



Imperial College

OF SCIENCE, TECHNOLOGY AND MEDICINE
MECHANICAL ENGINEERING DEPARTMENT

COMPUTER ASSISTED IMAGE-BASED TRAINING AND MONITORING SYSTEMS FOR MINIMALLY INVASIVE SURGERY

by

Adrian Barrett

Submitted for the degree of Doctor of Philosophy of the University of
London and for the Diploma of Imperial College.

LONDON, NOVEMBER 2002

Abstract

This thesis describes the research and development of a novel training and monitoring system for minimally invasive surgery using intra-operative imaging. A hybrid training system is presented which incorporates the benefits of existing virtual reality and physical-phantom based trainers. Proof-of-concept for the monitoring system is achieved using a laboratory-based mock-up of the in-vivo procedure. An optimized algorithm for detecting the outline of the prostate is described and is shown to perform better than previously reported methods.

Minimally invasive surgery is favoured over traditional “open” surgery techniques due to improved patient mortality and morbidity. However, the closed nature of minimally invasive surgery means that the surgeon loses direct visual feedback, thus making it difficult to maintain knowledge of the position and orientation of the surgical tool with respect to the patient’s anatomy. Visual information is provided to the surgeon, but it is restricted to a small field at the tip of the surgical instrument, so does not provide good feedback to aid in avoiding cutting in “no-go” areas.

The problems with minimally invasive surgery mean that training in the techniques is lengthy and difficult. Training aids are available, but even the most sophisticated systems suffer from a lack of realistic tactile feedback or do not provide enough navigational information to facilitate training. A novel hybrid training system is presented which combines previous approaches to tackle both these failings. An innovative extension of the training system is also described – a monitoring system which provides additional information to the surgeon during an actual procedure and can issue warnings of imminent surgical mistakes.

The monitoring system requires up-to-date knowledge of the position and shape of the soft tissue surgical targets as they move and change during the course of a procedure. The thesis investigates the potential of medical ultrasound imaging as a means of obtaining this information intra-operatively. Existing methods of automatically detecting soft-tissue organ boundaries in ultrasound images are

reviewed and implemented. A rigorous statistical analysis, comparing the algorithms' results to human expert opinions, has been undertaken to identify the best outlining method which is shown to be a genetic algorithm creating potential outlines from a deformable shape model. Two aspects of the outlining algorithm (namely the shape model and method of testing a candidate outline against the ultrasound image) have been optimized to produce an algorithm which provides significantly more accurate and reliable results than previous routines.

A laboratory-based training and monitoring system has been built to mimic the transurethral prostate resection procedure and to provide positional feedback to the surgeon. The system uses a novel prostate phantom which deforms under the action of the surgical tool, which may be resected and is pervious to ultrasound. A computer model of the prostate phantom is generated by automatically outlining a set of ultrasound scans pre-operatively. A motorized platform moves the ultrasound probe to image the phantom intra-operatively. These scans are processed to update the computer model and maintain registration during deformation and movement. An optical overhead tracking system allows monitoring of the cutting tool and the system gives warnings on proximity to vital organs at the limits of safe resection.

The result of the research shows that intra-operative ultrasound imaging is a potential method for tracking the motion and deformation of the prostate during surgery. However, several difficulties are highlighted which prevent ultrasound-based tracking at present. These problems are discussed and possible solutions are provided as directions for future research.

Acknowledgments

There are many people without whose generous help and support this thesis would never have been completed. First of all, I would like to thank my supervisor, Professor Brian Davies, for his guidance and patience throughout the course of my studies.

This research was partly funded by an Engineering & Physical Sciences Research Council grant (GR/L80577), with valued industrial support from Karl Storz Endoscopy (UK) and Limbs and Things Ltd.

Many thanks are due to the staff and students of the Mechatronics in Medicine Laboratory for making my time at Imperial interesting and rewarding. Particular mention goes to Dr Paula Gomes for her support, Dr Simon Harris for knowing all the answers and to Ferdinando Rodriguez y Baena for coffee and comedy! I would also like to thank the staff and students of the CIMI Lab, Nanyang Technical University for their continued friendship and for making my time in Singapore an enjoyable experience, in particular I owe a debt of gratitude to Professor Ng Wan Sing.

Dr Lewis Jones and Dr Stephen Breen have been great friends and have willingly provided much valuable discussion of my work over a few beers. Thanks also go to Amo, Phil, Noush, Nick P, Mork, Gareth, Jackie, Az, Antho, Tony, Ness, Patsy, Liz, Lavie, Gillian, Bex, Vandna and Nikki for their fantastic friendship and for providing much needed entertaining distraction from my work.

Finally, and most of all, I would like to thank my family and Katy for their love, support and encouragement.

Contents

Abstract	1
Acknowledgements	3
Contents	4
List of Figures	11
List of Tables	18
Notation	19
1 Introduction	21
1.1 Introduction	21
1.2 Training systems for minimally invasive surgery	23
1.3 Monitoring systems for MIS	25
1.4 Aims and objectives	27
1.5 Transurethral resection of the prostate	28
1.6 Overview of the thesis	31
2 Background	35
2.1 Introduction	35
2.2 Minimally invasive surgery	35
2.3 Surgical training	37
2.4 MIS training systems	38

2.4.1	MIS physical model training systems	39
2.4.2	MIS virtual reality training systems	42
2.4.3	VR versus physical models	47
2.4.4	Assessment of surgical skill	50
2.5	CASTS	52
2.5.1	Tracking technologies	52
2.6	MIS monitoring - CASMS	56
2.6.1	Intra-operative imaging	57
2.6.2	Ultrasound speckle	59
2.6.3	Ultrasound artefacts	60
2.7	Conclusions	64
3	CASTS TURP training system	65
3.1	Introduction	65
3.2	System components	65
3.2.1	Surgical equipment	67
3.2.2	Prostate phantom	69
3.2.3	Instrumenting the tools	69
3.2.4	Computer system	73
3.3	System trials	75
3.4	Improving the VR accuracy	77
3.4.1	Tracking the prostate phantom	78
3.4.2	Prostate ultrasound model	80
3.5	Conclusion	82
4	Prostate outline detection	83
4.1	Introduction	83
4.1.1	Requirements of an outlining algorithm for CASMS	85
4.1.2	Algorithm implementation	86

4.2	Edge based detection schemes	87
4.2.1	Edge definition	87
4.2.2	First derivative approximations	89
4.2.3	Second derivative approximation	90
4.2.4	Results of the first and second derivative approximations .	91
4.2.5	Noise smoothing filters	97
4.3	Outline detection in ultrasound images	98
4.3.1	Edge based detection schemes	98
4.3.2	Laplacian of Gaussian	101
4.3.3	Edge strength and second derivative in gradient direction .	103
4.3.4	Sticks and anisotropic diffusion filtering	105
4.3.5	Radial bas-relief	108
4.3.6	Edge selection	112
4.4	Model based detection schemes	114
4.4.1	General model based image processing	116
4.4.2	Snake based techniques	117
4.4.3	Statistical shape modelling	124
4.5	Region based detection schemes	127
4.5.1	General region-based image processing	128
4.5.2	Texture based segmentation	132
4.5.3	Neural networks	134
4.5.4	Gradient direction based region growing	141
4.6	Discussion	142
4.6.1	Edge based detection schemes	144
4.6.2	Model based detection schemes	146
4.6.3	Region based detection schemes	147
4.7	Conclusion	148

5	Assessment of outlining algorithms	150
5.1	Introduction	150
5.2	Image data	151
5.3	Reference outlines	153
5.3.1	Direct measurement of the outline	153
5.3.2	Cross-modality registration	155
5.3.3	Prostate phantoms	157
5.3.4	Ultrasound outlining by “experts”	161
5.4	Outline comparison	164
5.4.1	Derived parameters	165
5.4.2	Error metrics	165
5.5	Statistical analysis	170
5.5.1	Repeated measurements study	170
5.5.2	Joint agreement study	172
5.5.3	Percentage statistic	173
5.6	Results	174
5.6.1	Multiple experts results	176
5.6.2	Reference outline comparison	179
5.7	Conclusion	180
6	Optimizing the prostate outlining	182
6.1	Introduction	182
6.2	Feature modelling	183
6.2.1	Feature model formation	184
6.2.2	Feature modelling discussion	185
6.3	Automated shape alignment	188
6.3.1	Simulated annealing	189
6.3.2	Expert outlines alignment	190
6.4	Point distribution model for the aligned expert outlines	191

6.4.1	Selecting corresponding points	192
6.4.2	Model formation	193
6.4.3	Polar corresponding points	194
6.4.4	Equally spaced corresponding points	195
6.5	Optimizing the corresponding point positions	197
6.5.1	Modified Chalana and Kim corresponding points	202
6.5.2	Model information	205
6.5.3	Minimum description length	205
6.5.4	Outline parameterisation	207
6.5.5	Optimized corresponding points	208
6.6	Statistical analysis	211
6.7	Image pre-processing	213
6.7.1	Speckle reduction	213
6.7.2	Edge-enhancement	213
6.7.3	Statistical image input dataset	214
6.8	Statistical measures	215
6.8.1	Point, line and patch selection	215
6.8.2	Indicators	218
6.9	Analyzing the indicators	218
6.9.1	Kolmogorov-Smirnov test	218
6.9.2	Random points histogram	220
6.9.3	Random outlines histogram	222
6.10	Results	225
6.10.1	Multiple expert results	226
6.10.2	Reference outline comparison	227
6.10.3	Algorithm execution speed	229
6.11	Discussion	229
6.11.1	Shape model optimization	229

6.11.2	Objective function optimization	230
6.11.3	Algorithm speed	231
6.11.4	Additional outlining applications	232
6.12	Conclusion	233
7	CASMS TURP monitoring system	235
7.1	Introduction	235
7.2	System components	236
7.2.1	Tracking the resectoscope	237
7.2.2	Ultrasound probe platform	238
7.3	Initial operating theatre trials	240
7.3.1	Optical tracking system camera	241
7.3.2	Tracked instruments	242
7.3.3	Ultrasound platform	243
7.3.4	Ultrasound	243
7.3.5	OT trial outcomes	245
7.4	Lab based TURP monitoring/training system	247
7.4.1	Overcoming the major OT problems	247
7.4.2	Phantom outline detection	249
7.4.3	Modelling the treatment volume	251
7.4.4	Software architecture	265
7.4.5	Modelling the resection cavity	275
7.5	CASTMS system testing	276
7.5.1	CASTMS setup	277
7.5.2	Experimental procedure	278
7.5.3	Results	281
7.6	Discussion	284
7.7	Conclusion	289

8	Conclusions and future work	291
8.1	Introduction	291
8.2	Summary of conclusions	291
8.3	MIS training system	292
8.3.1	Future work	294
8.4	Prostate ultrasound image processing	294
8.4.1	Future work	297
8.5	MIS monitoring system	298
8.5.1	Further work	299
	Bibliography	301
A	Speckle reduction algorithms	322
A.1	Speckle smoothing algorithms	322
A.1.1	Anisotropic diffusion filter	322
A.1.2	Homogeneous region growing median filter	325
A.1.3	Median filter	327
A.1.4	Separable median filter	327
A.1.5	Variable length median filter	328
A.1.6	Double-window modified trimmed median filter	329
A.1.7	Adaptive weighted median filter	329
A.1.8	Directional median filter	330
A.1.9	Sticks filters	331
A.2	Edge enhancing algorithms	333
A.2.1	Order statistics filter	333
A.2.2	Rotating kernel transformation	334
A.2.3	Ratio of averages filter	335

List of Figures

1.1	Anatomy of the prostate.	28
1.2	Transrectal ultrasound system and typical scan.	29
1.3	Resectoscope components.	30
2.1	Bozzini's Lichtleiter. The first endoscope.	36
2.2	Kelling's "luft tampanade" for insufflating the abdominal cavity. .	36
2.3	Simple laparoscopic suturing trainer.	39
2.4	Basic training in resectoscope manipulation and co-ordination. . .	40
2.5	Sinus endoscopic trainer.	40
2.6	Arthroscopic trainer.	41
2.7	Torso and internal models for laparoscopic surgery training. . . .	41
2.8	Limbs and Things training systems.	42
2.9	MIST VR training system.	44
2.10	Knee arthroscopy simulator.	45
2.11	ESS VR sinus simulator.	46
2.12	VEST MIS trainer.	47
2.13	Anatomy and endoscopic view.	48
2.14	Simbionix system.	48
2.15	VEST gynaecological system.	49
2.16	Ultrasound artefacts.	62
3.1	The CASTS demonstrator for TURP.	66
3.2	The CASTS system hardware.	66

3.3	Resectoscope assembly ready for use.	67
3.4	Resectoscope electrodes.	68
3.5	Diathermy foot switch controls and base plate.	68
3.6	Mock-up abdomen.	70
3.7	Resectoscope motion.	71
3.8	The resectoscope IRED ring tracker.	72
3.9	Potentiometer attachment to detect the cutting loop motion.	72
3.10	Prostate phantom casts.	74
3.11	VR model of the Limbs and Things phantom.	74
3.12	User interface for the CASTS system.	75
3.13	Initial trials of the CASTS resection system.	76
3.14	CASTS user interface with deformed VR prostate model.	77
3.15	User interface following a resection training session.	78
3.16	Phantom casts and moulds.	80
3.17	Resectable prostate ultrasound phantom.	81
4.1	First and second derivatives for a one-dimensional edge.	88
4.2	Microscopic image of red blood cells.	92
4.3	First derivative approximations.	92
4.4	Edge maps a) Roberts, b) Sobel, c) Prewitt operators.	93
4.5	Second derivative approximation and binary result.	93
4.6	Ultrasound image of the prostate.	95
4.7	Sobel operator applied to ultrasound image.	95
4.8	Prostate ultrasound image with interpolated filled probe region.	96
4.9	Sobel operator applied to filled ultrasound image.	96
4.10	Noise smoothed edge detection on blood image.	99
4.11	Noise smoothed edge detection on prostate image.	100
4.12	Zero-crossings of the LOG filtered prostate ultrasound image.	102
4.13	The four gradient directions in the pixel neighbourhood.	102

