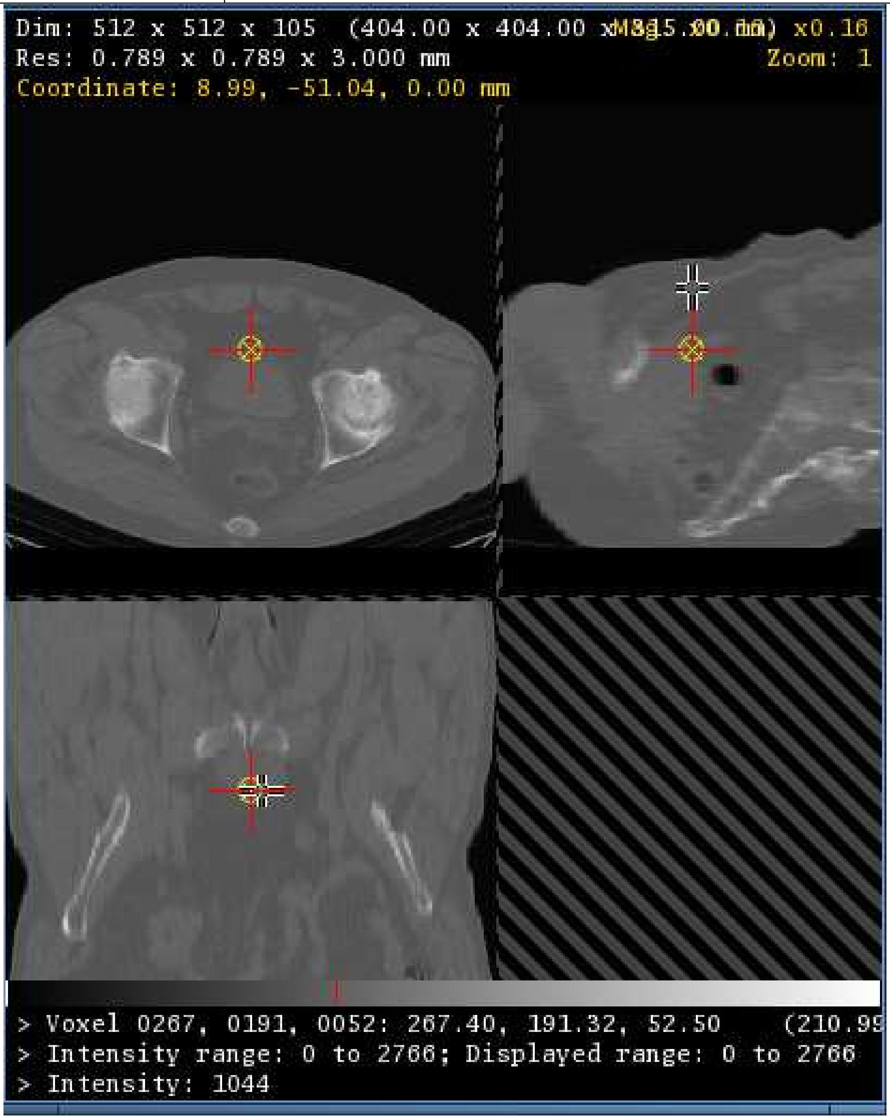	

```
Dim: 512 x 512 x 105  (404.00 x 404.00 xM3g5.00.mm) x0.16
Res: 0.789 x 0.789 x 3.000 mm                       Zoom: 1
Coordinate: 8.99, -51.04, 0.00 mm
```

```
> Voxel 0267, 0191, 0052: 267.40, 191.32, 52.50    (210.99
> Intensity range: 0 to 2766; Displayed range: 0 to 2766
> Intensity: 1044
```

A.10 Shape Model Patient J

Adult male pelvis

Name	SM-J
Description	Adult Male Pelvis
Modality	CT
Size (Voxels)	512x512x153
Voxel Dimensions (mm)	0.703x0.703x2.000
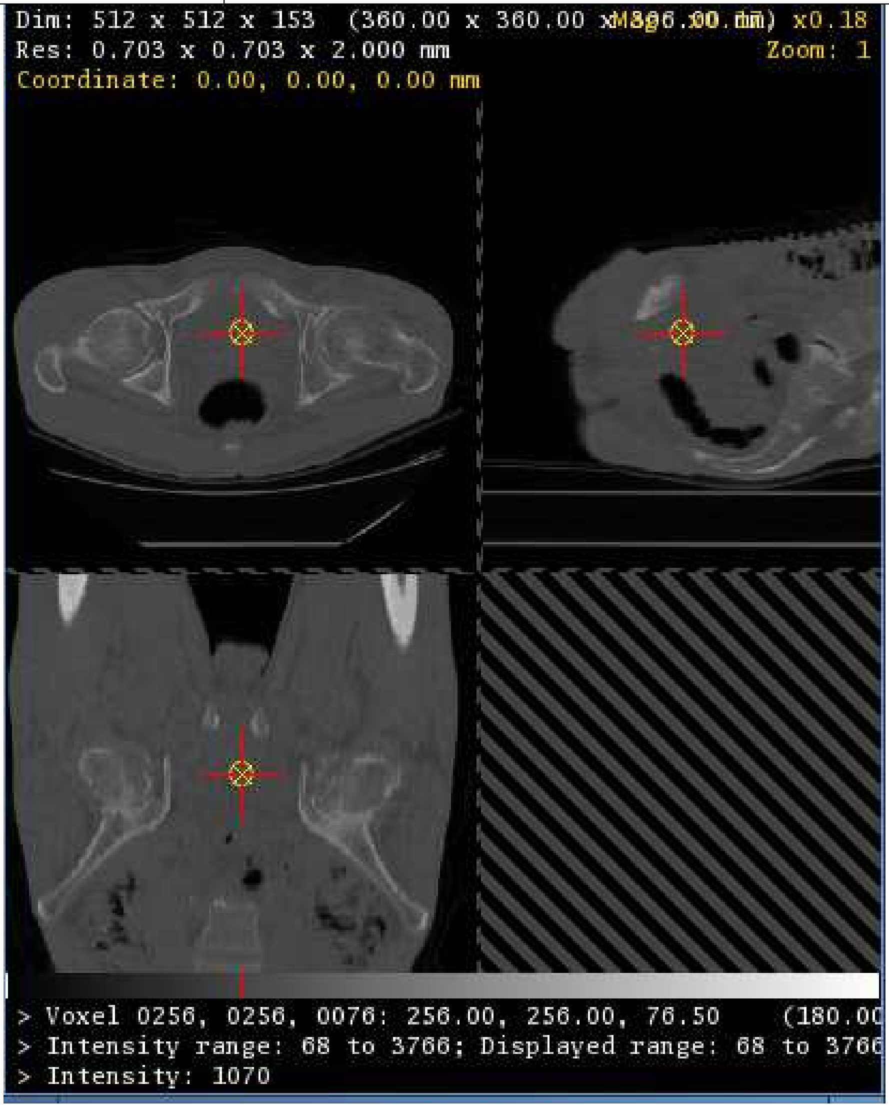	

```
Dim: 512 x 512 x 153  (360.00 x 360.00 x306.00 mm) x0.18
Res: 0.703 x 0.703 x 2.000 mm                     Zoom: 1
Coordinate: 0.00, 0.00, 0.00 mm
```

```
> Voxel 0256, 0256, 0076: 256.00, 256.00, 76.50    (180.0
> Intensity range: 68 to 3766; Displayed range: 68 to 376
> Intensity: 1070
```

A.11 Shape Model Patient K

Adult male pelvis

<table>
<tr><td>Name</td><td>SM-K</td></tr>
<tr><td>Description</td><td>Adult Male Pelvis</td></tr>
<tr><td>Modality</td><td>CT</td></tr>
<tr><td>Size (Voxels)</td><td>512x512x163</td></tr>
<tr><td>Voxel Dimensions (mm)</td><td>0.703x0.703x2.000</td></tr>
<tr><td colspan="2">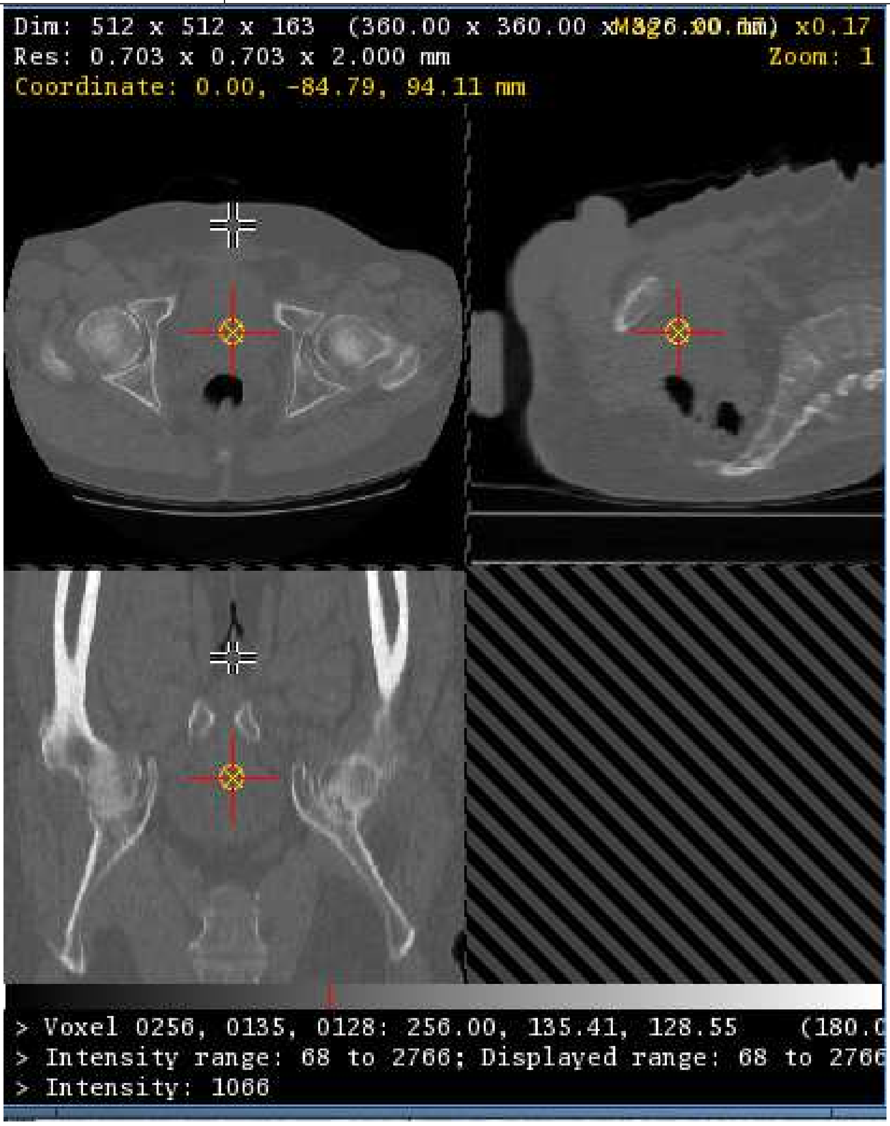</td></tr>
</table>

```
Dim: 512 x 512 x 163  (360.00 x 360.00 x326.00 mm) x0.17
Res: 0.703 x 0.703 x 2.000 mm                     Zoom: 1
Coordinate: 0.00, -84.79, 94.11 mm
```

```
> Voxel 0256, 0135, 0128: 256.00, 135.41, 128.55    (180.0
> Intensity range: 68 to 2766; Displayed range: 68 to 2766
> Intensity: 1066
```