4.14	Edge intensity result from the Zhao algorithm.	103
4.15	Edge intensity result from the Aarnink algorithm.	105
4.16	Edge map result from the Aarnink algorithm.	106
4.17	Sticks and anisotropic diffusion filtered image.	108
4.18	Edge map from the Pathak algorithm.	109
4.19	Binary result from the RBR algorithm.	110
4.20	Original and modified RBR	111
4.21	Modified RBR applied to the input ultrasound image.	111
4.22	Edge intensity result from the modified RBR technique.	112
4.23	Results of the edge based detection schemes.	115
4.24	Problems with pressurized snakes.	121
4.25	Gradient vector flow snake result on the test ultrasound image .	123
4.26	Snake applied to heart MRI image.	123
4.27	Problems with the GVF snake.	124
4.28	Feature modelling and genetic algorithm result.	127
4.29	Region growing result for the ultrasound image of figure 4.8. . . .	129
4.30	Merging algorithm applied to the prostate ultrasound test image.	130
4.31	Quadtree decomposition.	131
4.32	Quadtree decomposition applied to the test ultrasound image. . .	131
4.33	Split and merge algorithm result.	132
4.34	Texture based segmentation result.	135
4.35	Training image and its ideal segmentation.	135
4.36	Inputs for a given pixel in the 5-6-2 neural network	136
4.37	5-6-2 network result on the ultrasound image of figure 4.35. . . .	137
4.38	Inputs for the 132-30-2 neural network.	138
4.39	132-30-2 neural network result.	138
4.40	132-30-2 neural network result after extensive training.	139
4.41	Inputs for a given pixel for the 61-10-2 neural network.	140

4.42	Result of the 61-10-2 neural network.	140
4.43	Results for the 61-10-2 network applied to non-training images. . .	141
4.44	Local gradient compared to centre gradient.	142
4.45	Algorithm processing pipelines.	143
5.1	Examples from the ultrasound image data set.	152
5.2	A typical transurethral ultrasound image of the prostate.	154
5.3	MR image of the prostate.	156
5.4	CT image of the prostate.	157
5.5	Brachytherapy phantom and resulting ultrasound image.	158
5.6	Ultrasound phantom box and prostate.	160
5.7	Phantom ultrasound and optical image.	160
5.8	TRUS images and expert boundaries.	162
5.9	Problems in finding the average outline.	163
5.10	Finding the average of two outlines.	164
5.11	Hausdorff distance (thick black line) for two pairs of outlines . . .	166
5.12	Area between corresponding points	168
5.13	Outline area approximation using the trapezium rule	169
5.14	Results for the ultrasound image with close expert agreement. . .	175
5.15	Results for the ultrasound image with poor expert agreement. . .	176
5.16	Percentage statistic stringency.	178
5.17	Box plots of the average boundary distances.	180
6.1	Corresponding points of two outlines.	184
6.2	Expert outline which cannot be represented by polar points. . . .	184
6.3	Corresponding points and mean outline for the FM.	186
6.4	Effect of varying the weights on the FM.	187
6.5	"Un-prostate-like", valid shapes from the FM.	188
6.6	Simulated annealing outline matching algorithm testing.	191
6.7	Outline results from the simulated annealing automated alignment. .	192

6.8	Manually annotated landmark points.	193
6.9	Polar corresponding points from the aligned expert outlines. . . .	195
6.10	Effect of varying the weights on the polar PDM prostate model. .	196
6.11	Extremities of the polar corresponding points PDM.	197
6.12	Equally spaced corresponding points from the aligned expert outlines.	198
6.13	Weight variation in the equally spaced PDM.	199
6.14	Extremities of the equally spaced PDM.	200
6.15	Aligned modified corresponding points.	202
6.16	Effect of varying the modified corresponding points PDM.	203
6.17	Extremities of the modified corresponding points PDM.	204
6.18	Corresponding points parameterization.	208
6.19	Optimized corresponding points from the aligned expert outlines.	209
6.20	Effect of varying the weights on the optimized PDM.	210
6.21	Extremities of the optimized PDM.	211
6.22	Histogram result for an ideal statistical measure.	212
6.23	Bi-linear interpolation	216
6.24	Point, line and patch regions for the statistical analysis.	217
6.25	Expert boundaries and null hypothesis regions.	217
6.26	Prostate image and HRGMF result.	221
6.27	Histograms of random, expert and reference outlines.	223
6.28	Histograms of the FM and GA algorithm result.	224
6.29	Comparison of best and Wu's fitness function.	224
6.30	Input image for which best fitness function fails.	225
6.31	Box plots of the average boundary distances.	228
6.32	Outline results of the optimized algorithm.	228
7.1	Components of the CASMS TURP system.	237
7.2	New IRED tracking component.	238
7.3	Optical tracking IRED mount for monitoring the trigger motion. .	239

7.4	Ultrasound probe platform.	240
7.5	Flashpoint camera on the OT mechanical arm.	242
7.6	Testing the mock-up tools during TURP.	242
7.7	Mock-up of the ultrasound probe platform.	244
7.8	Acoustic shadowing due to the presence of the resectoscope. . . .	244
7.9	Artefacts due to the electrosurgical current.	245
7.10	Laboratory based CASTS/CASMS system.	247
7.11	Phantom ultrasound image after preprocessing and outline detection.	250
7.12	IRET tracker for the TRUS probe.	251
7.13	Image and ultrasound coordinate systems.	252
7.14	World and ultrasound coordinate systems	253
7.15	Relationship between the coordinate systems.	256
7.16	Problems of polar outlines.	260
7.17	Enlarged median lobe problems.	261
7.18	Delaunay triangulation.	263
7.19	Initial triangles around data points.	263
7.20	Software class hierarchy.	266
7.21	Frame Factory thread flow diagram.	268
7.22	Motion Controller flow diagram.	271
7.23	Frame Processor and Gland Processor threads flow diagram. . . .	272
7.24	Operation thread flow diagram.	274
7.25	The GUI of the CASTMS software.	275
7.26	CASTMS physical setup.	276
7.27	Gelatin phantom set into perspex base.	277
7.28	CASTMS in use in the laboratory.	278
7.29	Prostate phantom in the Perspex cutting box.	280
7.30	A slice of the resected phantom.	280
7.31	Phantom slice with overlaid outlines.	281

7.32 Finding the resection boundary.	282
7.33 Area and average distance between the tracked/untracked outlines.	283
7.34 Phantom slice with overlaid outlines.	283
7.35 Area and average distance between the tracked/untracked outlines.	284
7.36 Rendered resection volumes.	285
7.37 “Pistol” type transrectal ultrasound probe.	287
7.38 The Voluson 730 3D ultrasound system.	288
A.1 Transrectal ultrasound image of the prostate	323
A.2 ADFs applied to the test ultrasound image.	325
A.3 HRGMF result.	327
A.4 Standard MF result.	327
A.5 SMF result.	328
A.6 VLMF result.	328
A.7 DWMTMF result.	329
A.8 AWMF result.	330
A.9 Size five, one pixel thick stick masks	331
A.10 DMF result.	332
A.11 Sticks results using a) the MSF and b) the DDS filter.	333
A.12 Order statistics filter result.	334
A.13 RKT result.	335
A.14 ROAF result.	336

List of Tables

1.1	Summary of the components of a MIS training system.	24
1.2	Summary of the components of a MIS monitoring system.	27
2.1	Advantages of VR versus model based MIS trainers.	50
4.1	Initial outlining results.	149
5.1	Repeated observations on a set of subjects.	171
5.2	Algorithm comparison using Hausdorff distance.	177
5.3	Algorithm comparison using average distance.	179
6.1	Variance of the FM, polar PDM and equidistant PDM.	201
6.2	Variances of the modified corresponding points model.	204
6.3	Variances and description lengths for the models.	210
6.4	Speckle reduction algorithms and parameter values	214
6.5	Edge enhancement algorithms and parameter values	214
6.6	Statistical measures.	219
6.7	Algorithm comparison using Hausdorff distance.	226
6.8	Algorithm comparison using average distance.	227
7.1	Difficulties preventing a fully functional TURP CASMS system. .	246

Notation

2D	– Two dimensional.
3D	– Three dimensional.
A/D	– Analogue to digital.
BPH	– Benign prostatic hyperplasia.
CASMS	– Computer assisted surgical monitoring system.
CASTS	– Computer assisted surgical training system.
CASTMS	– Computer assisted surgical training and monitoring system.
CCP	– Correctly classified pixel.
CMM	– Coordinate measurement machine.
COD	– Computer observer disagreement.
CT	– Computed tomography.
DCP	– Distance to closest point.
DOF	– Degrees of freedom.
DRF	– Dynamic reference frame.
ENT	– Ear, nose and throat.
FEM	– Finite element modelling.
FM	– Feature modelling.
GA	– Genetic algorithms.
GVF	– Gradient vector flow.

HIFU	– High intensity focussed ultrasound.
HRGMF	– Homogeneous region growing median filter.
ICP	– Incorrectly classified pixel.
IGES	– Initial graphics exchange specification.
IOD	– Inter observer disagreement.
IREL	– Infra-red emitting diode.
KS	– Kolmogorov Smirnov test.
LOG	– Laplacian of Gaussian.
MFC	– Microsoft foundation classes.
MIS	– Minimally invasive surgery.
MR, MRI	– Magnetic resonance imaging.
NURBS	– Non uniform rational B-splines.
OT	– Operating theatre.
PC	– Personal computer.
PDM	– Point distribution modelling.
RBR	– Radial bas-relief.
RMS	– Root mean squared.
RO	– Reference outline.
SDGD	– Second derivative in gradient direction.
SSM	– Statistical shape model.
TGC	– Time gain compensation.
TRUS	– Transrectal ultrasound.
TURP	– Transurethral resection of the prostate.
TUUS	– Transurethral ultrasound.
US	– Ultrasound.
VR	– Virtual reality.

Chapter 1

Introduction

1.1 Introduction

Minimally invasive surgery (MIS) is rapidly becoming the preferred technique in many areas of surgery. Russell [186] defines MIS as:

“The achievement of a therapeutic objective by the least physiologically disturbing stimulus so that the metabolic, cardio-respiratory and psychological effects are minimal.”

MIS is characterized by the use of small incisions or natural body orifices to allow access to the surgical target. This leads to many advantages over conventional open surgical techniques. Tissue injury is lessened, resulting in reduced post-operative pain and shorter recovery periods benefiting both the patient and hospital. The likelihood of secondary infection is decreased since the patient is less exposed to the external environment. There are cosmetic benefits for the patient with the minimization of visible scarring. There is also evidence to show that patient mortality and morbidity are reduced [148].

Instruments for MIS are typically long thin tubes which provide access to the surgical target and through which a variety of tools may be inserted. Visual information of the surgical target is provided by an endoscope, often incorporated with a video camera so that the image is viewed on a monitor. In some techniques, such as laparoscopy¹, multiple incisions are made so that the surgical instruments

¹Laparoscopy involves surgery in the abdominal cavity. Typically four incisions are made; two for surgical instruments, one for the endoscope and one for insufflating the abdomen.

and endoscope occupy different positions, and the surgical instruments are visible in the endoscopic image. Other techniques, such as uroscopy and neuroscopy² use a single orifice – the endoscope and surgical instrument are introduced through the same tubular guide.

The closed nature of MIS results in difficulties for the surgeon. Although generally of excellent magnified quality, the visual field available through an endoscope is much smaller in comparison to open surgical techniques; the image is two dimensional so depth perception is difficult and in the case of single orifice MIS only the tip of the surgical instrument is visible. Excellent hand-eye co-ordination is required since the surgeon must observe the endoscopic image rather than looking at his hands. Excessive bleeding can lead to occlusion of the endoscopic view. The dexterity of the surgeon is reduced and the surgical tool position and orientation may become confused from the loss of direct visual information and the reversal of direction inside the patient. The surgeon's knowledge of the surgical target is often built up from preoperative scans which must be maintained as a three-dimensional mental map of the anatomy, and be registered with the tool position and orientation. Delicate structures which should not be damaged may lie outside the field of view - their positions must be maintained within the surgeon's mental map. Finally, MIS removes the direct tactile feedback available to the surgeon in open surgery.

The problems involved with MIS mean that training in the techniques is lengthy and difficult. In the past, surgical training was undertaken in a hands-on manner, with the trainee shadowing a consultant and learning under his observation. With MIS, the trainee has many more skills to acquire; they must first be competent in manipulating the surgical tools and maintaining knowledge of position and orientation with respect to the anatomy, before practicing on a patient. To this end, a variety of training systems have been devised. The simplest systems involve using tools to manipulate objects without direct vision and allow honing of such skills as cutting and suturing, *e.g.* "lunch box" trainers which use an endoscope to visualize tools passing through a phantom abdomen to resect parts of a chicken or manipulate pieces of plastic [131, 194]. More advanced systems are generally procedure specific and are based around one or a combination of anatomical phantoms, computer graphics and virtual reality, haptic feedback and finite element modelling. However, there is still a wide gap between such systems and the real surgical procedure. A brief discussion of surgical training systems is

²Uroscopy involves insertion via the urethra. Neuroscopy requires a small incision in the skull to allow access to the brain.

given in section 1.2, and the background to surgical training systems is discussed in full in chapter 2, section 2.4.

The term “*monitoring*” is generally associated with systems which provide status reports on the patient during the course of the operation, such as blood pressure, respiratory rate, heart rate and anaesthetic levels. In the context of this thesis, the term is used to describe a system which monitors the progress of the operation and the position of the surgical tools with respect to the anatomy. Such a system stems the gap between training and the real operation by providing assistance to the newly trained surgeon. The components required for a MIS monitoring system are introduced in section 1.3, and discussed more fully in chapter 2, section 2.6 and chapter 7.

1.2 Training systems for minimally invasive surgery

The generic training system discussed in this section falls into the sophisticated end of the range of products, where the training is directed towards a specific MIS procedure. Therefore, the objective is to enable the trainee to acquire and practice the necessary skills required to perform the operation safely and accurately. The components required for a successful trainer are discussed over the following paragraphs.

The manipulators which the trainee uses to practice the surgery should perform in the same manner as the real surgical instruments. These manipulators can be specific products designed for simulating surgical instruments, or be the actual surgical instruments themselves. In either case, they must be encoded in some manner which provides information on their positions³. An important requirement for these manipulators is that they should have the same feel as that during the real procedure, both in terms of their actual physical composition and their interaction with the surgical target, since tactile sensation provides a significant degree of feedback to the surgeon whilst operating.

The visual stimulus, given to the trainee in the form of an endoscopic image, should convey the same information which is received in the real surgery. Two possibilities exist for the display; either a computer generated virtual reality image of the surgical target and instruments, or a real endoscopic image of a phantom

³Examples of several manipulators for MIS training systems are given in chapter 2, which also provides a discussion of methods for tracking real surgical instruments.

Component	Requirement
Manipulators	Realistic tactile sensation
Display	Realistic visual stimuli
Evaluation	Accurate assessment of performance
Additional information	Graduated level of difficulty
Recording	Performance record and assessment of learning rate

Table 1.1: Summary of the components of a MIS training system.

model of the surgical anatomy. Each option has associated advantages and disadvantages. Virtual reality displays can be varied so that different anatomical and pathological scenarios are included, although current virtual reality based trainers lack realism (see section 2.4.2, chapter 2). Phantom models can provide a more realistic display, but it is difficult to model the differences in anatomy and pathology, or to simulate side-effects such as bleeding. The choice also governs the method of providing tactile feedback through the manipulators; a phantom model implies direct tactile feedback through interaction of the tools and model, virtual reality requires force feedback in the manipulators and sophisticated force models for tactile sensations to be simulated.

The training system should provide an evaluation of the surgical task so that performance can be monitored over time. However, the criteria on which a procedure may be judged remains a subject under current research. Various factors are relevant depending on the particular operation, although time taken and the number and nature of errors made are universal. The topic of surgical skill assessment is discussed in more detail in section 2.4.4, chapter 2.

The final component which can be present within a MIS trainer is the ability to provide different levels of difficulty during the surgical procedure. Additional information, such as external views giving extra navigation cues, can reduce the degree of complexity in the initial stages of training. Alternatively, complications which occur rarely during a real procedure can be included in the training system once the user reaches a more advanced level of experience.

Table 1.1 provides a summary of the components of a MIS training system and their requirements.

1.3 Monitoring systems for MIS

The primary objective behind a monitoring system for MIS is to prevent mistakes being made during the course of a surgical procedure. This can only be achieved if the monitoring system has knowledge of the relative positions of the surgical tools with respect to the patient's anatomy. In particular the delicate structures in the vicinity of the operative site must not be damaged. The secondary objectives of such a system are to provide additional information to the surgeon so that navigation is facilitated within the confined space of the operative site, and to provide the possibility of a post-operative review of all the actions made during the procedure.

A strict requirement for a MIS monitoring system is that it should cause as few changes as possible to the operating procedure, and the design should be such that any required changes should have minimal impact. This requirement will enable surgeons who have trained whilst using the system, to still perform the surgery in situations where the monitoring system is not present.

The surgical instruments used during a procedure must be modified so that their positions are known to the monitoring system. This can be achieved through the use of a tracking system which measures the positions of markers attached to the tools. Several different tracking options are available which are discussed in section 2.5.1, chapter 2. The tracking system should provide timely, accurate positional information using tracking markers which do not overly affect the feel of the surgical tools.

Obtaining the relative positions of the tools with respect to the patient's anatomy presents a more difficult challenge, with the degree of difficulty depending on the surgical task being performed. In some applications, such as orthopaedic procedures, the position of the surgical targets are either stationary (using external fixation devices) or can be ascertained by attaching a few fiducial markers to the anatomy. The knowledge of the complete shape of the anatomy can be determined pre-operatively by imaging the patient, processing the images to demarcate the permitted regions for treatment (either automatically or by hand) and constructing a model of the treatment volume from these regions. The model could then be used during the procedure, either directly or after affine transformation according to the position of the tracked fiducial markers. However, the majority of MIS operations have surgical targets which partially or wholly consist of soft tissue. Soft tissue tends to move and deform according to the patient's position and

under the influence of the surgical tools during the procedure. In addition, the surgical target may change considerably during the course of the operation, due to such factors as tissue response to injury (causing swelling) and the relaxation of internal stresses as tissue is removed.

The only solution to the problem of tissue movement during the surgery is to acquire intra-operative images and process them automatically to find the current position and shape of the surgical target. Three imaging modalities exist which have been previously used for intra-operative applications (they are discussed in section 2.6, chapter 2), however they have only been used to provide images to the surgeon during the procedure, rather than being automatically processed to gain an understanding of the anatomy. The image processing must accurately detect the boundaries of the treatable region in each image, at a rate which is sufficiently fast to allow the monitoring system to compensate for changes in the surgical target. Therefore, the image processing must perform without any human input, with the additional requirement that the resulting outline information must be accurate, since there is no time for manual correction.

Once the outline information is obtained from the intra-operative images, it must be aggregated into a complete computer model of the surgical target. Generation (or update) of the model should be sufficiently rapid so that registration is maintained to the surgical anatomy. Interrogation of the model must also proceed at interactive rates, since the model must provide positional information of the surgical boundaries given the position of the surgical instruments. This information must be accurate so that only valid warnings are generated by the monitoring system.

During the course of an operation, there are many factors which demand the surgeon's attention. MIS often involves the complex coordination of sophisticated surgical tools, using both hands and feet, while observing and interpreting video displays of the surgical site, and remembering patient-specific pre-operative information. The additional information presented by a monitoring system must therefore be presented in a clear and concise manner which does not require interpretation on the part of the surgeon. In particular, warnings of imminent mistakes must be clear and timely, but the priority is that they should be accurate.

Table 1.2 provides a summary of the major components required for a monitoring system, and the relevant factors which should be considered for each.

Component	Requirements
Surgical instrument tracking	Timely, accurate positional information
Intra-operative imaging	Fast data acquisition across the surgical volume
Image processing	Rapid, automated detection of surgical boundaries
Anatomical model	Fast generation and precise interrogation
Information display	Real-time simple display and accurate warnings

Table 1.2: Summary of the components of a MIS monitoring system.

1.4 Aims and objectives

The research contained in this thesis is directed towards providing a novel system for training and monitoring in the field of MIS. The project aims are twofold: (i) the development of a training system which bridges the gap between the training and the real surgical environment, and (ii) the development of an associated monitoring system which eases the trainee into the real operation by providing extra information on the operation progress and surgical instrument position, and warnings of imminent surgical mistakes.

Although the research is applicable across the discipline of minimally invasive surgery, a specific implementation for transurethral resection of the prostate (TURP) has been used in this research. The choice of the surgical application relates to the considerable experience in prostate surgery applications by the Mechatronics in Medicine Laboratory (MiMLab), Imperial College. The TURP procedure is highly localized to one area of the anatomy (the prostate) and therefore requires only one computer model. Additionally, the operation uses a single tool thus the system need only track a single tool.

The research has been directed towards the following six objectives:

- To develop a training simulator for TURP which provides additional information to the trainee and enhances the learning process, with the ability to evaluate training performance.
- To investigate current image processing methods which could be used to detect the prostate outline in ultrasound images so that information on the position and shape of the prostate may be obtained intra-operatively.
- To analyze the underlying features of transrectal prostate ultrasound images, and identify the indicators which best correlate to the presence of the prostate outlines.




Figure 1.1: Anatomy of the prostate.

- To assess the hypothesis that ultrasound is a feasible imaging modality for tracking tissue movement during surgery.
- To design an algorithm for fast generation and interrogation of a three-dimensional treatment volume constructed from the surgical boundaries.
- To develop a proof-of-concept, laboratory-based in-vivo monitoring system which provides additional navigation feedback to the surgeon, including warnings on the proximity of the cutting loop to delicate structures, on a mock-up of the TURP procedure.

1.5 Transurethral resection of the prostate

The prostate is a chestnut sized gland found in males at the base of the bladder and just in front of the rectum, as shown in figure 1.1. It is roughly conical in shape, with the base of the cone at the bladder, and the apex pointing downward towards the bulb of the penis. The prostate functions as part of the ejaculatory system, secreting a milky, alkaline fluid that contributes to the semen [145]. The urethra passes through the prostate from base to apex, and characteristically bends through a 35° angle at about two thirds of the length of the gland [179]. This bend often coincides with the position of the verumontanum - a protrusion on the posterior wall of the urethra where the ejaculatory ducts open into the urethra, thought to prevent semen from passing into the bladder during ejaculation [179]. The prostate is enveloped by a fibrous capsule, which varies in thickness from a few millimetres to as little as 0.25 mm [159].




Figure 1.2: Transrectal ultrasound system and typical scan.

Benign prostatic hyperplasia (BPH) is an enlarging of the prostate as part of the natural ageing process. This can cause constriction of the urethra inhibiting the flow of urine from the bladder. Various treatments exist for the relief of the obstruction including drug therapy, resection of the tissue by electrosurgery or laser, microwave therapy, cryosurgery, high intensity focused ultrasound and interstitial laser coagulation. Of the surgical interventions, resection by electrosurgery is the gold standard and is commonly known as transurethral resection of the prostate (TURP).

The initial stage in the treatment of BPH is imaging of the prostate using transrectal ultrasound (TRUS). Figure 1.2 shows a transrectal ultrasound system and a typical TRUS scan. The TRUS procedure allows an estimate of the volume of the prostate to be made; this volume is a key factor in the clinical decision making process and choice of treatment. The TRUS scan also gives the surgeon an initial view of the patient's anatomy with the chance to note the positions of the prostatic capsule, and any potential areas of surgical difficulty, such as calcifications within the tissue.

The TURP operation proceeds as follows. An endoscope is introduced into the urethra through the penis allowing the surgeon to view the urinary tract and bladder. Key anatomical sites such as the verumontanum and seminal vesicle openings are noted. Once the surgeon has a complete mental image of the anatomy, a resectoscope is used to cut away the excess tissue. The resectoscope is fitted with a camera at its distal end and the view from the resectoscope tip is shown on an adjacent monitor. An electrosurgical cutting loop at the proximal end of the resectoscope performs the actual cutting. Movement of the cutting loop is achieved through a trigger on the resectoscope body, and current to the loop is controlled



Figure 1.3: Resectoscope components: (from top) endoscope camera, endoscope, resectoscope, outer sheath, hot wire cutting element.

by a foot switch. A second foot switch allows a lower current supply to the loop which is used to coagulate bleeding arteries. Figure 1.3 shows the resectoscope components.

During the course of the resection, care must be taken not to damage delicate structures within the prostate. Damage to the sphincter mechanism is associated with post-operative incontinence – a condition which occurs in 4% of cases [72]. Resection of the verumontanum leads to post-operative loss of ejaculatory function, and should also be avoided. Puncturing of the prostate capsule (with the possibility of perforating the rectal wall) can lead to persistent bleeding and infection. The prostate capsule is identified on the endoscopic image by its fibrous texture once the covering of hyperplastic tissue has been cut away. However, the capsule can be as little as 0.25mm thick [159], and the most careful resectoscope cut removes 1mm of tissue [20], therefore it is not possible for the surgeon to continually cut away tissue until the capsule appears. Instead, the surgeon must maintain knowledge of the resectoscope position with respect to the prostate capsule and verumontanum, without seeing either on the endoscopic image.

Two objectives have been identified to reduce problems with the TURP procedure. The first is a training system which allows the trainee surgeon to become adept at manipulating the resectoscope and cutting tissue with extra guidance as to the resectoscope position with respect to the delicate structures. The second objective is to extend the training system into the operating theatre by providing intra-operative monitoring of the resectoscope position and orientation, with respect to an updated model of the patient model derived from intra-operative imaging.

1.6 Overview of the thesis

This chapter introduced the field of MIS and presented several advantages of the technique which have resulted in its continued adoption in preference to standard open surgical techniques. However, the benefits also come with some added cost to the surgeon, in terms of the increase in difficulty of performing the procedures and the requirement for additional training above that associated with conventional operations. Various systems exist which allow the skills of MIS to be practised; the components of a generic training system have been described. A novel concept for extending training into the operating theatre and for providing monitoring during the surgical procedure was introduced and the general features of such a system were characterized. The objectives of the research contained in this thesis in the areas of training and monitoring systems for MIS were described, with a specific target operation of TURP chosen for the prototypes. A brief discussion of the prostate was given, including the anatomy and pathology, and the symptomatic treatment via the TURP operative procedure.

Chapter 2, starting on page 35, gives a discussion of various aspects of the background to the research contained in this thesis. A short history of MIS is given, chronicalling the rise of the technique and the increase in complexity of the procedures and surgical tools involved. Various approaches to surgical training are examined with particular emphasis on the the additional difficulties apparent in MIS in comparison to conventional open surgery. Examples of existing MIS training equipment are reviewed and discussed; these include both virtual reality based and physical model based systems. The advantages and disadvantages associated with each type of training system are discussed, and a new hybrid approach is introduced which incorporates elements of both the virtual reality and physical model based systems. The major advantages of the novel training system all require tracking of the surgical tools during the training procedure; various tracking technologies are discussed and the most suitable chosen. The final part of the chapter develops the idea of bringing the training system into the operating theatre in an attempt to close the gap between training and performing an actual procedure on a patient. This extends the concept of a training system into a full blown monitoring system which not only provides additional guidance to the surgeon, but can also provide warnings if incorrect actions are imminent. The vital component of such a system is that it maintains a correct map of the patient's anatomy which is updated through the interpretation of intra-operative images. Several imaging modalities are discussed and evaluated as potential can-

didates for this intra-operative imaging. The chapter finishes with a discussion of the problems associated with the chosen modality – ultrasound imaging.

Chapter 3 (page 65) describes the implementation of a hybrid training system for MIS which is a specific simulator for TURP. Each component of the system is described, including the standard TURP surgical equipment which is identical to that used in the real operation, the additional instrumentation components required to monitor the positions and statuses of the surgical tools, the phantom surgical target and the computer system which provides additional navigation information to the surgeon. Initial results from using the system are presented and discussed, and the problems highlighted. Several potential solutions for these problems are discussed with ultrasound tracking chosen as the most suitable method. The choice of ultrasound tracking has the additional benefit of closing the gap between the training system and monitoring system. The final section of the chapter describes the formation of a novel prostate phantom which can be tracked using ultrasound and resected with the resectoscope in a similar manner to a real prostate.

The major work of the thesis, the automated detection of the prostate outline in ultrasound images, is contained in chapters 4 through 6. The motivation for the outline detection is to enable the tracking of the prostate as it is moved and deformed during the course of a TURP procedure – a feature which is crucial to the success of an intra-operative monitoring system for TURP. Chapter 4, starting on page 83, investigates many potential solutions to the problem of identifying the prostate outline. Examples of common image processing techniques for edge detection are described and shown to perform poorly on ultrasound images due to the low signal to noise ratio inherent in the imaging modality. A comprehensive review of the literature of edge detection in ultrasound images is presented, and candidate solutions for prostate outline detection are described. These solutions fall into three categories of image processing techniques; edge-based, region-based and model-based detection schemes. Existing algorithms from each category have been implemented (including standard and prostate specific) and the results of applying each to a high quality prostate ultrasound image are given. These results are discussed and promising candidates are presented for further analysis.

The topic of chapter 5, which begins on page 150, is an in-depth assessment of the promising algorithms from chapter 4. The algorithms are applied to a set of ultrasound images taken from different patients and covering the range in image quality typically observed in routine prostate ultrasound examinations. The con-

cept of a reference outline (the true outline of the prostate) is introduced, and various potential methods for finding the reference outline are discussed, with averaging of a set of expert drawn outlines chosen as the most appropriate. Several different error metrics are proposed for measuring the difference between two outlines. A thorough statistical analysis is presented which compares the results of the promising algorithms to the reference outlines and the expert outlines, and identifies the most promising algorithm to date.