A.12 Shape Model Patient L

Adult male pelvis

Name	SM-L
Description	Adult Male Pelvis
Modality	CT
Size (Voxels)	512x512x166
Voxel Dimensions (mm)	0.703x0.703x2.000
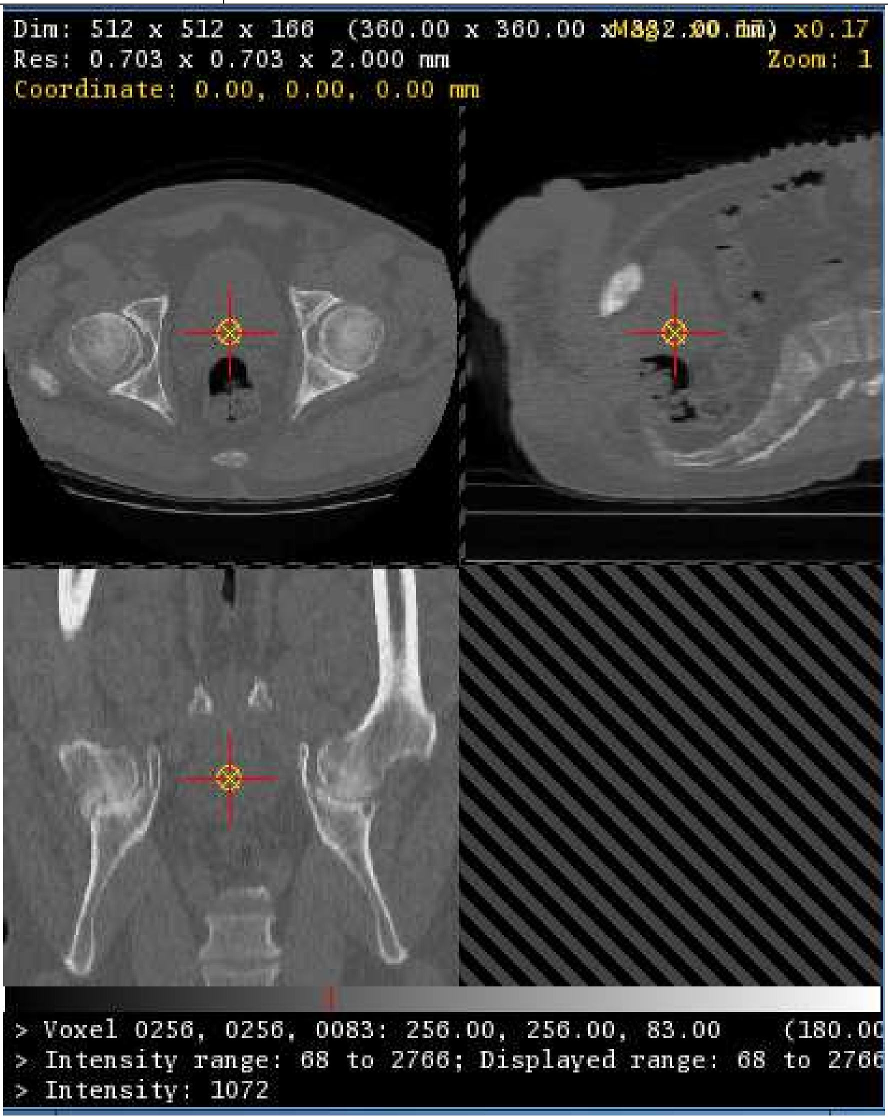	

```
Dim: 512 x 512 x 166  (360.00 x 360.00 x 332.00 mm) x0.17
Res: 0.703 x 0.703 x 2.000 mm                     Zoom: 1
Coordinate: 0.00, 0.00, 0.00 mm
```

```
> Voxel 0256, 0256, 0083: 256.00, 256.00, 83.00     (180.0
> Intensity range: 68 to 2766; Displayed range: 68 to 276
> Intensity: 1072
```

A.13 Shape Model Patient M

Adult male pelvis

Name	SM-M
Description	Adult Male Pelvis
Modality	CT
Size (Voxels)	512x512x160
Voxel Dimensions (mm)	0.723x0.723x2.000
	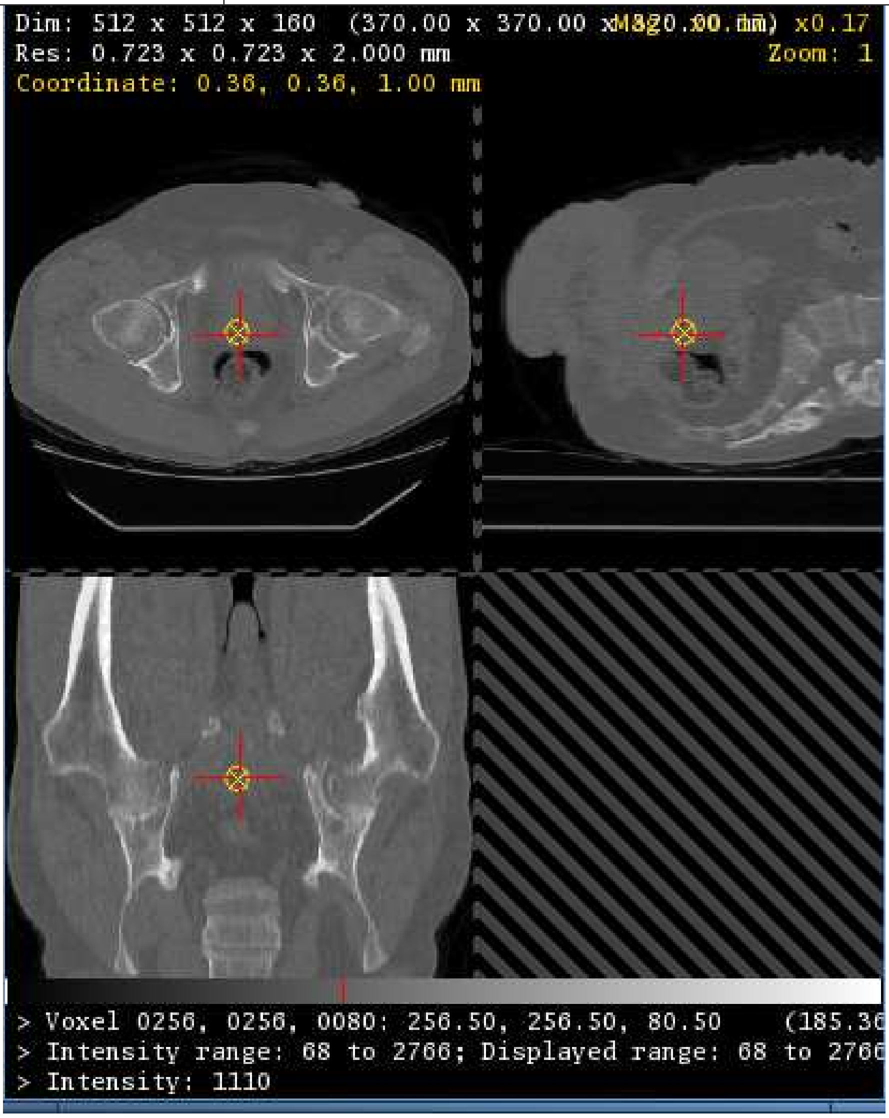 Dim: 512 x 512 x 160 (370.00 x 370.00 x 320.00 mm) x0.17 Res: 0.723 x 0.723 x 2.000 mm Zoom: 1 Coordinate: 0.36, 0.36, 1.00 mm > Voxel 0256, 0256, 0080: 256.50, 256.50, 80.50 (185.36 > Intensity range: 68 to 2766; Displayed range: 68 to 2766 > Intensity: 1110

A.14 Shape Model Patient N

Adult male pelvis

Name	SM-N
Description	Adult Male Pelvis
Modality	CT
Size (Voxels)	512x512x150
Voxel Dimensions (mm)	0.768x0.768x2.000
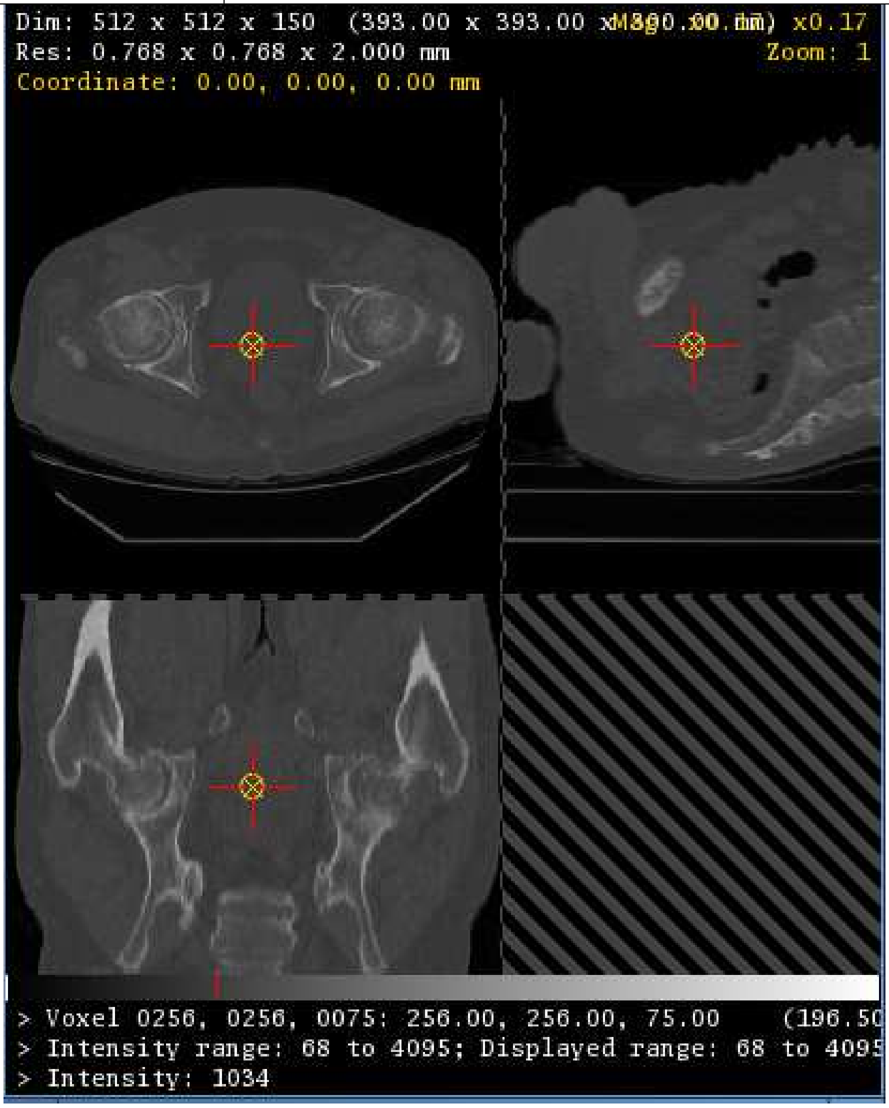	

```
Dim: 512 x 512 x 150  (393.00 x 393.00 x 300.00 mm) Mag x0.17
Res: 0.768 x 0.768 x 2.000 mm                         Zoom: 1
Coordinate: 0.00, 0.00, 0.00 mm
```

```
> Voxel 0256, 0256, 0075: 256.00, 256.00, 75.00     (196.5
> Intensity range: 68 to 4095; Displayed range: 68 to 4095
> Intensity: 1034
```

A.15 Shape Model Patient O

Adult male pelvis

Name	SM-O
Description	Adult Male Pelvis
Modality	CT
Size (Voxels)	512x512x158
Voxel Dimensions (mm)	0.703x0.703x2.000
	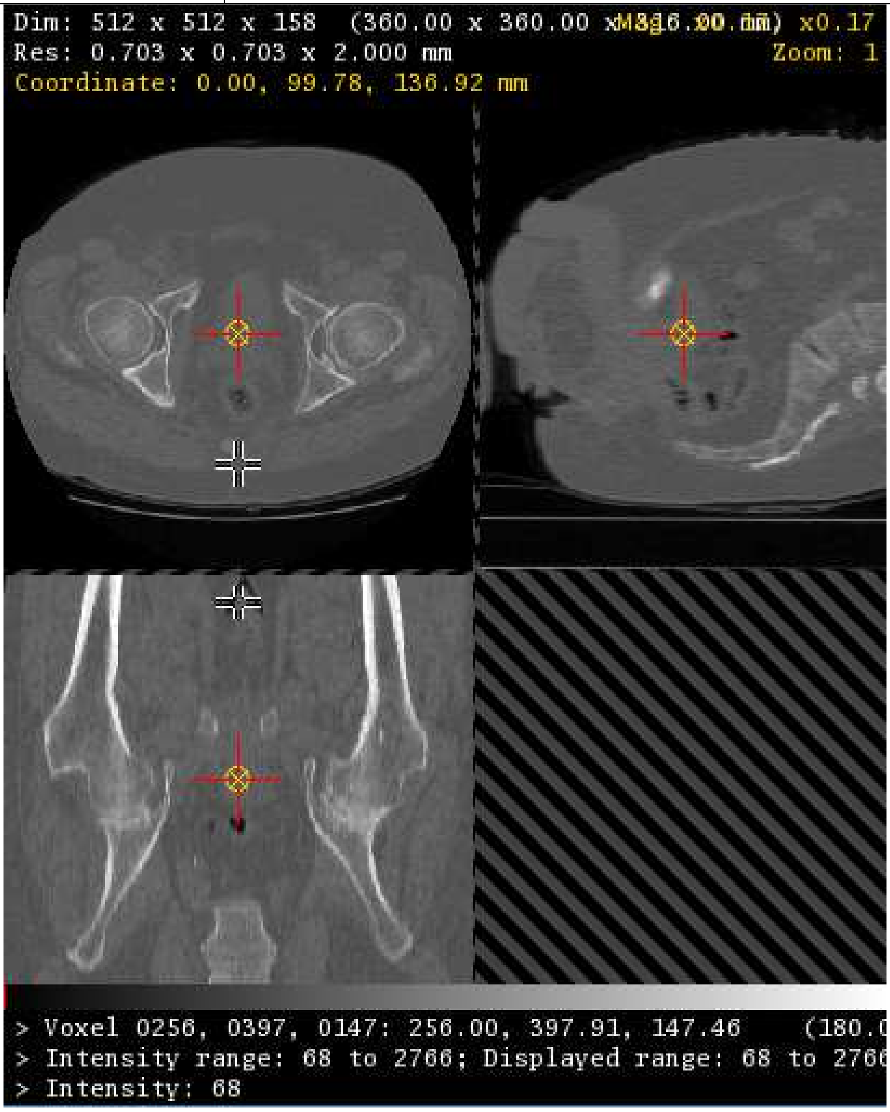

```
Dim: 512 x 512 x 158  (360.00 x 360.00 xM8g6.00.mm) x0.17
Res: 0.703 x 0.703 x 2.000 mm                        Zoom: 1
Coordinate: 0.00, 99.78, 136.92 mm
```

```
> Voxel 0256, 0397, 0147: 256.00, 397.91, 147.46    (180.0
> Intensity range: 68 to 2766; Displayed range: 68 to 2766
> Intensity: 68
```