Chapter 6 (page 182) provides an in depth analysis of the most promising algorithm, which identifies the components of the algorithm that are candidates for optimization. The essential features of this algorithm are a prostate model which is used to generate candidate curves that may match the outline of the prostate in an ultrasound image, and the mechanism by which each outline is assessed to determine if it is indeed a good fit. The first part of the chapter investigates various methods for forming the prostate model and evaluates the potential of each model to produce outlines which are always valid examples of prostate shapes. The work contained in the remainder of the chapter questions previous assumptions about which features of an ultrasound image give the best indication that a prostate outline falls within a given region of the image. Previous work has always started from the assumption that some feature of the image signifies the prostate outline – the image is then searched to find the points where the given feature is abundant and these points are taken to lie on the outline. The analysis reported here is based on a novel approach, where the process is reversed, i.e. starting from the expert delineated outlines for each image, many different indicators are checked to see which ones give a marked difference for outline and non-outline areas. The investigation is correlated across the full data set of ultrasound images, to determine the indicator which gives the best overall test. The final part of the chapter compares the optimized algorithm with previous outlining algorithms presented in the literature.

Chapter 7, which begins on page 235, begins with an explanation of the adjustments to the TURP training system equipment, and the additional components required, to implement a full scale operating theatre monitoring system for TURP. The results of initial experiments within the operating theatre are presented, and the difficulties highlighted which prevent the achievement of this objective. The remainder of the chapter focuses on a laboratory mock-up of the TURP monitoring system which enables the investigation of ultrasound imaging as a potential solution for tracking the prostate, both in the monitoring system and in the training system where registration between prostate phantom and model needs

improvement. The chapter concludes with the results and discussion of experiments conducted using the lab-based system.

The final chapter, (chapter 8 on page 291), discusses the overall conclusions for the work presented in this thesis, and provides directions for further work in the field.

Chapter 2

Background

2.1 Introduction

This chapter begins with the history of minimally invasive surgery and surgical training. A review of the literature in two categories of minimally invasive surgery training simulators is then given, highlighting the benefits of each and the potential for a new hybrid approach. The concept of a minimally invasive surgery monitoring system is introduced as a bridge between training and a real surgical procedure in the operating theatre. Fundamental components of the training and monitoring systems are examined and suitable technologies are discussed.

2.2 Minimally invasive surgery

The beginnings of minimally invasive surgery [135] (MIS) are rooted in the early 19th century when Philip Bozzini invented a light conductor called the “Lichtleiter” (figure 2.1) to illuminate the internal cavities of the body using a candle as a light source. However, medical rivalries and politics in Vienna at the time meant that Bozzini never had the chance to use his creation. Not until fifty years later was the Lichtleiter finally used by Antoine Jean Desormeaux for urological cases. Not surprisingly, burns were a common complication due to the source of illumination. Following Edison’s invention of the light bulb, Maximilian Nitze created the first optical endoscope using electricity as the illumination source, again only using it for urology.

George Kelling was the first surgeon to use an endoscope to explore the abdominal cavity. He invented a system of pumping filtered air (figure 2.2) into the




Figure 2.1: Bozzini's Lichtleiter. The first endoscope.




Figure 2.2: Kelling's "luft tampanade" for insufflating the abdominal cavity.

abdominal cavity of dogs with the objective of stopping intra-abdominal bleeding. Although the technique was not successful for preventing blood loss, the "coelioskope" became the precursor to modern laparoscopy.

In 1929, the German gastroenterologist Heinz Kalk developed a 135° degree lens system to provide a wider view, and the use of two trocars to allow the first surgical procedure in the abdomen using the minimally invasive technique. Ten years later he had performed over 2000 diagnostic liver biopsies without mortality. John Ruddock devised the first endoscope with built in forceps in 1934. He also added an electrocoagulation capability to allow the cauterization of internal bleeding.

During the mid 20th century many advances were made in the field with improved techniques and instruments, such as the rod lens and fibre-optics. In 1980, Patrick Steptoe became the first surgeon to conduct laparoscopic procedures under sterile conditions in the operating theatre (OT).

Up until 1982, the main limitation of MIS was the necessary use of one of the surgeon's hands for manipulating the endoscope, leaving only one hand to perform the surgical task. This restricted the operations which could be carried out to a few gynaecological, arthroscopic (joint) and cystoscopic (bladder) procedures.

The surgeon also had to operate under uncomfortable conditions, bending over to peer down the endoscope. The invention of the solid state camera opened up the potential of minimally invasive surgery, freeing the surgeons hands by allowing a second person to control the view of the surgical field. Further advances in rod lens technology and the introduction of the halogen light source were also significant advances.

The benefits of smaller scars, lower amounts of pain and shorter recovery periods made minimally invasive techniques appealing to patients. Largely driven by this patient demand, many procedures have now been adapted to the minimally invasive approach. Surgical specialities such as urology, gynaecology, thoracic surgery, sinus surgery and orthopaedics have been conducted without the large open procedures necessary in the past, and minimally invasive is now the preferred technique in many cases. As technology develops and new instruments are introduced, the field of minimally invasive surgery will continue to grow rapidly.

2.3 Surgical training

The development of MIS obviously requires trainee surgeons to acquire the necessary skills to become proficient in the new techniques. Historically, training in surgery has followed the apprenticeship model where education is obtained in the OT, with the trainee observing the progress of an operation undertaken by a skilled surgeon. After a certain period of passive observation, the trainee is allowed to perform the operation under close supervision, but without any transition period between observation and practice. For many reasons, the OT is unsuitable as a classroom for learning surgical skills. For the trainer, the priority in the OT must always be the patient's welfare with training taking a secondary role, with the result that the opportunity for teaching is underused [190].

The unpredictable nature of surgery means that training cannot always be well structured and predicted. Teaching of a particular procedure relies on a patient presenting with symptoms which require that procedure. This can lead to a situation where the trainee is not well suited to the case in hand, where either the technical requirement is too demanding or falls below her abilities. Since the patient is the priority, the operation cannot be altered to achieve educational goals. Tasks cannot be repeated and the consequences of failure are drastic.

The OT is a stressful environment. A number of factors contribute to a high level of stress in the OT, *e.g.* time constraints, technical difficulties, concern for the

patient, equipment failures, interpersonal issues, and the handling of telephone calls about other patient issues. The ability to learn has been shown to be reduced under such stressful conditions [192]. OT time is also extremely expensive, and on-the-job training extends average surgery time by 20 minutes.

Training for MIS makes the traditional apprentice approach even more difficult for both trainee and trainer, since the visual field of the operation and the physical interaction with the operative site are no longer in same viewing frame. The trainee must try to interpret the movements of the trainer's hands when manipulating the surgical tools and relate these to what is happening on the video display of the operation. This makes it extremely difficult for the trainee to obtain an understanding of the interaction between the surgeon's hand movements while operating the surgical tool, the active end of the surgical tool, the surgical site and the video display. The trainers themselves may be sometimes unaware of the problem that the trainee can encounter, and also have difficulty translating their knowledge into instructions which the trainee can follow.

Minimally invasive surgery requires an ability to visualize a 3D mental picture of the surgical anatomy from a combination of theoretical knowledge, experience, preoperative examinations and imaging, and a localized view of the surgical site from the endoscopic camera. The fulcrum created at the point of entry to the body results in a reverse relationship between the movement of the working end of the surgical tool and the surgeon's hands. The surgeon must be sufficiently skilled so that the instruments can be wielded dexterously without the surgeon's concentration wandering from the endoscopic view.

Although experience gained by training in the OT can never be fully replaced, much of the manual dexterity required for handling minimally invasive tools can be learnt prior to exposure on a patient. Training using live animal specimens is one approach used outside the OT but the considerable cost, ethical and legal issues, and differences to human anatomy mean frequent use of animals is not a solution. Human cadavers are used, but procedures can only be conducted once on each.

2.4 MIS training systems

Several inanimate systems are available for training in MIS outside the OT without the need for the presence of a patient or experienced surgical trainer. They range in complexity from simple object manipulation exercises which teach the




Figure 2.3: Simple laparoscopic suturing trainer. a) The training box. b) Simulated tissue and grippers. [194]

basic skills required to handle the surgical instruments proficiently, through to complete simulations of surgical procedures. The systems fall into two categories: physical models and virtual reality.

2.4.1 MIS physical model training systems

The simplest products in this category provide training in techniques such as cutting and suturing using endoscopic tools without direct visual guidance, such as the Simuvision¹ suture trainer (figure 2.3). The trainer consists of a box with holes for access of the instruments, a simulated tissue pad and a set of mirrors to provide an endoscopic view without the need for a real endoscope. More specific surgical tasks may be practiced on equipment such as the resection trainer² shown in figure 2.4.

Various anatomical models are available commercially which allow training in particular surgical applications, such as the sinus trainer³ (figure 2.5) which permits the internal passageways of the nose to be explored endoscopically as well as via direct viewing, although it is not designed to allow practice of surgical procedures. However, several more sophisticated products are available which do provide the ability to practice complete surgical procedures. For arthroscopic

¹Simulab Corporation, Washington, USA.

²Limbs and Things Ltd, Bristol, UK.

³Global Technologies, Florida, USA.




Figure 2.4: Basic training in resectoscope manipulation and co-ordination. [131]




Figure 2.5: Sinus endoscopic trainer. [78]




Figure 2.6: a) Arthroscopic shoulder trainer. b) Knee model. c) Endoscopic view of the knee. [95]



Figure 2.7: Torso and internal models for laparoscopic surgery training. [194]

procedures, Hillway Surgical⁴ manufacture a knee and shoulder model (figure 2.6) Both include accurate anatomical features and have sacrificeable, replaceable ligament components upon which the surgery is performed.

Several specialized training systems for soft tissue surgery are available on the market. Figure 2.7 shows a torso model⁵ which can be fitted with various replaceable internal models to allow practice of laparoscopic procedures such as cholecystectomy⁶ or the gastro-oesophageal reflux⁷ procedure. MIS gynaecological and urological techniques may be performed on the female pelvic trainer

⁴Hillway Surgical Ltd, West Sussex, UK.

⁵Simulab Corporation, Washington, USA.

⁶Cholecystectomy is the MIS technique for removal of the gall bladder.

⁷Gastro-oesophageal reflux is the back flow of stomach acid into the oesophagus, commonly known as heartburn. In severe cases, surgery may be required.




Figure 2.8: a) The Limbs and Things female pelvic trainer. b) The TURP training system in use. c) The resectable prostate model. [131]

and TURP trainer respectively (figure 2.8), both produced by Limbs and Things Ltd⁸ [196, 40]. The TURP trainer forms the basis of the CASTS system which is described in section 2.5 and chapter 3.

2.4.2 MIS virtual reality training systems

Virtual reality is the simulation of a real or imaginary environment generated using computer hardware. In its simplest form, the user experiences the simulation as a 3D image displayed on a 2D computer screen with additional components such as sound and tactile feedback to enhance the sense of realism. More complex systems use hardware such as stereo vision display goggles, powered gloves or even full body suits to provide a complete immersive environment. Much of the technology associated with virtual reality has resulted from the lucrative computer gaming market, where significant advances have been made in 3D rendering technology such as OpenGL [223].

The use of VR systems as an aid to teaching is an attempt to bridge the gap that exists between simplistic “lunch box” trainers and training on patients. The additional realism that results is considered to justify the increased cost and complexity of VR training systems in comparison to the lunch box type.

Virtual reality has many uses in the surgical arena [189]. During surgery planning, preoperative images may be rendered in 3D to allow the surgeons to explore

⁸Limbs and Things Ltd, Bristol, UK.

the surgical site, obtain an understanding of the anatomy and highlight areas of potential difficulty. Augmented reality systems use virtual reality to overlay a computer generated scene onto the surgeon's normal view, either direct or endoscopically. Information such as anatomical labels or the area of disease within an organ can be displayed to the surgeon during the course of the operation. [209].

Virtual reality as a tool for training has already found widespread use in the teaching and continued assessment of aircraft pilots. Flight simulators allow pilot candidates to spend many hours perfecting their skills without the danger associated with actually flying a craft, and also provide the ability to train for situations which would not occur outside of an emergency situation in real life. Experienced pilots must undergo a series of annual checks in a flight simulator to ensure that they still have the necessary skills for safe flying.

Several systems for training surgeons in MIS are currently available commercially or under development at research institutions. The ProCedicus MIST⁹ [147] system consists of a pair of laparoscopic instruments (figure 2.9b) which are fitted with encoders to monitor the tools' positions in five degrees of freedom (DOF). Positional information is relayed to a computer which updates the display accordingly. Initial versions of the system allowed the user to perform tasks such as picking up objects and placing them at certain positions around the screen to provide training in general manipulation of the laparoscopic tools. The graphical display was simple, based around spherical or cuboidal objects, such as shown in figure 2.9a) where the objective is to hold the tool tip position on the cubes for a certain period, to simulate electro-coagulation tasks. Recently, a more sophisticated visual interface has been developed which allows practice of surgical tasks such as anastomosis (figure 2.9b). A similar simple system is manufactured by Surgical Science¹⁰ [120].


More specific surgical procedures such as arthroscopy have also been simulated using virtual reality [156]. Figure 2.10 shows the knee arthroscopy simulator jointly developed by the American Board of Orthopedic Surgeons¹¹ and Boston Dynamics Inc.¹² The simulator consists of a physical leg model into which a mock-up arthroscope camera is inserted. The camera is motion tracked so that a VR display simulating the arthroscope view is updated on the computer screen. The leg model is flexible and also tracked so that manipulation of the knee causes

⁹Mentice Medical Simulation, Gothenburg, Sweden.

¹⁰Surgical Science Ltd, Gothenburg, Sweden.

¹¹American Board of Orthopaedic Surgeons, Pennsylvania, USA.

¹²Boston Dynamics Incorporated, Massachusetts, USA.



(a)



(b)

Figure 2.9: a) Graphical display from the MIST VR training system. One instrument is holding the blue sphere while the other coagulates the red cube. b) The laparoscopic instruments and anastomosis repair simulator. [149]




Figure 2.10: a) Boston Dynamics' knee arthroscopy simulator. b) Virtual reality generated endoscope display. [23]

corresponding changes on the VR display. Haptic feedback is provided enabling the trainee to feel the interaction of the arthroscope and knee structure. However, the system is not so advanced as to allow surgical changes to be performed and remains a purely exploratory teaching aid. Similar systems are being researched by groups at Hull University¹³ [221], Sheffield University¹⁴ [143] and Mitsubishi¹⁵ [77] who have incorporated additional DOF so that forces are sensed along the length of the arthroscope rather than just at the tip. Prosolvias Clarus¹⁶ manufacture an arthroscopic system which uses a shoulder model rather than the knee.