A.16 Shape Model Patient P

Adult male pelvis

Name	SM-P
Description	Adult Male Pelvis
Modality	CT
Size (Voxels)	512x512x141
Voxel Dimensions (mm)	0.719x0.719x2.000
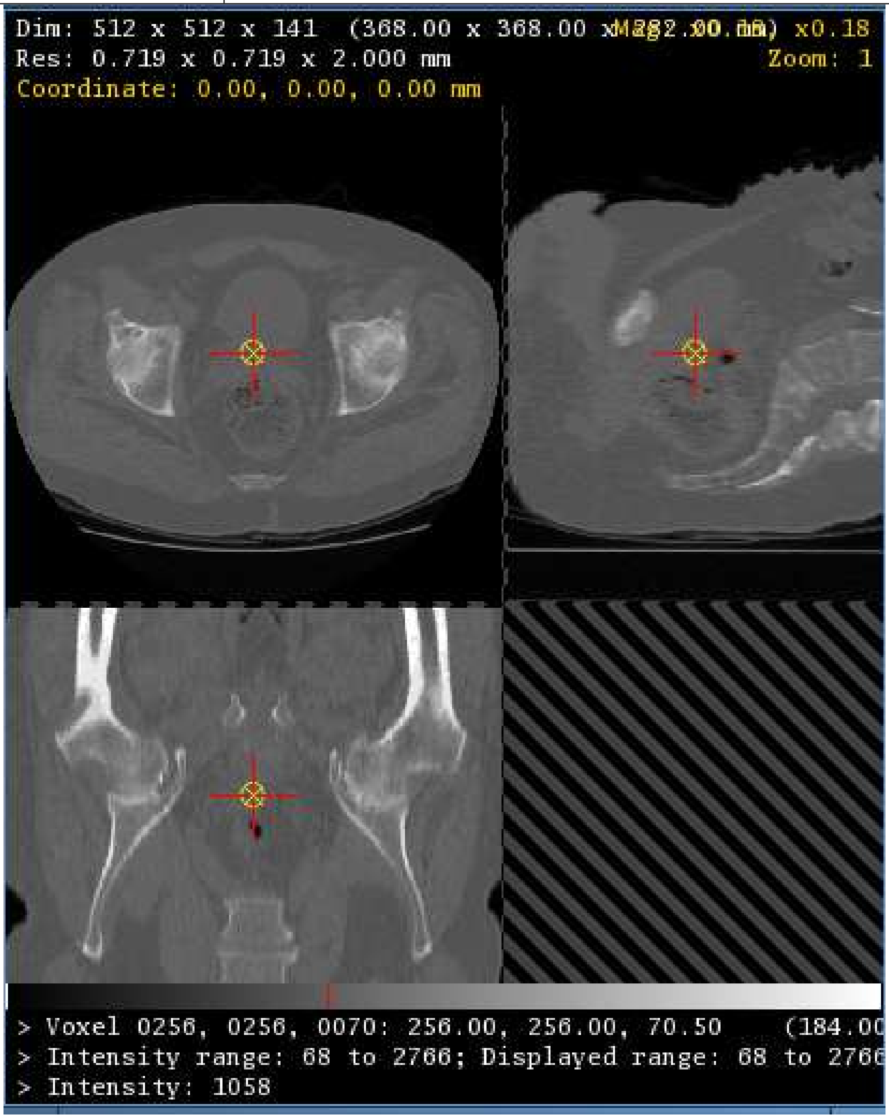 Dim: 512 x 512 x 141 (368.00 x 368.00 x 282.00 mm) Mag: x0.18 Res: 0.719 x 0.719 x 2.000 mm Zoom: 1 Coordinate: 0.00, 0.00, 0.00 mm > Voxel 0256, 0256, 0070: 256.00, 256.00, 70.50 (184.00 > Intensity range: 68 to 2766; Displayed range: 68 to 2766 > Intensity: 1058	

A.17 Shape Model Patient Q

Adult male pelvis

Name	SM-Q
Description	Adult Male Pelvis
Modality	CT
Size (Voxels)	512x512x139
Voxel Dimensions (mm)	0.746x0.746x2.000
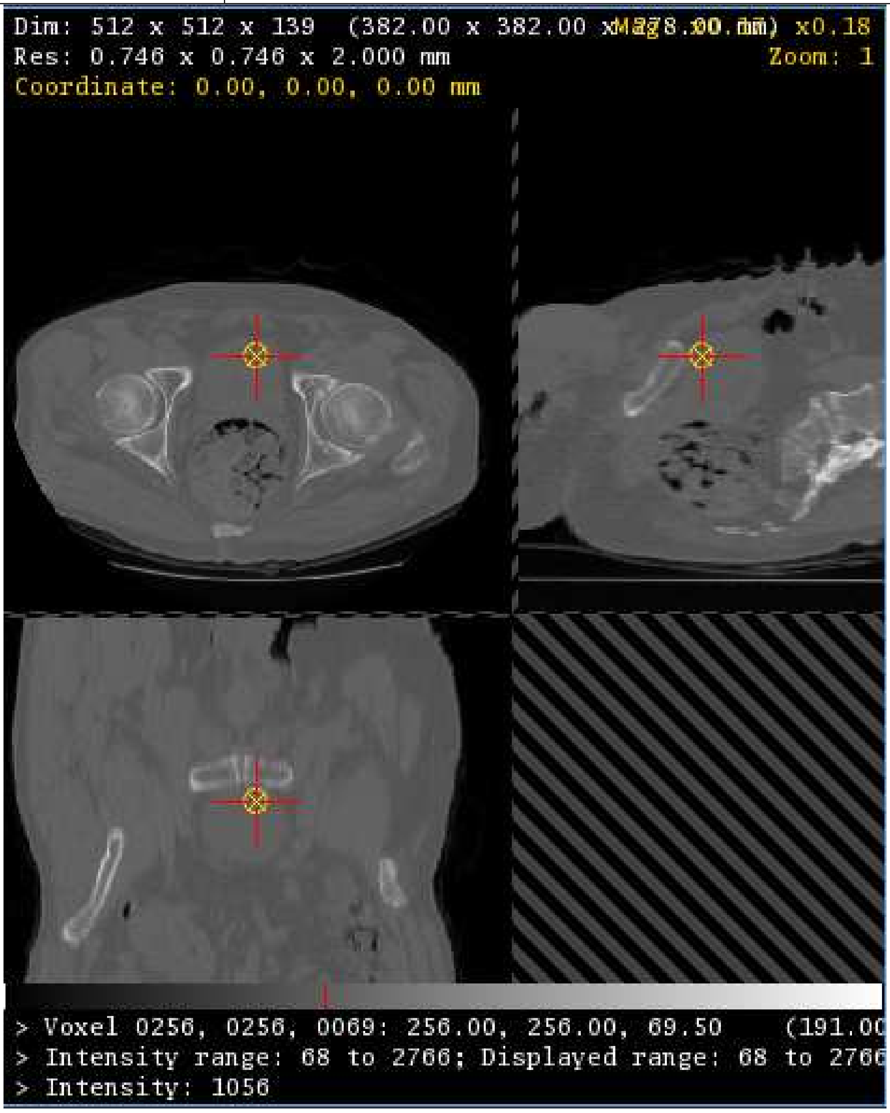	

```
Dim: 512 x 512 x 139  (382.00 x 382.00 x 278.00 mm)  Mag: x0.18
Res: 0.746 x 0.746 x 2.000 mm                             Zoom: 1
Coordinate: 0.00, 0.00, 0.00 mm
```

```
> Voxel 0256, 0256, 0069: 256.00, 256.00, 69.50    (191.0
> Intensity range: 68 to 2766; Displayed range: 68 to 276
> Intensity: 1056
```

A.18 Shape Model Patient R

Adult male pelvis

Name	SM-R
Description	Adult Male Pelvis
Modality	CT
Size (Voxels)	512x512x162
Voxel Dimensions (mm)	0.703x0.703x2.000

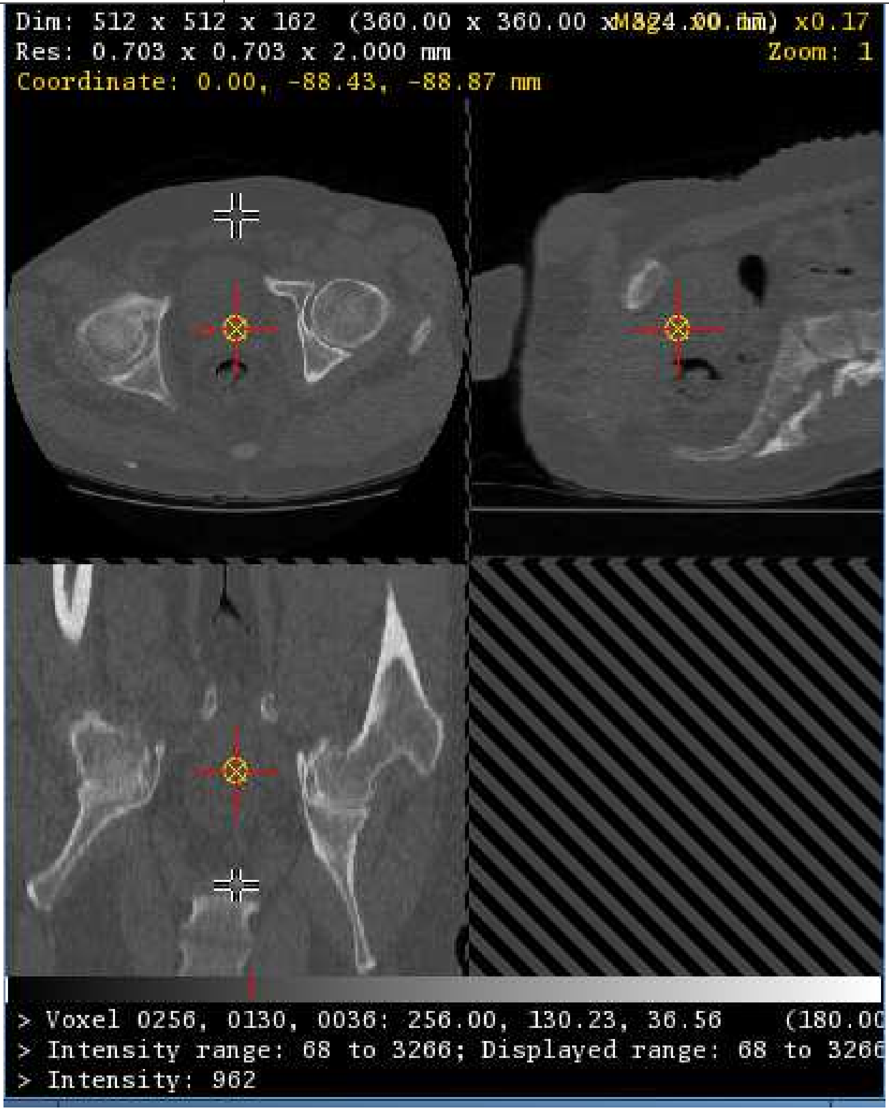

```
Dim: 512 x 512 x 162  (360.00 x 360.00 x 324.00 mm) x0.17
Res: 0.703 x 0.703 x 2.000 mm                      Zoom: 1
Coordinate: 0.00, -88.43, -88.87 mm
```

```
> Voxel 0256, 0130, 0036: 256.00, 130.23, 36.56    (180.00
> Intensity range: 68 to 3266; Displayed range: 68 to 3266
> Intensity: 962
```