The Madigan¹⁷ Endoscopic Sinus Surgery (ESS) simulator [219] uses a VR model of the nasal sinus derived from the Visible Human Database of the National Library of Medicine¹⁸. The system includes haptic feedback and soft tissue deformation, but does not allow surgical interventions to be practiced at present since the primary goal is to teach navigation. Future versions of the system will include the ability to perform procedures. The ESS simulator is shown in figure 2.11.

More complex VR MIS simulators provide the capability to perform complete

¹³Simulation and Visualization Research Group, Department of Computer Science, University of Hull, UK.

¹⁴Virtual Reality in Medicine and Biology Group, Department of Medical Physics and Clinical Engineering, University of Sheffield, UK.

¹⁵Mitsubishi Electric Research Laboratories, Massachusetts, USA.

¹⁶Prosolvia Clarus, Gothenburg, Sweden.

¹⁷Human Interface Technology Laboratory, University of Washington, Washington, USA.

¹⁸National Library of Medicine, Maryland, USA.




Figure 2.11: ESS VR sinus simulator. [96]

surgical procedures. The Virtual Endoscopic Surgery Training¹⁹ [31] (VEST) system allows practice of several minimally invasive laparoscopic techniques such as cholecystectomy and gynaecological surgery. The system comprises a hollow model of the abdomen into which mock tools may be inserted. The positions of the tools within the abdomen are monitored through a mechanical system which is fitted with potentiometers to measure the position of each degree of freedom. Position information is conveyed to a PC which updates a virtual model of the tools and anatomy accordingly, and displays the resulting endoscopic view. Sophisticated deformable models and collision detection are used to allow the user to manipulate the surgical anatomy and perform tasks such as stapling, cutting and suturing of tissue. The abdomen and mock tools are shown in figure 2.12a), while figure 2.12b) shows the endoscopic view of the gynaecological model. Several other centres are also currently researching VR for training in laparoscopic techniques. They include the Virtual Environments for Surgical Training and Augmentation system [210] from Berkeley University²⁰, the VIRGY simulation system for endoscopic training [13] from the Swiss Federal Institute of Technology²¹ and the commercially available LS500 system from Xitact²² [191].

Stanford University²³ has developed two simulation systems based around the

¹⁹Select-IT VEST Systems AG, Bremen, Germany.

²⁰Robotics and Intelligent Machines Laboratory, Department of Electrical Engineering and Computer Sciences, University of California at Berkeley, USA.

²¹Virtual Reality and Active Interfaces Group, Swiss Federal Institute of Technology, Lausanne, Switzerland.

²²Xitact SA, Morges, Switzerland.

²³Biocomputation Centre, Stanford University, California, USA.




Figure 2.12: a) VEST MIS trainer. b) Suturing of the uterus in the gynaecological surgery system. [193]

laparoscopic impulse engine from Immersion²⁴. The first is a training system for hysteroscopy and endometrial ablation [154], while the second provides a simulation of the female pelvic girdle to provide training in various laparoscopic tasks [26]. Both systems incorporate haptic feedback to increase the sense of realism experienced by the trainee. Figure 2.13 shows a rendering of the anatomical model and an example view from the endoscope.

At present no VR training simulator exists for the TURP procedure. Simbionix²⁵ market a urological training system [151] (figure 2.14) which allows practice in several procedures such as ureteroscopy and are currently researching a system for TURP.

2.4.3 VR versus physical models

Whilst significant advances have been and continue to be made in the realism involved with virtual reality surgical simulators, there is still a large gap between the simulated visual environment and reality. Rendering techniques and texture mapping enable static scenes which do resemble reality, however the calculations required to calculate realistic deformations between interacting bodies are too complex to be performed in real time. Simple models have been included in virtual reality training systems, but they do not provide animations which are

²⁴Immersion Medical, Maryland, USA.

²⁵Simbionix USA Corp, Ohio, USA.




Figure 2.13: a) Anatomical model for gynaecology. b) View from the endoscope during hysteroscopy. [16]




Figure 2.14: Simbionix system. [195]




Figure 2.15: a) Cutting an artery in the VEST gynaecological system. b) Simulated pulsatile bleeding is not realistic. [193]

convincing, particularly where the model tries to display cutting, coagulation, suturing and other such surgical interactions with the tissue (figure 2.15).

A large degree of information feedback to the surgeon comes from the tactile sensations of the surgical tools interacting with their surrounds. Incorporating haptic feedback has improved realism in this area and virtual reality systems no longer suffer from the floating feeling associated with systems without force-feedback. Once again however, the calculations required for modelling realistic force reactions are too lengthy to be performed in real time, and even the most advanced systems fall short of giving the same feel as cutting through real tissue.

VR offers advantages over physical model based systems in a number of areas. The surgical tasks may be restarted or repeated as required, without the need for replacing used parts of a physical model. Using a software based model of the anatomy allows several different procedures or disease models to be performed using the same hardware. Running costs are reduced since there are no sacrificial parts of the system, as opposed to physical model simulators. Capital costs may be increased for a VR training system (particularly due to the cost of the haptic systems), however they do not require the use of a full set of surgical instruments, including light sources, cameras and electrosurgical current generators as well as the endoscopic tools. This could be construed as an advantage for the physical model systems since gaining familiarity with the real equipment is obtained prior to exposure in the OT.

One area where MIS trainees have difficulty is maintaining knowledge of the po-

Physical models	Virtual reality
Improved visual realism	Tasks may be repeated or restarted
Realistic tactile feedback	Many operations simulated
Use real surgical instruments	Lower running costs
Capital cost of haptic system not necessary	Improved navigation
	Adjustable difficulty level
	Skill assessment

Table 2.1: Advantages (current and potential) of VR based versus physical model based MIS training systems.

sitions of the surgical tool with respect to significant anatomical regions. The VR systems described in section 2.4.2 all provide a simulated view of the endoscope, which does not provide the trainee with any more navigational information than she would obtain from real surgery or a model based training system. However, VR reality models have the capability to allow the “camera” to be placed anywhere which could be exploited to provide an overall view of the surgical anatomy and the relative position of the surgical instruments. This is a failure of all currently available VR MIS training systems, but one which could easily be remedied. Another potential feature of VR systems is the ability to adjust difficulty levels from straightforward cases through to surgical situations involving significant complications, including problems which arise due to differences in anatomy/pathology.

One other area where VR offers significant benefits is in its potential for assessment of surgical skill which is discussed in detail in section 2.4.4. Table 2.1 gives a summary of the potential advantages of VR MIS training systems or physical model based simulators.

2.4.4 Assessment of surgical skill

In the current climate of medical liability with patients having ever increasing expectations upon the outcome of medical procedures, there is significant pressure on surgeons to demonstrate that they can perform operations to a high standard, consistently and with acceptable results. Various high media profile cases in the UK such as the Bristol heart surgery affair [68] have resulted in legislation requiring outcome data collection after every surgical procedure. Such data may be used to assess and rank the performance of an individual surgeon against his

peers, however, statistically significant poor performance will only be determined after a number of failures, with potentially dire consequences for the patients concerned in the meantime.

Spencer [198] estimates that the expertise required to conduct a successful operation is based 75% on correct decision making and 25% on the motor skills of the surgeon. For MIS, manual dexterity is considerably more important [52]. Surgeons undergo rigorous examinations to test their knowledge on surgical techniques and operative procedures, but at present there is no requirement for a surgeon to demonstrate competency in the complex motor skills necessary for successful operations. One of the difficulties in assessing the dexterity of a surgeon is the lack of an obvious criteria to use as a performance indicator. Various methodologies have been adopted, such as the assessment by experienced surgeons of endoscopic videos of trainees performing the operations [144]. Although this method was shown to perform reliably, it required assessment of each trainee by multiple experts to eliminate observer bias. Providing such assessment, outside the scope of a research project, for general appraisal of trainee surgeons would impose severe demands on the time of already stretched consultants. Other studies have attempted to evaluate surgical skill based upon the time required to complete a particular procedure [180], or on the total number of procedures undertaken [70], however neither may be considered a reliable indicator; faster operations are not necessarily beneficial and the total number makes no judgement on the operations' successfulness.

Virtual reality training systems offer the potential for an objective assessment of surgical skill, since all movements of the surgical tools performed by the trainee may be easily logged. Success in achieving surgical goals may be monitored and performance improvement over time evaluated. Taffinder et al [205, 206] have measured performance based on a number of factors such as the number and economy of movements whilst performing specific tasks, using the MIST trainer system described in section 2.4.2.

One system in the physical model based category has been used to assess surgical performance - the ENT (ear, nose and throat) training system [229] researched by the Japanese Institute of Advanced Industrial Science and Technology²⁶. This trainer uses a model of the head to allow sinus endoscopic training. The model is fitted with sensors to measure the forces applied by the endoscope to the tissue. The study shows that novices perform the training session using decreased levels

²⁶Institute for Human Science and Biomedical Engineering, National Institute of Advanced Industrial Science and Technology, Tsukuba Ibaraki, Japan.

of force after a period of practice, while experienced surgeons perform the task using a constant low level of force. However, the trainer only simulates endoscopic exploration rather than surgical intervention.

Further work is required in the area of MIS training to pinpoint the best methods for assessment of skill and define standards for performance evaluation.

2.5 CASTS

Current systems in the field of MIS training are based upon either VR or physical models of the surgical anatomy. However, there are several disadvantages associated with each approach as described in section 2.4.3. This has led to the adoption of a hybrid approach to MIS training simulators, incorporating the benefits of both types. The general hybrid training system is named CASTS (computer-assisted surgery training system) and the prototype model provides training in the TURP procedure. TURP has been described as the most difficult operation to learn and teach [100], therefore it provides a demanding test of the CASTS approach. The CASTS-TURP system (referred to simply as CASTS throughout the remainder of the thesis) is described fully in chapter 3.

CASTS uses a physical model of the prostate²⁷ to provide realistic tactile feedback during resection using the same surgical equipment and procedure as used in the OT. The endoscopic view seen by the trainee is the internal view of the physical model at the tip of the resectoscope. CASTS provides an additional VR display showing an external view of the physical model and the position of the resectoscope in relation to it. Objective data such as the amount of tissue resected, perforation or damage of no-go areas and time taken can be logged for use in assessing performance. These additional features of the hybrid system all require tracking of the resectoscope to provide the relative position of the tool with respect to the prostate model.

2.5.1 Tracking technologies

Tracking a rigid body in space is the basis of the CASTS system. The position of a point in space can be defined by three coordinates. However, the orientation of the rigid body must also be considered, so an additional three values – the

²⁷Limbs and Things prostate phantom - see figure 2.8c).

angles of rotation around each axis of the coordinate system are required. There are a number of factors to be considered when selecting an appropriate tracking technology for use in CASTS:

- **Timeliness** - CASTS must show the position and status of the resectoscope in relation to the prostate in real time, hence the latency of the tracking system should be less than the frame rate of the VR display (25 frames per second).
- **Accuracy** - There are several potential sources of error, such as manufacturing variability in the physical model, bending of the resectoscope cutting loop, registration between the physical model and the tracking system and deformation of the model during resection. Errors in readings from the tracking system should be minimized and a threshold of sub-millimetre accuracy has been selected as an acceptable level [39].
- **Reliability** - The output of the tracking device must not drift over time and should not be affected by surrounding conditions and equipment.
- **Working volume** - The system must be capable of tracking the resectoscope through the full working volume. During TURP the resectoscope is constrained to move within a volume of less than 500 cm^3 , however the rotation around the resectoscope shaft can exceed 360° .
- **Ease of interface** (with the computer and surgical equipment) - The tracking system must be easily integrated with the other components of the CASTS system.
- **Ergonomics** - Any fixture attached to the resectoscope should not significantly affect the feel of the tool or change the operation procedure.

Several tracking technologies are available which could be used for obtaining tool position and orientation information in CASTS. They are briefly reviewed over the proceeding sections.

2.5.1.1 Acoustic devices

The principle behind acoustic tracking is to measure the time of flight of an ultrasound pulse from an emitter to a receiver. For full tool position and orientation information, three acoustic emitters must be attached to the resectoscope and

monitored by three acoustic receivers. The emitters are small and light in weight, so attachment to the resectoscope is possible without adverse effects, however, the positional information is affected by atmospheric conditions where variations cause changes in the speed of sound. Acoustic tracking systems also require a clear line of sight between emitter and receiver and are easily deceived by reverberating echoes and are therefore not considered appropriate for tool tracking in CASTS. An example application of ultrasound tracking is the Natural Writing Board from Pegasus Technologies²⁸.

2.5.1.2 Electromagnetic devices

Electromagnetic trackers comprise an emitter which is attached to the tracked object, and a receiver, both of which contain three orthogonal coils. The emitter transmits an electromagnetic field which is detected by the receiver. Variations in the received signal from each coil in the receiver allow the calculation of six DOF for the emitter. The size and weight of the emitters are low enough to allow them to be fitted to a resectoscope and no direct line of sight to the receiver is required, however electromagnetic tracking systems are susceptible to interference which causes erroneous readings. CASTS uses a diathermy unit for resection which generates significant electromagnetic interference.

Two commercially available electromagnetic tracking systems are the Fastrak system from Polhemus²⁹ and the Flock of Birds from Ascension Technology³⁰. The latter claims to be less sensitive to interference from surrounding metal since it uses pulsed d. c. magnetic fields rather than a. c. Both systems have been used in surgical simulators. However, neither system is considered reliable enough for use in the OT [17] (tracking in the OT is required for CASMS, described in section 2.6).

A new product which could improve the potential for electromagnetic tracking in the OT is the Aurora from Northern Digital³¹. The Aurora system is capable of tracking miniature coil trackers which are small enough that they may be included in the tip of flexible endoscopes, giving an improved positional accuracy. The manufacturers also claim that the system is capable of detecting electromagnetic interference, and warning of its presence. However, at the time of writing the system is not commercially available.

²⁸Pegasus Technologies Ltd, Holon, Israel.

²⁹Polhemus Incorporated, Vermont, USA.

³⁰Ascension Technology Corporation, Vermont, USA.

³¹Northern Digital Inc, Ontario, Canada.

2.5.1.3 Inertial devices

Intersense³² produce an inertia based tracking system using piezo-electric transducers. The manufacturers claim accuracy in the order of a millimetre, however inertial systems are susceptible to position drift over time, and the transducer size is too large for easy attachment to a resectoscope.