A.19 Shape Model Patient S

Adult male pelvis

Name	SM-S
Description	Adult Male Pelvis
Modality	CT
Size (Voxels)	512x512x150
Voxel Dimensions (mm)	0.738x0.738x2.000
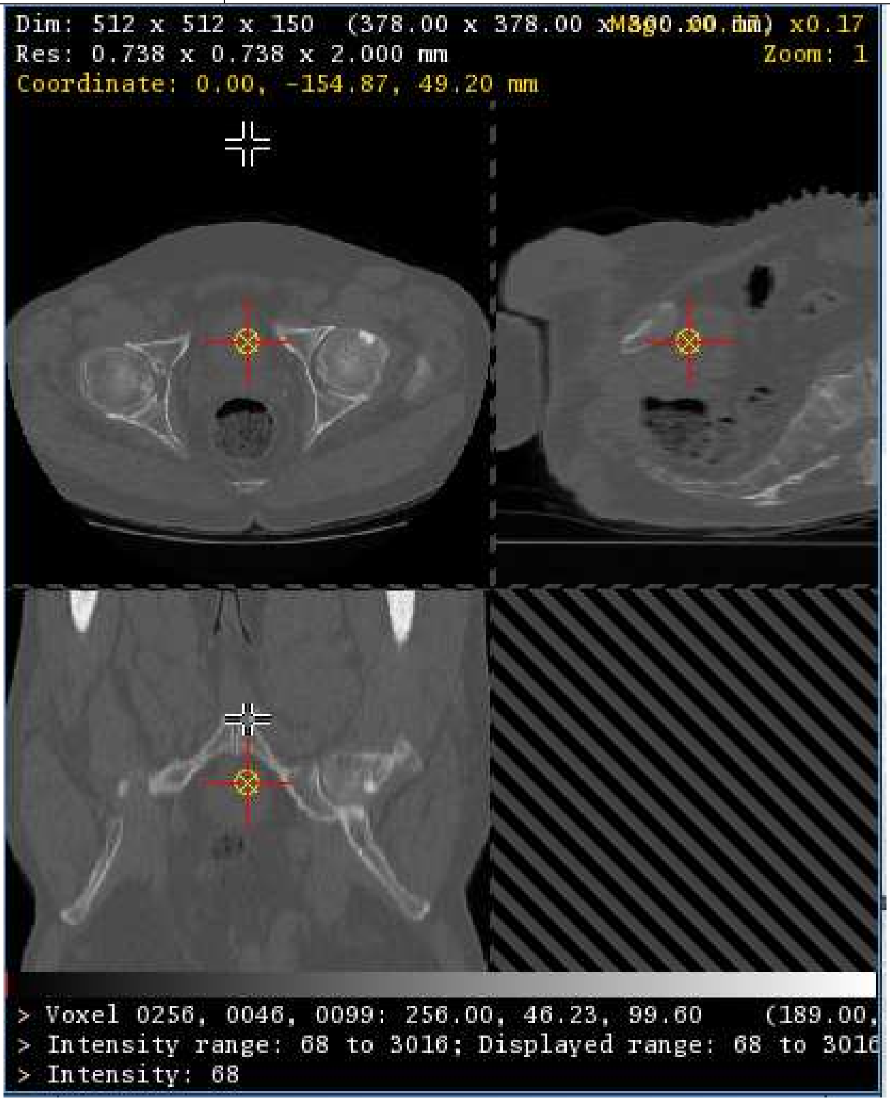	

```
Dim: 512 x 512 x 150  (378.00 x 378.00 x 300.00 mm) x0.17
Res: 0.738 x 0.738 x 2.000 mm                       Zoom: 1
Coordinate: 0.00, -154.87, 49.20 mm
```

```
> Voxel 0256, 0046, 0099: 256.00, 46.23, 99.60    (189.00,
> Intensity range: 68 to 3016; Displayed range: 68 to 3016
> Intensity: 68
```

A.20 Shape Model Patient T

Adult male pelvis

Name	SM-T
Description	Adult Male Pelvis
Modality	CT
Size (Voxels)	512x512x154
Voxel Dimensions (mm)	0.703x0.703x2.000
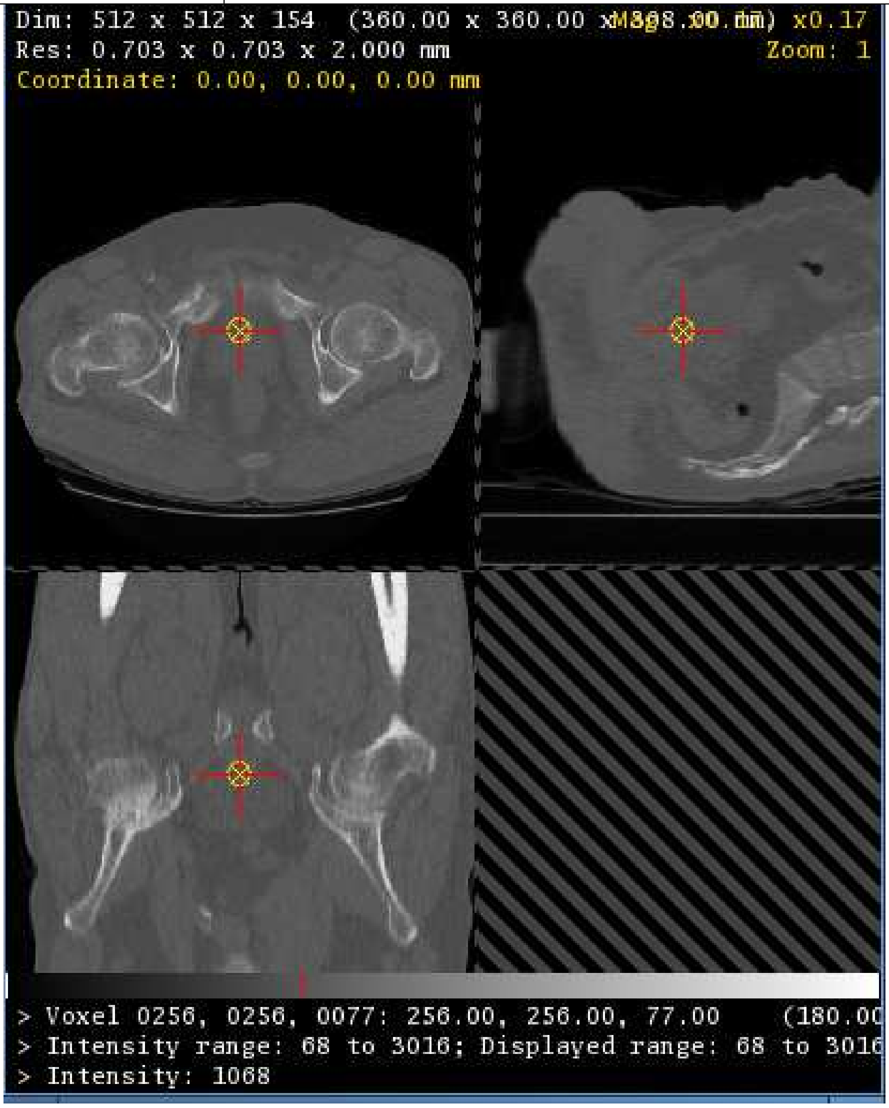	

```
Dim: 512 x 512 x 154  (360.00 x 360.00 x 308.00 mm)  Mag: x0.17
Res: 0.703 x 0.703 x 2.000 mm                          Zoom: 1
Coordinate: 0.00, 0.00, 0.00 mm
```

```
> Voxel 0256, 0256, 0077: 256.00, 256.00, 77.00     (180.0
> Intensity range: 68 to 3016; Displayed range: 68 to 3016
> Intensity: 1068
```

A.21 Shape Model Patient U

Adult male pelvis

Name	SM-U
Description	Adult Male Pelvis
Modality	CT
Size (Voxels)	512x512x150
Voxel Dimensions (mm)	0.738x0.738x2.000

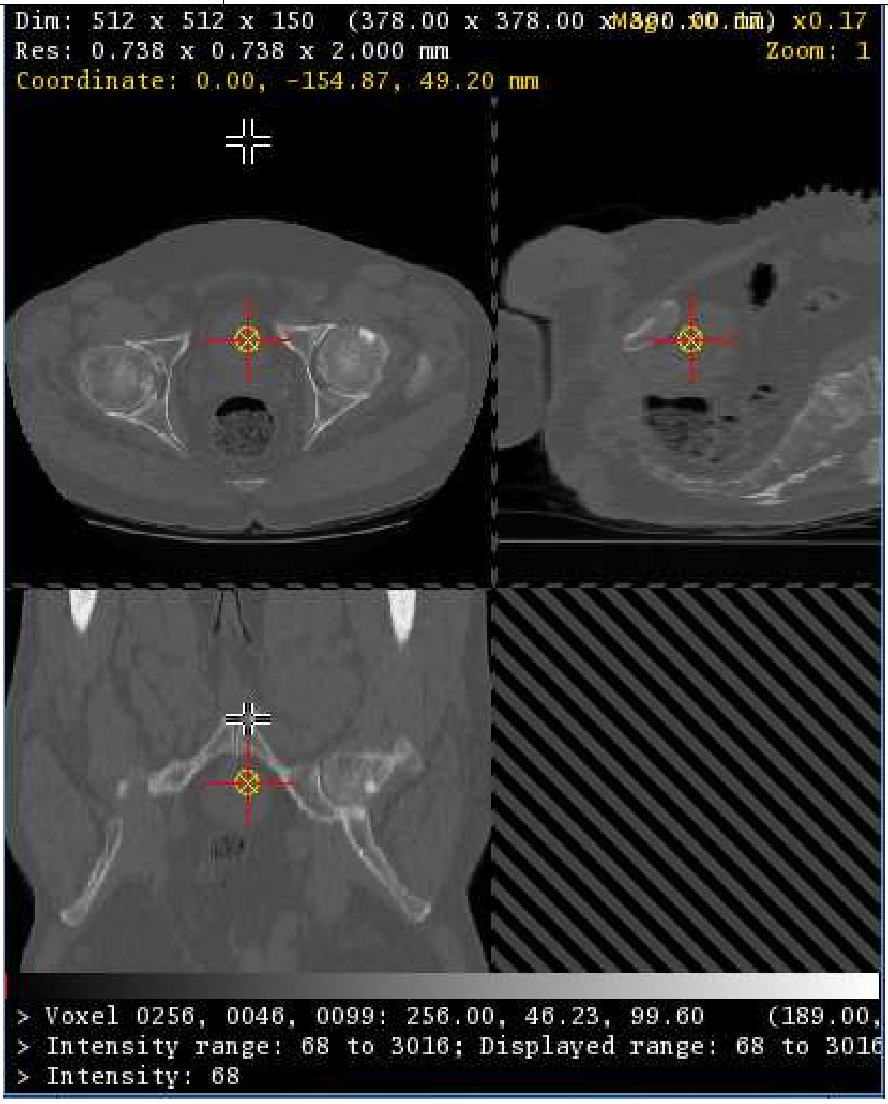

```
Dim: 512 x 512 x 150  (378.00 x 378.00 x 300.00 mm)  Mag: x0.17
Res: 0.738 x 0.738 x 2.000 mm                         Zoom: 1
Coordinate: 0.00, -154.87, 49.20 mm
```

```
> Voxel 0256, 0046, 0099: 256.00, 46.23, 99.60    (189.00,
> Intensity range: 68 to 3016; Displayed range: 68 to 301
> Intensity: 68
```

A.22 Prostatectomy Patient 01

Patient for Radical Prostatectomy

Image Description	Size	Resolution	Slice Spacing	Image date (Days Before Surgery)
t2 Sagittal	512:512:13	0.49:0.49:5	6	127
t2 Transverse	320:320:30	1.19:1.19:5	6.25	127
t1 Sagittal (Oblique)	512:512:28	0.72:0.72:5	5.5	127
t1 Transverse	320:320:23	0.63:0.63:3	3.45	127
t2 Transverse	320:320:23	0.59:0.59:3.2	3.84	127
t2 Coronal	512:512:23	0.59:0.59:5	5.75	127

Name	Patient-01 Anatomy
Description	
Modality	T2 MRI
Size (Voxels)	320x320x23
Voxel Dimensions (mm)	0.594x0.0.594x3.840

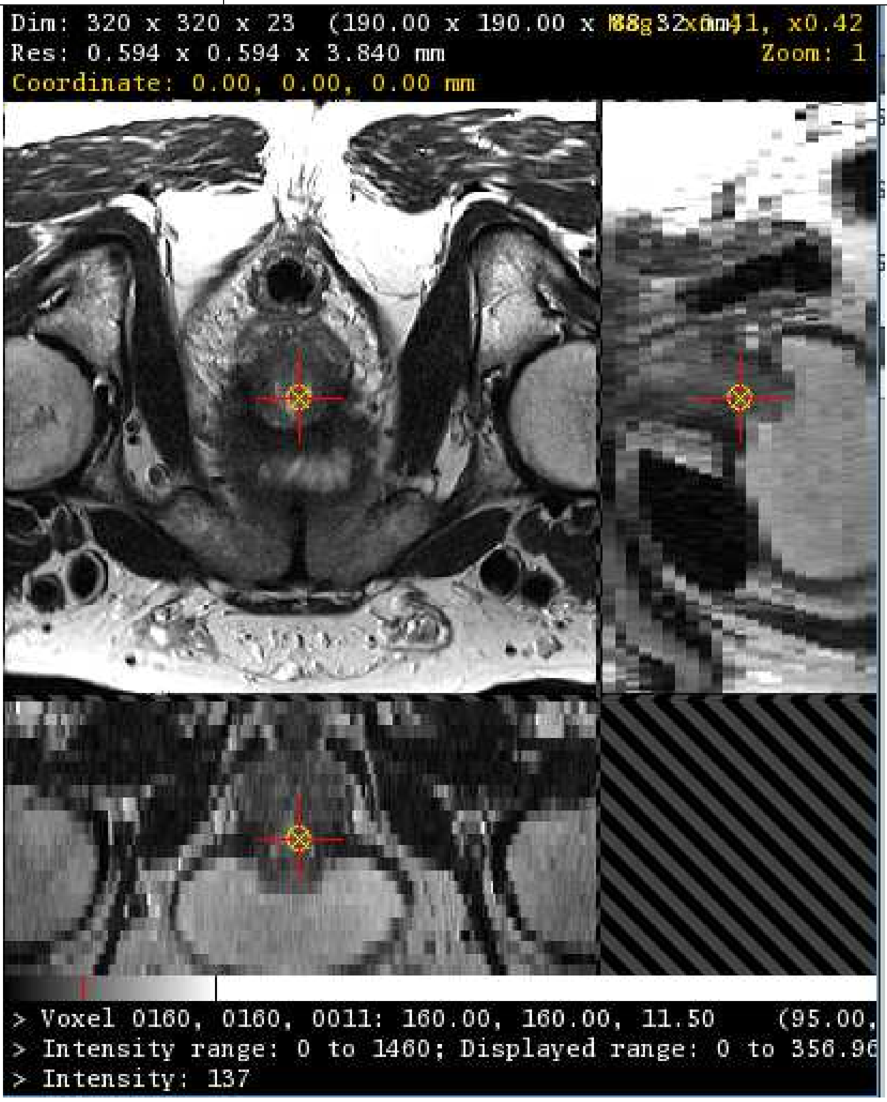

```
Dim: 320 x 320 x 23  (190.00 x 190.00 x M8g32xmm41, x0.42
Res: 0.594 x 0.594 x 3.840 mm                    Zoom: 1
Coordinate: 0.00, 0.00, 0.00 mm
```

```
> Voxel 0160, 0160, 0011: 160.00, 160.00, 11.50    (95.00,
> Intensity range: 0 to 1460; Displayed range: 0 to 356.96
> Intensity: 137
```

A.23 Prostatectomy Patient 02

Patient for Radical Prostatectomy

Image Description	Size	Resolution	Slice Spacing	Image date (Days Before Surgery)
t2 Sagittal	512:512:25	0.78:0.78:10	15	21
Diffusion Weighted	128:128:125	1.95:1.95:6	7.2	21
t1 Transverse	512:512:30	0.49:0.49:6	7.2	21
t2 Axial (Transverse)	384:384:24	0.65:0.65:3	3.9	21
t2 Coronal	384:384:20	0.52:0.52:3	3.6	21

Name	Patient-02 Anatomy
Description	
Modality	T2 MRI
Size (Voxels)	384x384x24
Voxel Dimensions (mm)	0.651x0.651x3.900

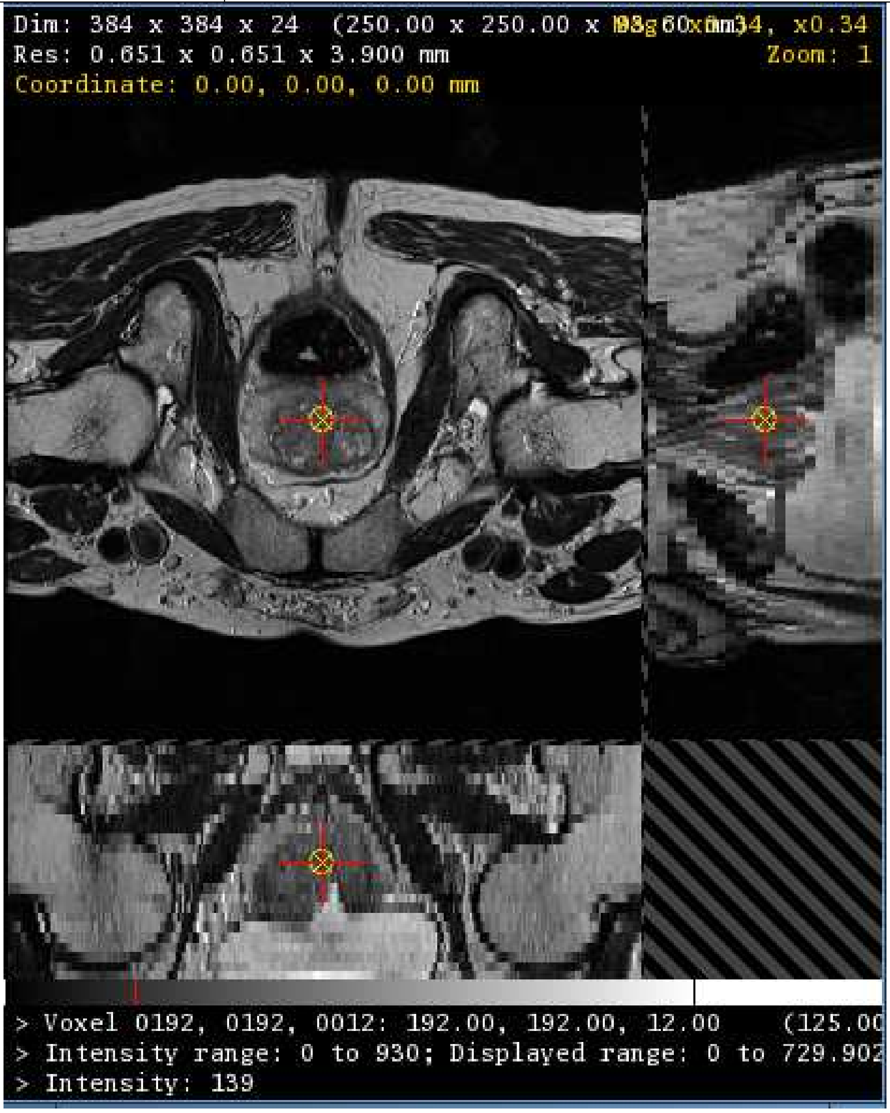

```
Dim: 384 x 384 x 24  (250.00 x 250.00 x M8g60xmm34, x0.34
Res: 0.651 x 0.651 x 3.900 mm                     Zoom: 1
Coordinate: 0.00, 0.00, 0.00 mm
```

```
> Voxel 0192, 0192, 0012: 192.00, 192.00, 12.00    (125.00
> Intensity range: 0 to 930; Displayed range: 0 to 729.902
> Intensity: 139
```

A.24 Prostatectomy Patient 03

Patient for Radical Prostatectomy

Image Description	Size	Resolution	Slice Spacing	Image date (Days Before Surgery)
t2 Sagittal	512:512:13	0.49:0.49:5	6	106
t2 Transverse	320:320:30	1.19:1.19:5	6.25	106
t2 Coronal	512:512:26	0.59:0.59:5	5.75	106
t2 Transverse	320:320:23	0.59:0.59:3.2	3.84	106
t1 Transverse	320:320:23	0.63:0.63:3	3.45	106
t1 Sagittal (Oblique)	512:512:28	0.72:0.72:5	5.5	106

Name	Patient-03 Anatomy
Description	
Modality	T2 MRI
Size (Voxels)	320x320x23
Voxel Dimensions (mm)	0.59x0.59x3.8
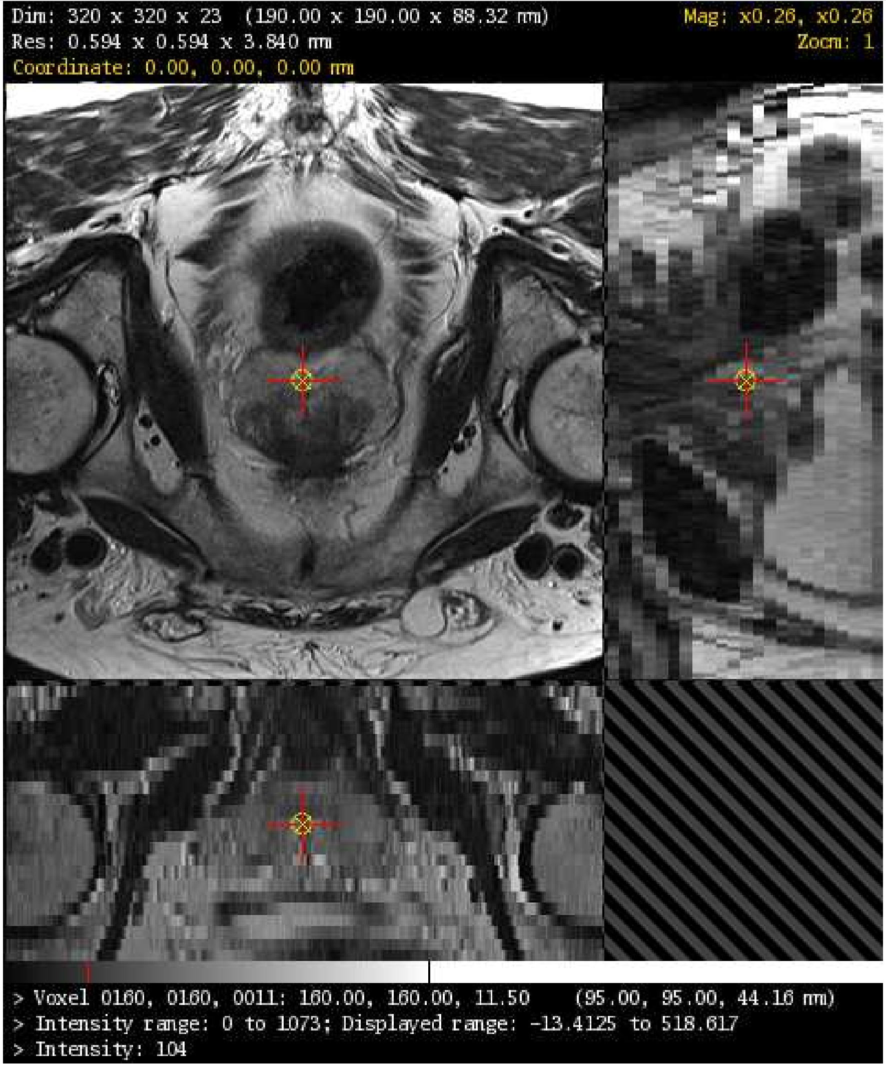	

```
Dim: 320 x 320 x 23  (190.00 x 190.00 x 88.32 mm)          Mag: x0.26, x0.26
Res: 0.594 x 0.594 x 3.840 mm                                       Zoom: 1
Coordinate: 0.00, 0.00, 0.00 mm
> Voxel 0160, 0160, 0011: 160.00, 160.00, 11.50    (95.00, 95.00, 44.16 mm)
> Intensity range: 0 to 1073; Displayed range: -13.4125 to 518.617
> Intensity: 104
```

A.25 Prostatectomy Patient 04

Patient for Radical Prostatectomy

Image Description	Size	Resolution	Slice Spacing	Image date (Days Before Surgery)
Sagittal	256:256:23	0.70:0.70:3	3	177
t1 Axial	512:512:39	0.68:0.68:5	5.5	177
t1 Coronal	512:512:23	0.74:0.74:5	6.05	177
t2 Coronal	256:256:23	0.70:0.70:3	3.3	177
t2 Transverse	256:256:23	0.63:0.63:3	3.3	177

Name	Patient-04 Anatomy
Description	
Modality	T2 MRI
Size (Voxels)	256x256x23
Voxel Dimensions (mm)	0.63x0.63x3.3