2.5.1.4 Mechanical devices

Mechanical devices are used in the majority of MIS simulators, since they are low-cost, insensitive to their environment and can provide haptic feedback. The normal configuration is a linked arm system, with a potentiometer or encoder arrangement measuring the position of each joint. Popular mechanical tracking devices are manufactured by Immersion³³ and SensAble Technologies³⁴ and products from both companies have been used in surgical training simulators. However, use of a mechanical tracker for CASTS would alter the feel of the resectoscope and hinder its movement, therefore the devices are not suitable.

2.5.1.5 Optical devices

Optical tracking systems use an array of two or more cameras to calculate the positions of visible markers. Some systems use passive markers which have the advantage of not requiring a direct wire link to the tracking system, but the disadvantage is that passive markers are not uniquely distinguishable. The second class of optical trackers use active markers, normally infra-red emitting diodes (IRED). Active optical tracking systems, such as the Optotrak³⁵ and Flashpoint³⁶, are commonly used in the OT for image-guided surgery applications [130], since they have high accuracy, good reliability and timeliness, and are largely unaffected by environmental factors such as electromagnetic interference. However, two drawbacks are the requirement for direct line of sight between the markers and the cameras, and the higher cost in relation to other tracking technologies.

For CASTS the adopted tracking solution is a Flashpoint 5000 optical tracker from Image Guided Technologies. Optical tracking was chosen for its superior

³²Intersense Inc, Massachusetts, USA.

³³Immersion Medical, Maryland, USA.

³⁴SensAble Technologies Inc, Massachusetts, USA.

³⁵Northern Digital Inc, Ontario, Canada.

³⁶Image Guided Technologies Inc, Colorado, USA.

accuracy, reliability and timeliness and easy interface. Having selected an optical tracker, cost dictated this particular manufacturer.

2.6 MIS monitoring - CASMS

The term *monitoring*, when applied in a surgery context, usually refers to the assessment of patient condition during the course of an operation. Surgical monitoring equipment measures physiological signals such as pulse and breathing rates, blood pressure and anaesthetic levels in the blood stream. However, in the context of this thesis, monitoring refers to the assessment of the surgical task rather than the status of the patient.

The MIS training systems discussed in section 2.4 allow trainees to practice surgical tasks and acquire necessary skills before exposure to a real patient in the OT. However, there is a significant gap at present between performing on a simulator and the real surgical procedure. The major motivation behind the work in this thesis is to bridge this gap by providing an in-vivo monitoring system which provides additional information to the surgeon, in much the same way as the CASTS system gives additional help when training. The system is referred to as CASMS (computer-assisted surgical monitoring system), and the concept has been investigated using the TURP procedure as a prototype target operation. The CASMS-TURP concept is referred to as CASMS throughout the remainder of this thesis.

There are several potential benefits from using CASMS during TURP. The fundamental role of CASMS is to monitor the relative positions of the surgical tools and the anatomy, and provide warnings on proximity to surgical “no-go” areas such as the verumontanum and prostate capsule. Additional potential functions include monitoring and display of the progress of the surgery, indicating the extent of the prostate resection and highlighting regions requiring additional tissue removal. The system could also provide a navigational view indicating the relative position of the resectoscope within the prostate. Finally, the system could provide a post-operative record of the surgery which could potentially be used in any litigation proceedings which may result following problems in the surgery.

CASMS requires two pieces of information, namely knowledge of the position of the resectoscope and position of the prostate. Various technologies were discussed in section 2.5.1 for resectoscope tracking in the CASTS training system. CASMS

has the same requirements for tracking with the additional necessity that any elements attached to the resectoscope are sterilizable.

2.6.1 Intra-operative imaging

The VR model used for CASTS is a computer generated version of the physical Limbs and Things prostate model. However, the VR model required in CASMS must be specific to the anatomy of the patient. In addition, any changes which occur to the patient anatomy must be updated on the corresponding VR model. For CASMS, these changes comprise the resection of the prostate during the course of the surgery, and movement and deformation of the prostate under the action of the resectoscope.

Generating models of patient anatomy from preoperative images is a common task in many areas of computer assisted surgery. For example the PROBOT [4, 54, 55, 157] TURP robotic system uses a set of hand delineated prostate outlines from preoperative³⁷ transurethral ultrasound images, to generate an allowable cutting zone for the resection apparatus. The ACROBOT [57, 89, 90, 91, 97, 98] system uses pre-operative computed tomography (CT) images to build a model of the patient's fibula and tibia from which the active constraint [56] regions are derived. These regions define the allowable cutting zone of the robot when milling the bone ends to take prosthetic knee components. Several other robotic surgery systems use CT images to generate patient specific anatomical models [7, 201].

For orthopaedic applications, it can be generally assumed that once a computer assisted surgery system is registered to relevant bony anatomy, the pre-operative model will remain valid for the rigid bones throughout the course of the surgery, provided bone motion is prevented. However, for soft tissue applications, such as TURP, no such assumption can be made. One of the difficulties involved with the PROBOT project was that the system could not track any movement or deformation of the prostate resulting from the resectoscope action. In order to provide accurate information to the surgeon, CASMS must have the ability to track prostate motion and deformation through the procedure.

Several imaging modalities are potential candidates for providing intra-operative images. They have all been used in conjunction with MIS, computer-assisted surgery or image guided surgery applications, however, the author was unable

³⁷The transurethral ultrasound scans are obtained in the OT just prior to the surgical procedure to minimize the possibility of the soft tissue deforming or changing position.

to find any references describing intra-operative imaging with image processing designed to gain an understanding of the anatomy without human intervention.

2.6.1.1 Computed tomography

CT scanners consist of an x-ray tube and detector which rotate around the body, capturing projected views of the same anatomical region from different angles. These views are compounded in software to form a complete composite image which shows the anatomical features in a slice through the body. In contrast to conventional x-rays which only show bone clearly, CT scans provide high contrast images of different tissue types and structures in the imaging plane.

Modern ultrafast CT, or electron beam CT scanners use a rotating electron beam in place of the x-ray tube. Since no mechanical parts are involved, ultrafast CT scanners are capable of acquiring images at a frame rate of up to 10Hz. However, there are several disadvantages of CT systems including radiation exposure risks and equipment capital cost. The patient must be positioned inside the bore of the machine during imaging, making access to the prostate difficult – during conventional TURP the patient's legs held open in stirrups to a wide angle.

CT has been used as an intra-operative imaging modality in the Minerva neurosurgery robot [28], which is designed to work within the scanner's bore. However, for the CASMS application CT imaging during TURP surgery is not a feasible option.

2.6.1.2 Interventional MRI

Magnetic resonance imaging (MRI) is a noninvasive diagnostic technique that uses nuclear magnetic resonance to produce cross-sectional images of the anatomy. The patient lies inside a large, hollow cylinder containing a strong electromagnet, and the magnetic field causes alignment in the spins of protons throughout the body. Radio frequency coils around the imaging area emit radio waves which deflect the proton spins within the coil; once the applied radio waves are withdrawn, the deflected protons realign and emit radio waves in the process. These RF waves are detected and are used to form of an image slice through the body which distinguishes tissues types as a result of differences in their proton density.

Interventional MRI scanners use an open bore magnet system which permits access to the patient so that imaging may be performed during surgical procedures.

MRI is a non-ionizing modality and therefore poses no risk to either the patient or the surgeon. The scans obtained provide the highest level of detail among all the radiology imaging systems. However, the extremely powerful magnetic fields mean that metallic objects can be hazardous in the proximity of an MRI scanner, therefore specialized MRI compatible surgical implements are required. MRI systems are also prohibitively expensive and very noisy.

2.6.1.3 Ultrasound

Ultrasound scanners use a ultrasonic transducer to transmit high frequency sound waves into the scanning region of the body. As the waves travel through the body they cross tissue boundaries (such as between tissue and fluid regions) where some energy is reflected back towards the transducer. These echoes are detected by an acoustic receiver and processed into an image based on the time of flight and intensity of the signal.

Ultrasound imaging is commonly used as a diagnostic tool in many areas including urology. The scanners are much cheaper than MRI or CT systems, and are portable, allowing easy entry to the OT. Ultrasound is non-invasive and non-ionizing, and the bulk of equipment can be located at a distance, with only a small probe required to be in the vicinity of the patient. Therefore, ultrasound allows intra-operative imaging with the least disruption to the surgeon of all the imaging modalities. However, the images are poor in quality with low contrast, making automated outlining much more difficult than for CT or MRI.

For CASMS, the adopted intra-operative imaging modality is ultrasound, with the choice dominated by system cost and the minimization of changes to the standard TURP procedure. Several companies produce medical diagnostic ultrasound scanners, the choice for the CASMS project is a B&K³⁸ scanner dictated by the available hardware in MIMLab³⁹ for the project.

2.6.2 Ultrasound speckle

Speckle is a phenomenon observed in all coherent imaging systems such as laser interferometry, synthetic aperture radar and ultrasound. Speckle is a random, deterministic, interference pattern [49] which is the result of the interaction between

³⁸B&K Medical, Copenhagen, Denmark.

³⁹Mechatronics in Medicine Laboratory, Imperial College, London, UK.

the ultrasound beam and small (on the scale of the wavelength) scatterers in the medium being imaged. Many scatterers are present within any resolution cell of the ultrasound system and the echoes from these scatterers are summed by the transducer, therefore a granular texture is apparent in the image as the scattered waves interfere.

Speckle is often referred to as noise, however it cannot be considered noise in the classical sense of signal corruption (such as electrical noise) since it is a fundamental characteristic of the image formation process and without it there would be no image. The problem with speckle is that the observed texture does not correspond to the underlying structure of the tissue, although on a local scale the brightness does reflect the local tissue echogenicity. Goodman [79] describes speckle as 2D random walk with the reflection from each scatterer in the resolution cell representing a step of random magnitude and direction. If a target (variation in underlying structure, such as the interface between two tissue types) is present within the resolution cell, then the random walk is offset by a coherent component with a strength determined by the variation in ultrasound speed across the target. Large velocity changes, such as those which occur at the interface between tissue and fluid, cause a high degree of reflection of the incident ultrasound so that a strong signal is received by the transducer, and the interface shows relatively clearly on the final ultrasound. Boundaries between soft tissues such as across the prostate capsule have lower associated velocity variation, therefore they are less distinguishable in the ultrasound image.

Speckle is responsible for the poor quality observed in ultrasound images. Bamber and Daft [10] have quantified this impact as an eightfold reduction in the detectability of lesions within the imaging region. Such a radical reduction in contrast is responsible for the poorer effective resolution of ultrasound, when compared to MRI or CT. Much research has been conducted into algorithms for reducing speckle in images, however, the price for speckle reduction is loss of fine detail within the ultrasound image.

2.6.3 Ultrasound artefacts

Speckle in ultrasound images causes difficulties because it masks the underlying grey level variation caused by changes in tissue structure, with the result that low intensity edges between tissue structures are difficult to detect. A further problem is caused by the possibility of artefacts - visible structures in the ultrasound

image which do not correspond in location or intensity to real interfaces within the imaging region [119]. Artefacts are inherent to the ultrasound modality and occur due to the underlying physics principles of the imaging system. They can cause confusion for experienced sonographers and can lead to mis-diagnosis [113].

Modern ultrasound systems, which provide direct digital images, have improved image quality. Despite the advances in medical ultrasound technology, the underlying ultrasound image formation process is simple, based on a few assumptions. The ultrasound pulse travels from the transducer into the tissue and reflects back to the transducer from structures within the tissue. The “time of flight” for the pulse is used to calculate the depth of the structure, based on an assumed constant velocity of the speed of sound in all tissues of 1540ms^{-1} . The received echoes are assumed to arise from the centre of the ultrasound beam, and a further assumption is that their intensities are only affected by the size and acoustic properties of the interface.

Any deviation from these assumptions cannot be recognized by the ultrasound machine, causing artefacts to occur in the displayed ultrasound image. Several artefacts, the conditions under which they occur and the likelihood of their appearance in prostate transrectal ultrasound images, are discussed in the following sections.

2.6.3.1 Reverberation

Reverberation artefacts occur when the ultrasound pulse bounces back and forth between the transducer and a strongly reflecting perpendicular interface within the tissue. The first returning echo will show as a strong feature in the ultrasound image. However, some of the returning ultrasound can reflect back off the transducer, into the tissue and be reflected again by the interface. This secondary echo will cause the ultrasound machine to think that there is a second interface at twice the depth of the actual interface. The resulting ultrasound image will show multiple images of the tissue interface at equal spacings through the depths of the tissue. Figure 2.16a) shows an example ultrasound image with reverberation artefacts, with the transducer at the top of the image.

Reverberation artefacts can cause problems when the secondary (or tertiary) reflection is mistakenly thought to have come from a real interface in the tissue, or when the echo hides a feature which does lie at the same location. For TRUS prostate scans, severe reverberation artefacts are unlikely because there are no




Figure 2.16: Ultrasound artefacts.

large reflecting interfaces between the prostate and the transducer. However, smaller artefacts may occur if hard calcifications are present within the prostate tissue. The resectoscope will also cause reverberation artefacts, but the knowledge of both the resectoscope and ultrasound position provide sufficient information on the intersection of the ultrasound beam and resectoscope, such that these artefacts can be identified and ignored if necessary.

2.6.3.2 Mirror images

Mirror image artefacts are similar to reverberation except the additional echo comes from a reflecting interface within the tissue, rather than from the transducer. Mirror artefacts can occur along the line of the ultrasound beam, however they commonly occur at an angle and can have a different size and shape depending on the angle and curvature of the reflecting interface. Figure 2.16b) shows a mirror image artefact which results from the angled reflecting plane within the imaging region.

Mirror image artefacts can occur in TRUS prostate images due to the large acoustic difference at the bladder wall between the tissue and urine. However, the resulting false targets lie inside the bladder region, so do not cause a problem in terms of prostate boundary identification.

2.6.3.3 Ring down artefacts

Ring down artefacts show as a line which runs distal to a strong reflecting object in the imaging region, along the path of the ultrasound beam. An example of

ring down artefact is shown in figure 2.16c). Ring down artefacts are unusual in ultrasound images because they require a very strong acoustic reflector in the image region, such as a surgical clip or lead shot, although a less pronounced effect can occasionally occur due to the presence of gas [236]. The artefact takes its name from the similarities with striking a bell – the ultrasound beam causes the object to resonate.

Ring down artefacts will occur in CASMS due to the presence of the resectoscope in the imaging region. Any part of the prostate boundary which lies directly behind the resectoscope will not be visible in the image. However, the resectoscope position is known, and the ring down artefact does not increase in width behind the resectoscope (as can be seen in figure 2.16c) therefore ring down artefacts are not likely to result in false boundary detection. Ring down does not occur at calcifications [44] which can be present within the prostate.