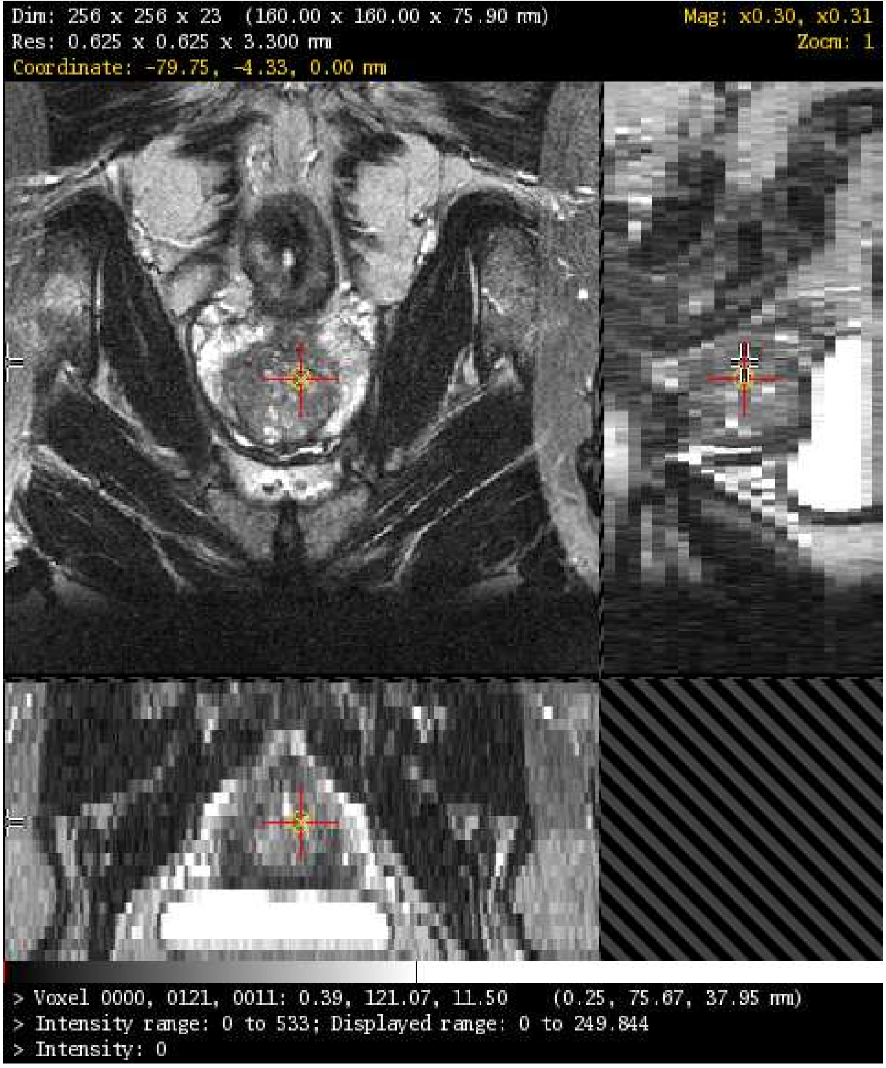

```
Dim: 256 x 256 x 23  (160.00 x 160.00 x 75.90 mm)          Mag: x0.30, x0.31
Res: 0.625 x 0.625 x 3.300 mm                                       Zoom: 1
Coordinate: -79.75, -4.33, 0.00 mm
```

```
> Voxel 0000, 0121, 0011: 0.39, 121.07, 11.50    (0.25, 75.67, 37.95 mm)
> Intensity range: 0 to 533; Displayed range: 0 to 249.844
> Intensity: 0
```

A.26 Prostatectomy Patient 05

Patient for Radical Prostatectomy

T1 Transverse	512:512:32	0.78:0.78:6	7.2	101
T2 Sagittal	512:512:30	0.39:0.39:3.5	3.84	101
T2 Transverse	512:512:30	0.35:0.35:3.5	3.84	101
T2 Coronal	512:512:50	0.35:0.35:3.5	3.84	101

Name	Patient-05 Anatomy
Description	
Modality	T2 MRI
Size (Voxels)	512x512x30
Voxel Dimensions (mm)	0.35x0.35x3.8
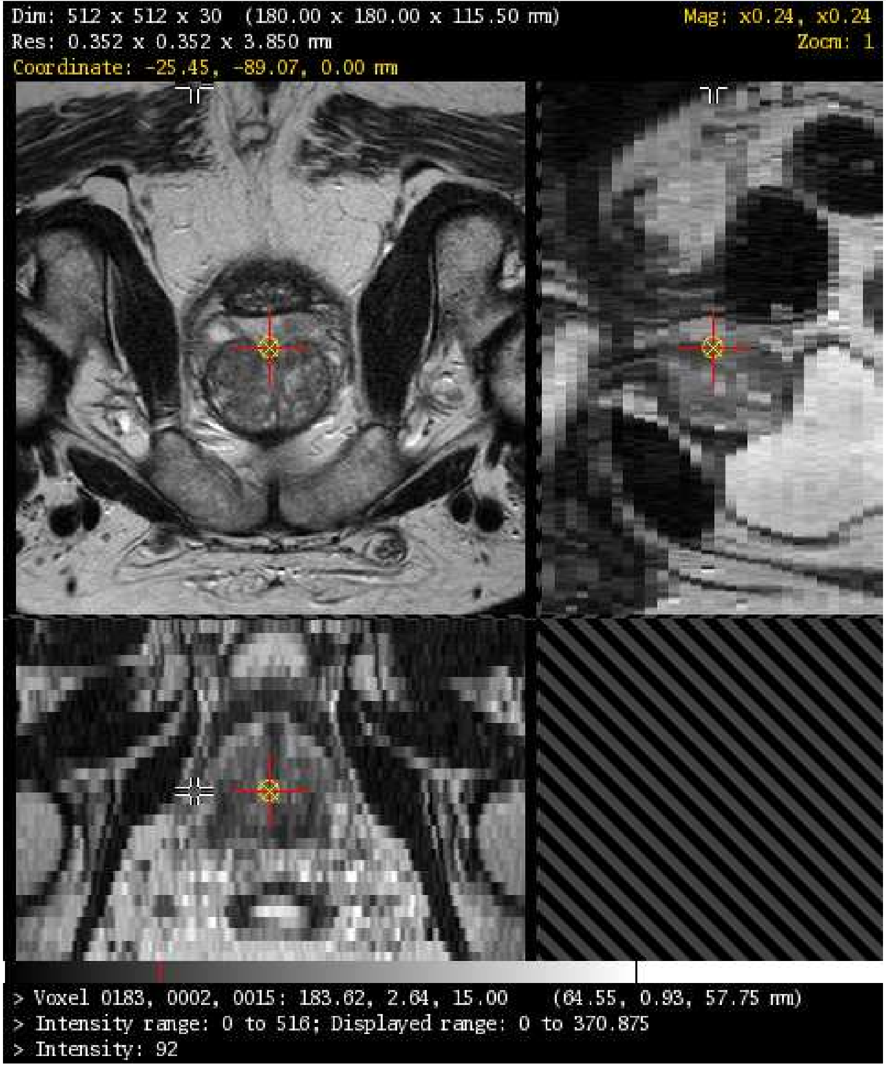	

```
Dim: 512 x 512 x 30  (180.00 x 180.00 x 115.50 mm)                  Mag: x0.24, x0.24
Res: 0.352 x 0.352 x 3.850 mm                                                Zoom: 1
Coordinate: -25.45, -89.07, 0.00 mm
> Voxel 0183, 0002, 0015: 183.62, 2.64, 15.00    (64.55, 0.93, 57.75 mm)
> Intensity range: 0 to 516; Displayed range: 0 to 370.875
> Intensity: 92
```

A.27 Brachytherapy Patient 01

Patient for Prostate Brachytherapy

Name	XMR-01
Description	Structural MRI of prostate and surrounds
Modality	
Size (Voxels)	256x256x20
Voxel Dimensions (mm)	0.703x0.703x3.300

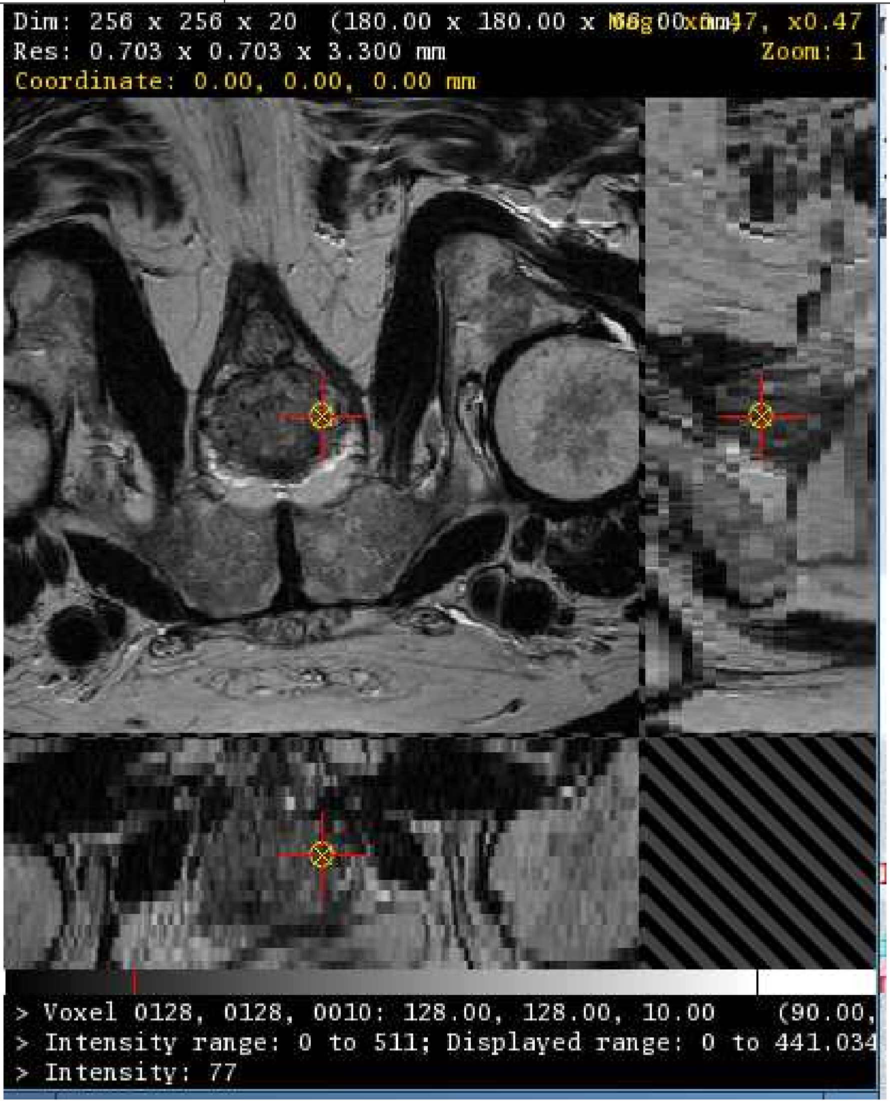

```
Dim: 256 x 256 x 20  (180.00 x 180.00 x 66.00 mm)  Mag: x0.47, x0.47
Res: 0.703 x 0.703 x 3.300 mm                               Zoom: 1
Coordinate: 0.00, 0.00, 0.00 mm
```

```
> Voxel 0128, 0128, 0010: 128.00, 128.00, 10.00    (90.00,
> Intensity range: 0 to 511; Displayed range: 0 to 441.034
> Intensity: 77
```

A.28 Brachytherapy Patient 02

Patient for Prostate Brachytherapy

Name	XMR-02
Description	Structural MRI of prostate and surrounds
Modality	
Size (Voxels)	256x256x25
Voxel Dimensions (mm)	0.703x0.703x3.850
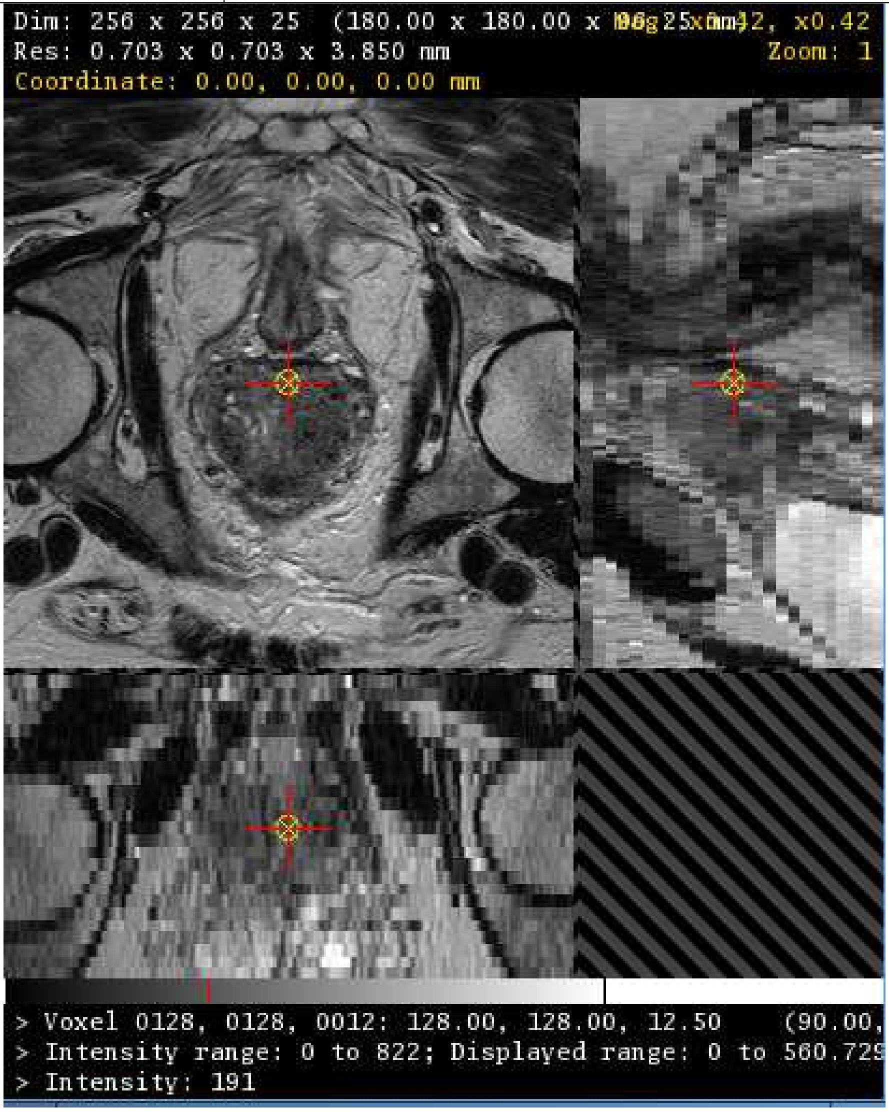	

```
Dim: 256 x 256 x 25  (180.00 x 180.00 x 96.25 mm)  Mag: x0.42, x0.42
Res: 0.703 x 0.703 x 3.850 mm                            Zoom: 1
Coordinate: 0.00, 0.00, 0.00 mm
```

```
> Voxel 0128, 0128, 0012: 128.00, 128.00, 12.50    (90.00,
> Intensity range: 0 to 822; Displayed range: 0 to 560.729
> Intensity: 191
```

A.29 Cadaver Patient 01

Adult Female Pelvis

Name	Cadaver-01
Description	Adult Female Pelvis
Modality	CT
Size (Voxels)	441x252x114
Voxel Dimensions (mm)	0.715x0.715x2.000
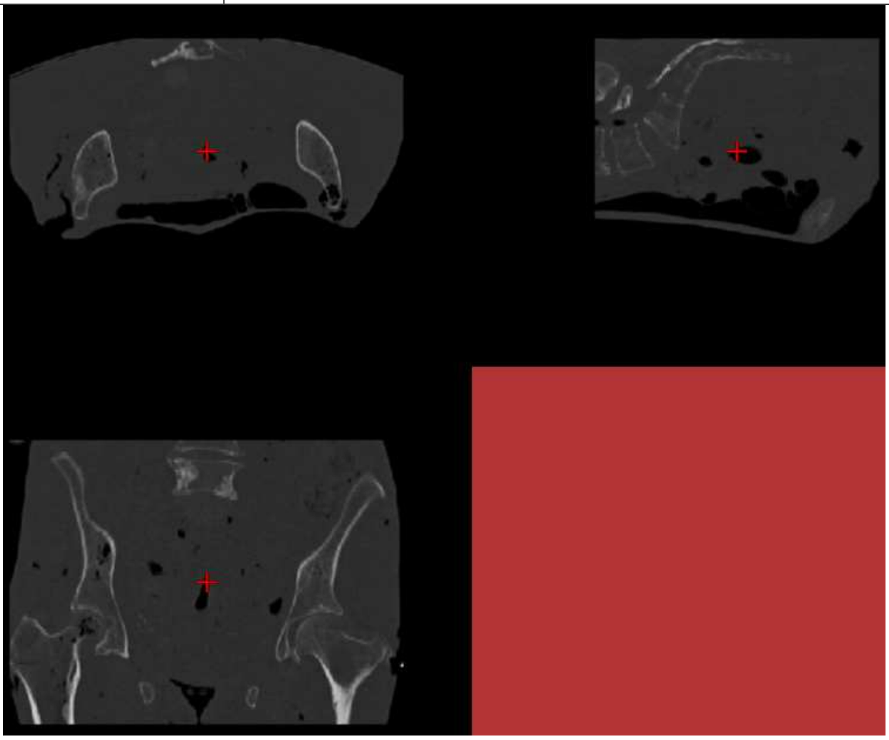	

A.30 Cadaver Patient 02

Adult Female Pelvis

Name	Cadaver-02
Description	Adult Female Pelvis
Modality	CT
Size (Voxels)	373x222x111
Voxel Dimensions (mm)	0.793x0.793x2.000
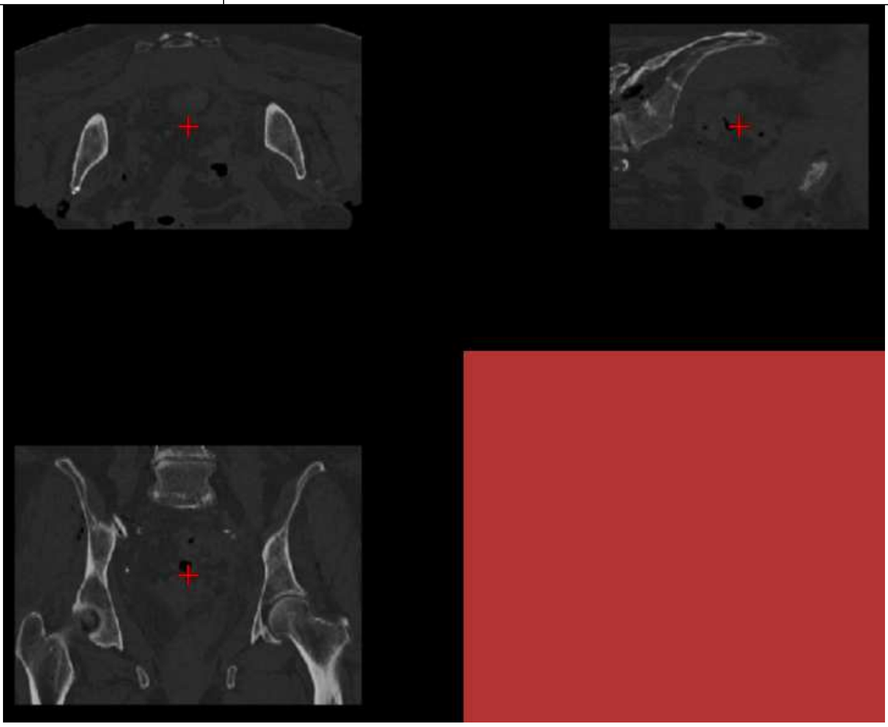	

A.31 Cadaver Patient 03

Adult Female Pelvis

Name	Cadaver-03
Description	Adult Female Pelvis
Modality	CT
Size (Voxels)	373x203x119
Voxel Dimensions (mm)	0.793x0.793x2.000
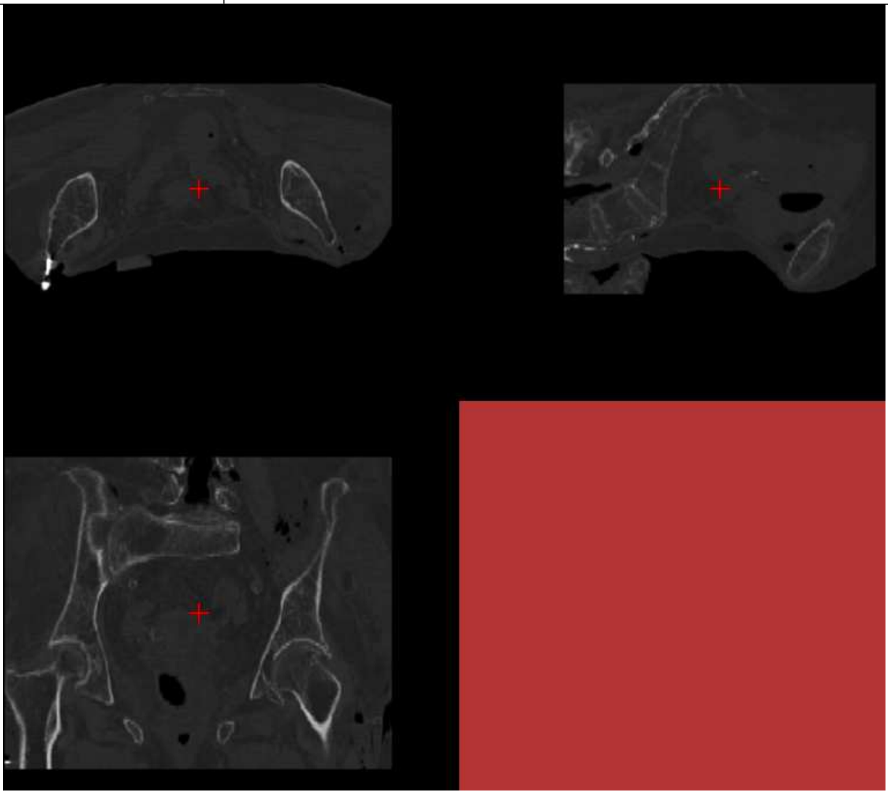	

A.32 Plastic Anatomy Phantom

Plastic Adult Male Pelvis

Name	Phantom
Description	Plastic Adult Male Pelvis
Modality	CT
Size (Voxels)	373x137x776
Voxel Dimensions (mm)	0.914x0.914x0.5
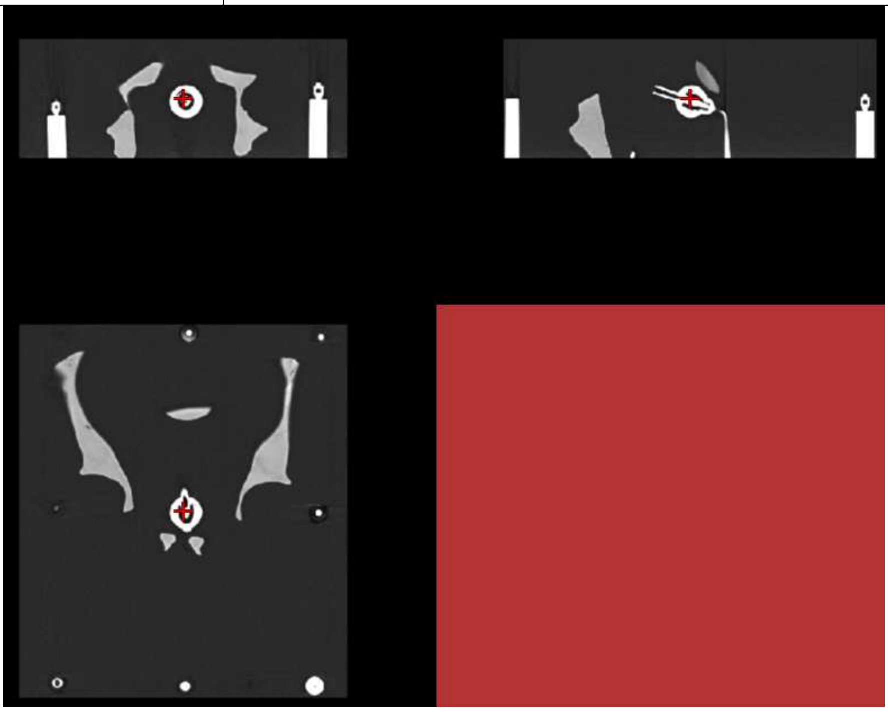	

Appendix B

Fitting a Male Shape Model to a Female Pelvis

The cadaver data used in Chapter 4 was for female pelves, raising the obvious question of whether a model based on male pelves can be applied to a female pelvic bone. The pelves are registered with a 9 DOF affine transform prior to the model being applied, so the difference in general size should not affect the model performance. To assess whether the model was finding suitable shapes for the female pelves the component coefficients for each of the twelve principal components were compared against the component coefficients found by fitting the shape model to the male CT data used to build the shape model (on a leave one out basis) data, see Chapter 3. Figure B.1 overlays the component coefficients for the three cadavers and the phantom with the distribution of training set parameters (approximately normal). The same data is quantitatively examined in Table B.1 using a Kolmogorov-Smirnov test to examine the hypothesis that the coefficients for the cadaver data are drawn from the same distribution as the results of the male pelves in the model.