2.6.3.4 Shadowing and enhancement

Shadowing and enhancement artefacts occur due to the variation in expected attenuation of the ultrasound beam as it passes through tissue. All tissues cause some degree of attenuation, therefore the ultrasound machine uses time gain compensation (TGC) to adjust the image brightness according to the depth in the image. For example, when a uniform medium is viewed using ultrasound, the TGC results in a uniform image even though the echo signal from deeper regions is weaker. However, if the expected degree of attenuation does not occur then the echo from a deeper structure will be incorrectly over or under compensated. Enhancement occurs when a region of low echogenicity (e.g. if a fluid filled cyst) lies between the transducer and a target, causing the target to appear brighter on the final image. Shadowing results when a higher echogenic region lies before the target.

Enhancement artefacts can occur in TRUS images, since fluid filled cysts may sometimes be present within the prostate. Shadowing can occur due to large calcifications within the prostate, causing loss of boundary information. Calcifications which are smaller than the beam width do not cause shadowing since parts of the beam can pass around the stone.

2.7 Conclusions

This chapter has discussed the various fields relevant to the technology of training systems for minimally invasive surgery. It has reviewed existing systems, both under research and commercially available, and grouped them into two categories according to the type of anatomical model at the heart of each simulator, namely virtual reality or physical model. Current and potential advantages of each type have been presented. A new hybrid MIS trainer architecture has been introduced which combines the advantages of both categories. Transurethral resection of the prostate has been introduced as a specific MIS application suitable for the implementation of a prototype hybrid training system, and forms the basis of the work presented in chapter 3.

A novel extension of the trainer has been proposed which bridges the current gap between surgical training systems and real surgical procedure. This *in vivo* monitoring system will use intra-operative ultrasound scans of the patient's prostate to create a virtual reality model of the specific 3D anatomy. Intra-operative ultrasound would then allow the condition of the prostate to be monitored during the procedure, and the VR model to be updated accordingly. Both formation and update of the model require knowledge of the prostate ultrasound boundary in the ultrasound image to be obtained automatically. The remainder of the work in this thesis is motivated by the search for an automated image processing algorithm for detecting the prostate outline and thus enable the tracking of prostate movement during TURP.

Chapter 3

CASTS TURP training system

3.1 Introduction

This chapter describes the CASTS training system which aids the teaching of surgeons in the techniques of TURP. The system is based on a commercial training aid, but includes computer VR displays to provide additional guidance to the trainee. The various components which form CASTS are described in section 3.2. Initial trials of the system are described in section 3.3 which show the problems of mismatch between the modelled operation and reality. First attempts at solving the discrepancies are described and shown to fail. Section 3.4 describes potential solutions and highlights ultrasound imaging as the most promising method. Finally, section 3.5 gives conclusions for the chapter.

3.2 System components

The CASTS system uses a hybrid approach to provide the most realistic setup. Resection is performed on a physical prostate phantom which can deform and be resected in a similar fashion to human prostate tissue. A VR display provides positional feedback with a 3D rendering of the phantom prostate overlaid with the real-time position of the surgical tool. In addition to aiding navigation, the system can also act as a watchdog, providing audio and visual warnings when the resection path approaches critical areas of the prostate, such as the verumontanum and prostate capsule. Figure 3.1 gives an overview of the CASTS components and figure 3.2 shows the CASTS hardware.




Figure 3.1: The CASTS demonstrator for TURP.




Figure 3.2: The CASTS system hardware. A – Flashpoint base unit. B – CASTS PC. C – Diathermy unit. D – Camera base unit. E – Xenon light source. F – CASTS user interface monitor. G – Endoscopic image monitor. H – Resectoscope and ring. I – Mock-up abdomen and prostate phantom. J – Flashpoint camera.




Figure 3.3: Resectoscope assembly ready for use.

3.2.1 Surgical equipment

CASTS uses a complete set of surgical equipment, identical to that required for the TURP procedure in the OT. This allows the trainee to become fully adept in the use of the equipment before performing an actual procedure. The surgical components are described over the following sections, and the complete resectoscope with all components connected is shown in figure 3.3.

3.2.1.1 Resectoscope

The surgical tool used for CASTS is a standard Karl Storz¹ 26 French resectoscope. The resectoscope comprises an outer sheath, housing the working element which has a trigger mechanism allowing the retraction of the resectoscope cutting electrode. Various electrode designs are available as shown in figure 3.4; experimentation on the prostate phantom showed that the wire loop electrodes provide the best cutting performance. The working element has a hollow shaft which contains a Hopkins II rod endoscope, also manufactured by Karl Storz.

3.2.1.2 Diathermy system

The diathermy unit provides current to the resectoscope electrode. The diathermic current is activated by stepping on foot switches; one switch delivers a higher current for cutting tissue while the other is used when coagulating bleeding arteries². A wire connects the diathermy to a terminal on the resectoscope working

¹Karl Storz GmbH & Co. KG, Tuttlingen, Germany.

²Only the cutting pedal is used for CASTS.




Figure 3.4: Resectoscope electrodes including cutting loops, coagulating ball-end and vaportrode versions.




Figure 3.5: a) Diathermy foot switch controls. b) Diathermy base plate.

element, and the circuit is completed through a plate electrode which is inserted into the base of the prostate phantom. The diathermy unit used for CASTS is a Pfizer Valleylab Force 40 electrosurgical generator and foot switch³. The diathermy unit is shown in figure 3.2(C)), while the foot switch and plate are shown in figure 3.5.

3.2.1.3 Camera system

Endoscopic images are displayed to the surgeon on a video monitor located above the prostate phantom box. The images are captured by a Karl Storz endoscopic

³Valleylab Division, Tyco Healthcare Group LP, Boulder, Colorado, USA.

camera which attaches to the rear of the endoscope, and plugs into a Telecam base unit (shown in figure 3.2(D)). The surgical scene is illuminated by a xenon light source (figure 3.2(E)) which attaches to the top of the endoscope through a fibre optic cable.

3.2.1.4 Irrigation

Inflow and outflow ports on the resectoscope body allow delivery of irrigant to the surgical site. The irrigant prevents charring of the tissue (or phantom) and transports chips of resected material away from the cutting loop. During TURP the irrigant can be pumped through the resectoscope to further aid the removal of chips; CASTS uses a more simple gravity feed with drip bags held by a drip stand. The irrigant for CASTS is water.

3.2.2 Prostate phantom

The phantom used in CASTS is a Limbs and Things diathermy prostate, made of a proprietary material which allows resection and provides a similar feel to human tissue. The phantom includes anatomical features of a real prostate, including the lateral and median lobes, urethra, verumontanum, bladder neck and ureteric orifices. The prostatic tissue is surrounded by material of a different colour to highlight the prostate boundary, and define the limits of resection. The diathermy plate fits into a slot in the base of the phantom.

The phantom fits into a latex sleeve which includes a representation of the bladder to provide a reservoir for the irrigating fluid. This is housed in a custom built mock-up abdomen which provides a frame of reference for the tracking system described in section 3.2.3.1. The mock-up abdomen, prostate phantom and latex bladder are shown in figure 3.6.

3.2.3 Instrumenting the tools

To be able to track the progress of the TURP training session, CASTS requires full knowledge of the interactions between the surgical tool and the prostate phantom. Various aspects of the surgical tools are instrumented to provide this information. The predominant criteria when designing the instrumentation have




Figure 3.6: a) Mock-up abdomen with un-sleeved prostate phantom. b) Latex bladder sleeve.

been to minimise the differences when CASTS is in use or not. Several components within the surgical framework are monitored as described in the following sections.

3.2.3.1 Resectoscope tracking system

The tracking system used in CASTS is a Flashpoint 5000 optical tracker from Image Guided Technologies⁴, shown in figure 3.2(A). The system's receiver is a three IRED camera array (figure 3.2(J)) providing a working volume of 0.5m by 0.5m by 1m at a distance of 0.5m below the cameras. The camera array is mounted above the perspex mock-up abdomen so that all movements of the surgical tool during resection are inside the tracking zone. Information from the camera array is fed to a base unit which interfaces to the CASTS PC through an RS232 serial connection. The base unit powers the IRED markers seen by the cameras; Flashpoint has the ability to track up to 128 IRED concurrently. Each attached IRED is fired in sequence so that their positions can be determined uniquely and individually. Flashpoint can report the position of each IRED, or more commonly, the position and orientation of a user-defined point relative to a defined IRED geometry, such as the tip of the resectoscope. The Flashpoint coordinate system is defined relative to a *dynamic reference frame* (DRF) containing three IRED, so that the slight motions of the camera array do not cause errors in the reported IRED positions. The DRF can be seen attached to the mock-up abdomen in

⁴Image Guided Technologies Inc, Colorado, USA.




Figure 3.7: Resectoscope motion.

figure 3.6a).

Flashpoint requires knowledge of the position of at least three IRED to give full resectoscope tracking in six DOF. During the TURP procedure, the resectoscope moves through a large range of motions⁵, including rotation around the major axis of greater than 360° , as shown in figure 3.7. Optical tracking requires line of sight between emitter and receiver, therefore at least three IRED must be visible to the camera system during all resectoscope motions.

Figure 3.8a) shows the geometric arrangement of the IRED for the resectoscope ring tracker (shown in figure 3.8b). The ring has a diameter of 100mm and contains 12 IRED, equally spaced around the circumference. The ring is clamped to the resectoscope working element in between the irrigation inlet and outlet ports. Careful design considerations have reduced the impact of the ring on the “feel” of the surgical tool, and it does not restrict the surgeon’s movement or interfere with the procedure.

3.2.3.2 Cutting loop position

The most crucial knowledge required by CASTS is the position of the cutting loop with respect to the prostate phantom, since this is the point at which resection occurs. The cutting loop can move with respect to the resectoscope body through motion of the trigger, therefore additional instrumentation of the trigger position is required. This has been achieved by attaching a 10K Ω precision servo mount potentiometer which is coupled to the trigger by a gear mechanism, as shown in figure 3.9. The potentiometer is connected to a DT302 A/D⁶ board inside the

⁵The fulcrum point is governed by the prostatic sphincter in the real operation and the entrance to the phantom in CASTS, both at a point approximately one quarter of the resectoscope length.

⁶DT302 data acquisition board from Data Translation, Massachusetts, USA.




Figure 3.8: a) Geometrical arrangement of the IRED for the ring. The position and orientation reported by Flashpoint is a point 218mm along the normal which passes through the ring centre and is perpendicular to the plane of the ring, and corresponds to the end of the resectoscope outer sheath. b) The resectoscope IRED ring tracker.




Figure 3.9: Potentiometer attachment to detect the cutting loop motion.

CASTS PC.

3.2.3.3 Diathermy activation

The diathermy current to the cutting loop is activated by pressing the foot pedals. This action is sensed by external switches attached to the foot pedals and connected to the CASTS PC via additional inputs on the DT302 A/D board.

3.2.4 Computer system

The core of the CASTS system is a computer program which runs on a standard, low-cost PC⁷. The program was written in C++ using the Microsoft Visual Studio integrated development environment and the Microsoft foundation classes (MFC) framework for the user interface.

On execution the CASTS software performs several system checks and initializes communication with the A/D card and the Flashpoint tracking system through the serial port. The user then loads the VR model of the prostate phantom stored in an IGES file (section 3.2.4.1). Various settings may be adjusted before resection begins, such as the model view (wire frame or rendered and the viewing location), the positions of the verumontanum and bladder neck resection limits, and the number of cross sectional view planes through the model.

Once resection begins the CASTS software repeatedly polls the serial port for the Flashpoint readings of the resectoscope position and orientation, and polls the A/D card for the cutting loop position and the status of the foot pedals. This information is then used to update the user interface. Once the resection is complete various information is stored, such as the cutting data (*i.e.* the position of the cutting loop when the diathermy foot pedal is activated), the resection volume and the time taken to perform the procedure.

3.2.4.1 VR prostate phantom model

Hard casts of the prostate phantom components were supplied by Limbs and Things Ltd. The three casts have the shape of the whole phantom, the prostatic part of the phantom and the internal cavity, as shown in figure 3.10. A rough model of each of the casts was acquired by manually guiding a coordinate measuring machine (CMM)⁸ through a series of selected points. This data was then input to a program to command the coordinate measuring machine to automatically acquire detailed positional information on the various surfaces.

The raw xyz point data from the CMM was converted to non-rational uniform b-splines (NURBS) surfaces using the Rhinoceros software package⁹. The NURBS

⁷PIII 400MHz processor DELL Dimension XPS T500 PC, running Microsoft Windows 98 operating system.

⁸GC-90 CMM from LK Metrology Systems Inc, Brighton, Michigan, USA.

⁹Robert McNeel & Associates, Seattle, Washington, USA.




Figure 3.10: Prostate phantom casts.




Figure 3.11: Rendering of the VR model of the Limbs and Things prostate phantom.

model data is stored as an IGES file¹⁰, and a rendering of the complete model is shown in figure 3.11.

3.2.4.2 User interface

During TURP, the surgeon must interpret the endoscopic view while using both hands to operate the resectoscope and one foot to control the diathermy pedals. Therefore, the user interface for CASTS has been designed to keep the information display simple, and to a minimum. The user interface provides the following features:

- A 3D rendering of the VR prostate phantom model and the resectoscope. The user can decide between displaying the whole model or a partial view by cutting away the top part, which lies above the resectoscope.

¹⁰Initial graphics exchange specification (IGES) standard. National Institute of Standards and Technology, Gaithersburg, Maryland, USA.




Figure 3.12: User interface for the CASTS system (resectoscope not shown).

- A 2D view of the section closest to the cutting loop position, perpendicular to the base of the phantom, showing the phantom's outline, the cutting loop's position and the area resected.
- A collection of 2D views (similar to the one described above), shown as thumbnail images, at pre-specified positions equally spaced at intervals between the bladder neck and verumontanum.

The user can choose which views are shown at any one time, including the ability to turn off all displays so that resection is performed using only endoscopic guidance, but with CASTS monitoring progress. The display of the resectoscope changes colour to provide simple and immediate visual feedback about the state of the tool (green in normal mode; red if cutting; yellow if coagulating; grey if obscured).

The user interface is shown in figure 3.12.

3.3 System trials

Initial trials of the CASTS system have been performed by two surgeons, one being an urological trainee and the other a consultant urologist. Figure 3.13 shows the system in use during a resection training session. The system was found to be helpful and easy to use, and a beneficial addition to training. The computer display provides additional visual confirmation of the resectoscope cutting relative to the critical prostatic features. The use of a resectable physical phantom provides a realistic sense of “feel” on the handle of the resectoscope when pressing and cutting the prostate.



Figure 3.13: Initial trials of the CASTS resection system.