For the most part the cadavers and data sets conform to the expected model shape. The results for mode 3 are significantly different. The results for mode 1 also appear unusual. These modes may capture some shape information separating male from female pelves.

As the optimisation algorithm allows the generation of a solution outside the model space it seems likely that a reasonable solution will be generated for the female data, however further investigation is required to validate this.

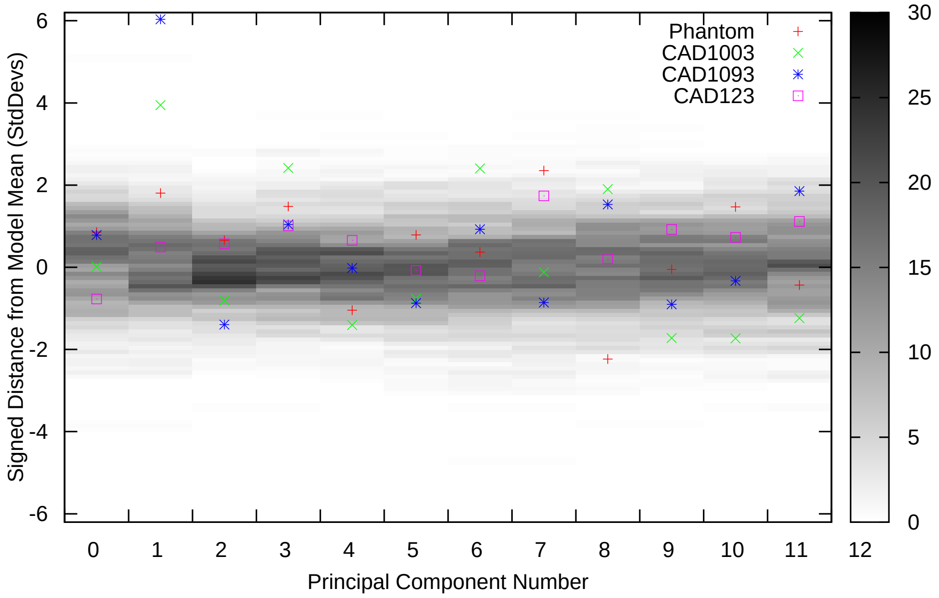

Figure B.1: The solution component coefficients found for the 3 cadaver and 1 phantom data sets, overlaid on histograms of the component coefficients from the male data sets used to build the model. The phantom falls within the expected model shape. As shown in Table B.1, the female cadavers do not.

Mode	KS Test Stat. (D)	P-Value	Accept Null at 5 %?
0	0.2643	0.9576	Yes
1	0.6643	0.07795	Yes
2	0.5214	0.2909	Yes
3	0.8786	0.003961	No
4	0.281	0.9315	Yes
5	0.5071	0.3227	Yes
6	0.5024	0.3337	Yes
7	0.2929	0.9076	Yes
8	0.6024	0.1449	Yes
9	0.5143	0.3066	Yes
10	0.3381	0.7699	Yes
11	0.5048	0.3282	Yes
All	0.1871	0.1636	Yes

Table B.1: The results of a Kolmogorov-Smirnov to test whether the observed values of model mode coefficients for the cadavers come from the same distribution as that of the 420 points tested using the model training data. The null hypothesis is that the observed cadaver coefficients (3 for each mode) are drawn from the same distribution observed for the model training data (420 for each mode). The p value is the probability that the KS test statistic value would occur if the null hypothesis is true. Except for mode 3 the null hypothesis is not rejected at a significance level of 5%. This raises the interesting possibility that mode 3 contains information that differentiates male from female pelves.

Appendix C

Estimating the Tracking Synchronisation Error Using Image Information

C.1 Aim

Using the method described in Chapter 5 there will be a tracking synchronisation error due to the processing time of the frame grabber versus the processing time of the Optotrak. The frame times of images grabbed by the frame grabber will not match those reported from the Optotrak. For tracking of static endoscope poses this should not be an issue, unless the lag is very large, which would require the endoscope to be held still for a prolonged period before image overlay commenced. The aim here is to quantify this lag and its stability. If the lag is stable then it should be trivial to adjust for it by adjusting the recorded times. Here we quantify the lag by recording tracked video of the workstation's own internal clock. We also present a method for estimating the frame lag using in theatre data by correlating two estimates of the endoscope tip speed, one based on the tracking data and the other based on the image data.

C.2 Method

The frame synchronisation error is here quantified by taking video images of the workstation's internal clock and comparing the times shown with those reported by the tracking system.

An alternative to this is to estimate the lag by measuring the correlation of two estimates of the endoscope tip speed, one based on the tracking data and the other on the image information. The endoscope speed was first estimated from the tracking data by calculating the displacement at the endoscope tip and dividing it by the frame rate to get an estimate of endoscope speed. Each image in the endoscopic video was then compared with the image preceding it using NMI. If the anatomy being imaged is relatively static (ie only the endoscope is moving) then the value of NMI should be an indicator of endoscope tip speed. This gives two one dimensional signals for the endoscope tip speed. By moving the image derived signal back and forwards through time relative to the tracking data derived signal it is possible to find the point at which they are maximumly correlated. This should correspond to the frame lag.

C.3 Data

A dedicated video of the workstation clock was used for the first method. See Figure C.1. Dynamic

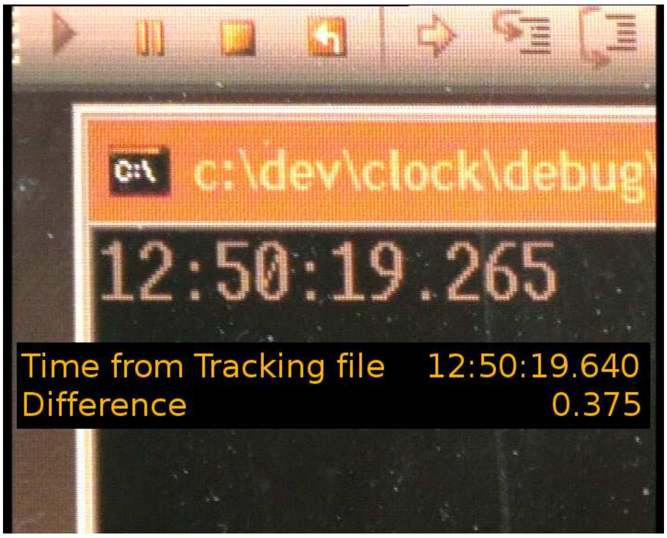

Figure C.1: A video frame of the workstation clock overlaid with the corresponding frame time recorded by the tracking software.

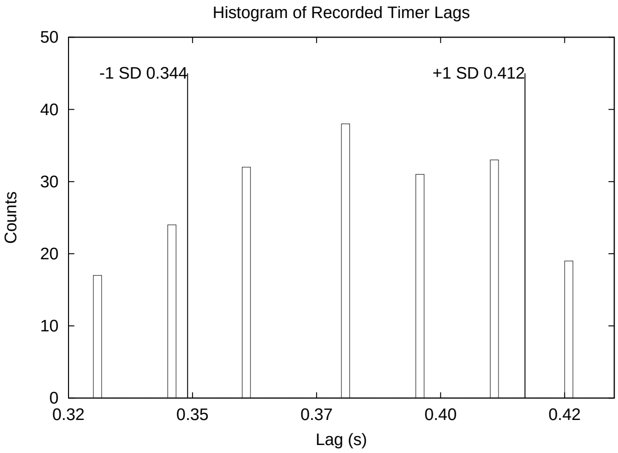

Figure C.2: A histogram of the frame lags recorded using images of the workstations own clock.

video data of a plastic phantom was used for the second approach.

C.4 Results

Using the data from the workstation clock gave a frame lag of 378 ms with a standard deviation of 34 ms. Figure C.2 shows a histogram of the recorded frame lags. The resolution of the workstation clock was 5 ms, while the resolution of the tracking clock was 1 ms. It is interesting that Figure C.2 indicates that the measured frame lag is quantised in units of roughly 20 ms. Presumably this is related to the performance of either the frame grabber or the tracking software.

Figure C.3 shows the correlation coefficient of the two endoscope tip speed signals. Peak correlation occurs at a frame lag of 369 ms, which is very close to the figure of 378 determined using the clock. This suggests that this method may be useful for estimating frame lag retrospectively for in theatre data. This

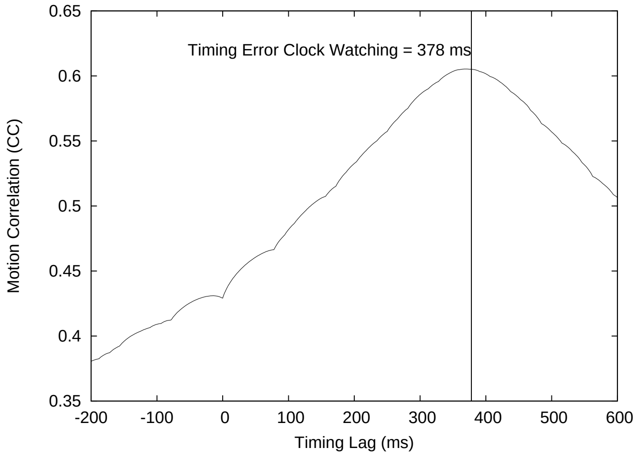

Figure C.3: The correlation coefficient of the two estimates of endoscope tip speed for the video data of the static phantom.

may be complicated by the presence of moving objects (anatomy and end effectors) in the surgical data.

Appendix D

Error due to using a Nominal Prostate

As discussed in Section 1.6.1 the RMS errors at the prostate surface referred to throughout this report do not in fact refer to the actual error for a given data set. This is because rather than using the actual prostate location for a given data set, a nominal prostate location was used for all data sets to allow direct comparison of the results. The potential effect of this standardisation on the actual error for a given data set is discussed here.

The degree to which the actual TRE is under or overestimated will depend on where the actual prostate lies in relation to the nominal prostate and what the actual TRE is. For TRE s in the range 0.8 to 4 mm it was found experimentally that the RMS error in TRE due to mis-positioning of the prostate was as shown in Figure D.1. At present the typical distances of the actual prostates from the nominal prostate are unknown. This could be determined relatively quickly. However, it is likely to be significantly less than 50 mm as the data sets are all rigidly aligned and scaled to match. Therefore the mismatch caused by using the nominal prostate should be significantly less then 0.4 mm, which is negligable compared to the expected system error.

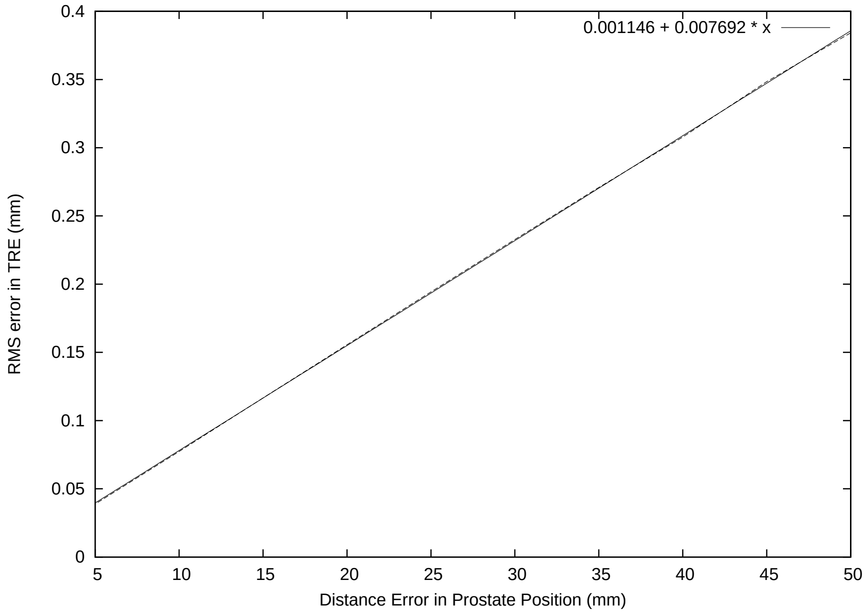

Figure D.1: Errors in prostate position due to the use of nominal rather than actual prostate surface points. This is based on a set of data files with RMS TRE s between 0.8 and 4 mm. At present the distance error for the prostate is unknown, however it should be significantly less than 50 mm. Therefore it is assumed that the error due to using the nominal prostate is small compared to the overall error.