However, several problems were highlighted which cause mismatches between the resection progress data in the VR display and the true resection in the prostate phantom. Some shape variability was observed in the prostate phantoms, possibly as a result of shrinkage over time as the material loses moisture. The phantom is constructed of a spongy material which deforms and moves under the action of the resectoscope. The wire cutting loop has a tendency to bend out of shape during cutting so that its reported position to CASTS does not match reality. As resection progressed the phantom was found to change shape due to the release of internal forces inherent in the manufacturing process. The phantom material is hydrophilic and the phantom swells due to the absorption of the irrigant. The system also suffers a time lag between movements of the resectoscope and corresponding update of the user-interface.

Several improvements have been implemented in an attempt to reduce the movement of the phantom during resection. A set of Perspex profiles, which slot into the mock-up abdomen box and surround the phantom, were constructed to provide a mechanical constraint to movement. Two new VR NURBS models were defined, each with a pre-calculated degree of deformation (either to the left or to the right). The deformed VR models are displayed when the appropriate degree of lateral movement of the resectoscope is achieved. Problems of cutting loop distortion were overcome by using a loop constructed of thicker wire – this is a slight departure from the constraint of not wishing to change the procedure, but the slight change haptic sensation of the new cutting loop was not deemed to be problematic.

An example view of the user interface from the improved system is shown in figure 3.14. The deformation caused by the resectoscope is simulated in the VR model. The closest plane to the cutting loop is shown on the right hand side, with the cross hairs depicting the current cutting loop position, and the blue circles




Figure 3.14: CASTS user interface with deformed VR prostate model.

show the positions of previous cuts.

The CASTS system using the pre-calculated deformation VR models and Perspex profiles gave more realistic results than the initial system. However, the mechanical constraints inhibit the resectoscope's range of motion so that achieving a full resection of the phantom prostate is difficult. The pre-calculated deformations models are used once the resectoscope moves laterally by more than threshold, instead of continuous deformation with movement of the resectoscope, so the visual appearance is unrealistic since the model jumps suddenly between the relaxed and deformed states. In addition, the swelling and internal force release during the course of the procedure could not be compensated.

Figure 3.15a) shows the user interface following the completion of a resection, where the model of the resection cuts (in blue) extends beyond the capsule of the prostate. In the corresponding prostate phantom (figure 3.15b), the resection does not overlap the capsule.

3.4 Improving the VR accuracy

The CASTS system does not provide an accurate visualization of the resection process, due to the swelling, movement and deformation of the phantom, despite the addition of Perspex profiles and pre-calculated deformation. The solution to improving accuracy lies in closing the loop, *i.e.* in verifying what tissue has actually been resected.




Figure 3.15: a) User interface following a resection training session. The right hand side of the VR prostate model is cut through with resection strokes. b) The prostate phantom from the training session does not exhibit the capsule penetration.

3.4.1 Tracking the prostate phantom

Additional tracking must be incorporated in CASTS so that any changes to the phantom are correctly updated on the user interface. There are three methods worth considering as potential solutions to the problem of poor accuracy, namely finite element modelling, optical tracking and ultrasound based tracking.

3.4.1.1 Finite element modelling (FEM)

A finite element model of the prostate phantom could be incorporated into the software component of CASTS. FEM has been as the basis for modelling tissue in several MIS training systems [204, 228, 233]. However, FEM has a large overhead in terms of computational power; to provide real-time modelling many simplifications are required. This has been one of the drawbacks of FEM based training systems. To provide accurate results for CASTS a FEM would need to model the material properties of the phantom prostate, and how those properties vary as irrigant is absorbed and the model is cut. Such a complex model could not be interrogated in real-time to provide timely display information.

3.4.1.2 Optical tracking

Optical tracking is already used in CASTS for tracking the resectoscope position, so initially looks attractive as a potential solution for tracking the phantom movement. Tracking would require placing IRED at points on the surface of the latex bladder sleeve. The position information would be used to update the control points on the NURBS prostate model, and give an accurate real-time display.

In practice, it is not feasible to use Flashpoint to track the phantom movement. Securely attaching the IRED to the latex bladder would be problematic; the bladder must be inverted and stretched when it is fitted to the phantom. The latex bladder fills with water during the resection so preventing short circuits would be difficult. These two problems could be solved by replacing the active infrared Flashpoint system with a combined active/passive infrared system such as the Polaris¹¹. The Polaris system would track the resectoscope in an identical fashion to the Flashpoint, but also provide the capability to track the position of passive infrared reflecting markers which have no trailing wires and can be securely fixed to the latex.

Passive optical tracking is a feasible method for tracking the movement of the phantom, however it does not provide any information regarding the deformation. Therefore, optical tracking does not provide a complete solution to the CASTS deficiencies.

3.4.1.3 Ultrasound based tracking

The ability to follow the changes to the prostate phantom requires knowledge of both the shape and position of the prostate capsule and the shape and position of the internal resection cavity. Therefore, the tracking method must be able to “see” through the phantom and determine changes on the inside. Ultrasound imaging offers the only feasible solution to the problem of tracking changes to the prostate capsule, and verifying the resection volume.

Ultrasound tracking provides an additional benefit in that the differences between CASTS and CASMS are reduced. Both will use ultrasound as the basis of tracking the prostate, with the difference that CASTS uses a phantom prostate, while CASMS reports on the TURP procedure in the prostate of a patient. However the common ground of both systems is sufficient to allow CASTS to become a laboratory based proof of concept system for CASMS.

¹¹Northern Digital Inc, Ontario, Canada.




Figure 3.16: a) Phantom casts and moulds. b) Modified latex bladder with access for ultrasound.

3.4.2 Prostate ultrasound model

Initial tests with the Limbs and Things phantom prostate showed that it was impervious to ultrasound. Therefore a new phantom was required which was visible in the ultrasound image but can still be resected. The new phantom is based upon the Limbs and Things model but with an extra section in the base which has an opening to allow insertion of the ultrasound probe.

The moulds for the prostate phantom were constructed in two pieces using T28 silicone rubber¹². Three moulds are required representing the outer model including the ultrasound orifice, the prostate capsule and the internal urethral cavity, shown left to right in figure 3.16a). A hollow cast of the internal urethral cavity was made from latex by repeated pouring into the silicone mould (right hand mould in figure 3.16a), which is shown in figure 3.16b). Figure 3.16b) also shows a modified version of the latex bladder sleeve for the new phantom which allows access of the ultrasound probe. Several materials were tested as potential candidates for the phantom material and ultrasound scatterers, with the best results found using a gelatin¹³ matrix containing psyllium fibre¹⁴, as described in section 5.3.3. The gelatin/fibre solution contained 500gl^{-1} of raw gelatin and 150gl^{-1} of psyllium fibre, which gives good acoustic scattering and a firm gelatin matrix when the solution is cooled below 30°C .

The phantom is constructed in stages. The inner latex cast is placed in its rubber mould and the mould is closed with a thick, square metal plate which contains

¹²Alec Tiranti Ltd, London, UK.

¹³Supercook Ltd, Leeds, UK.

¹⁴Regucol dietary fibre, Holland and Barrett Retail Ltd, Burton Upon Trent, Staffordshire, UK.




Figure 3.17: Resectable prostate ultrasound phantom.

pour holes. The latex is filled with water and frozen, thus forming a solid copy of the internal urethral cavity. The metal plate and the spur at the end of the latex cast allow registration with the other moulds; the plate fits into the moulds at the points marked A in figure 3.16a), while the latex spur fits at the points marked B. The prostate capsule is formed by placing the latex/ice cavity mould into the centre mould of figure 3.16a), and filling the mould with gelatin/fibre solution coloured with food colouring agent. The filled mould was placed in a freezer to set the gelatin and keep the ice frozen in the latex mould.

Once set, the prostate capsule is placed in the largest mould, along with a metal rod which forms the access hole for the ultrasound probe. The metal rod can be seen inside the wooden part of the mould former in figure 3.16a). The mould is filled with gelatin solution which does not contain psyllium fibre or colouring agent, so that it does not scatter the ultrasound beam and is visually distinguishable from the prostate capsule. The finished prostate ultrasound phantom is shown in figure 3.17a) and a ultrasound image taken with the phantom in a water bath is shown in figure 3.17b).

One initial purpose perceived for the phantom was to generate clinically realistic ultrasound images under controlled circumstances, with the ability to vary the types of artefact (see section 2.6.3) present in the image and with accurate knowledge of the size and shape of the prostate. This is discussed in more detail in section 5.3.3. The phantom is used in the system trials presented in chapter 7.

3.5 Conclusion

This chapter discussed various aspects of the CASTS training system for TURP. The surgical tools and associated hardware were described including the adaptations required so that the equipment is instrumented for position and state tracking. The Limbs and Things prostate phantom was reviewed and the formation of a VR model of this phantom reported. Several features of the software system for CASTS were introduced.

Initial results using the CASTS system show discrepancies between the computer system and reality. Different attempted solutions to reduce errors provide some degree of improvement. However, these improvements were not considered sufficient for the system to be tested in a full experimental trial to test the efficacy of CASTS as a surgical training aid.

Various potential methods for tracking the prostate phantom are given, and ultrasound highlighted as the most suitable method. The construction of an ultrasoundable, resectable prostate phantom has been given. An example ultrasound image of the phantom is given. Use of ultrasound to track the prostate phantom closes the gap between CASTS and CASMS so that the only difference between the usage is the resection target – patient’s prostate or phantom prostate. Therefore, the CASTS system with ultrasound can be used as a laboratory based “proof-of-concept” test for the in-vivo TURP CASMS system.

Chapter 4

Prostate outline detection

4.1 Introduction

The goal of image processing is to provide an understanding of the image – to identify the underlying structures which are present and deduce information about the objects which constitute the image. Such a task is by no means easy. Some success has been achieved in identifying man-made objects in images which are characterized by their uniformity. In contrast, natural objects such as the organs within the human body, display considerable variation in shape, size and visual appearance, making automated recognition extremely difficult.

One approach to outlining objects of interest in images is to use human interaction to manually or semi-automatically draw the outline. For example, a mouse or light pen could be used to draw the complete outline of the boundary by hand, or to specify salient points on the boundary which are then joined by automated contour tracking algorithms or more simple geometric curves. However, the use of manual outlining is both tedious and time-consuming, and certainly unfeasible for the current goal of tracking the prostate movement during the course of an operation.

The work described in this chapter is an initial investigation into the possible methods of automatically determining the contour of the prostate in ultrasound images. The adopted approach is holistic in nature – *i.e.* many approaches covering a wide range of image processing techniques have been investigated for a wide range of applications. The techniques fall roughly into three separate categories. For each of these categories, the initial outline detection trials were performed

using standard¹ image processing techniques. The remaining techniques are those described in the literature which have been previously used for boundary detection in ultrasound images.

Ultrasound images of the heart [14, 86, 150, 155, 174, 214], blood vessels [11, 153, 197], bladder wall [213], gall bladder [128], ovaries [114, 188, 235] and uterus [187] have all been processed with the aim of segmentation or boundary extraction. A feature common to all of these applications is that the objects of interest are either surrounded by or contain fluid. Tissue-fluid interfaces are more easily discernible in ultrasound images due to the relatively large sound velocity variation across the boundary. Fluid-filled cavities show up as constant grey level regions (usually black) in the image. In contrast, ultrasound images of the prostate are prone to poor distinction between the prostate boundary and surrounds, and suffer from noise artefacts inside the boundary, in addition to the poor resolution associated with ultrasound as an imaging modality due to speckle and artefacts (as described in sections 2.6.2 and 2.6.3). For this reason the emphasis of the search has been placed on techniques applied to processing of prostate ultrasound images, or ultrasound images where the target has prostate-like ultrasound characteristics (*e.g.* kidney).

The methods from the literature which have been applied to ultrasound images for outline detection of the prostate can be divided into three categories. Algorithms in the first category, follow similar principles to the standard image processing edge detection techniques. Following smoothing of the image using a speckle reduction algorithm, local gradient or second derivative operators (or a combination of both) are applied to the image. The resulting outputs contain edge segments corresponding to the strongest edges in the input image. These algorithms are reviewed in section 4.3.1.

Techniques from the second category use a model to represent the boundary, often based on a priori knowledge of the likely organ shape. The model is changed until some criterion is met which matches the model to the boundary in the image. An advantage of the second category techniques is that a complete boundary is inherent. Section 4.4 reviews algorithms from the second category.

The third category contains algorithms that attempt to segment the image into regions. Although not strictly boundary detection techniques, they can be used to find the boundary if they reliably segment the image into prostate and non-

¹Standard in this context covers those techniques which may be found in introductory image processing texts, such as [88] or [168].

prostate regions. They are discussed in section 4.5.

4.1.1 Requirements of an outlining algorithm for CASMS

The purpose of finding the prostate capsule in the ultrasound images is to enable CASMS to construct a VR model specific to the patient's prostate. CASMS forms the VR model using the capsule outlines from a set of ultrasound images which cover the resectable region within the prostate. The model formation process is discussed in section 7.4.3, chapter 7. Creating a computer model of patient anatomy from pre-operative scans is a common goal in computer assisted surgery, and had been used before for TURP by the PROBOT [54, 55] prostate robot. However, CASMS requires further knowledge than previous systems, since it must update the VR model as the course of the operation proceeds.

The following list details the demanding requirements of an automated algorithm which is suitable for use in CASMS:

- Precision – Accurate knowledge of the prostate outline is required so that CASMS can provide correct positional information and warnings on proximity to the limits of resection. Careful resection strokes remove 1 mm of tissue, therefore the accuracy of the algorithm must be of the order of 1mm or less.
- Automated – CASMS could use surgeon knowledge during the initial construction of the VR prostate model in much the same way as the PROBOT (where the surgeon uses a light pen to mark the resection boundary on a series of TUUS scans). However it is not possible for the surgeon to provide information (such as salient outline points) to the algorithm once the operation is in progress. Therefore, the algorithm must perform without human input.
- Unsupervised – The automated algorithm must produce accurate results without the need for user confirmation that the result is indeed correct.
- Timeliness – CASMS must perform at interactive rates (10 frames per second or better), *i.e.* the VR model is updated without a severe time lag between the display and reality. Therefore, the algorithm must provide outlining results in a fraction of a second.

- General – The prostate varies in size and shape in different patients. Several factors, such as the presence of calcifications in the prostate, can result in artefacts in the ultrasound image. The algorithm must give valid results despite these variations in the ultrasound image.
- Poor image quality – Modern ultrasound machines provide better quality images than their predecessors, although the improvements are not so great that the problems (see sections 2.6.2 and 2.6.3) which make outlining difficult are removed. However, while the ultrasound machines available for this project are not state-of-the-art, they are comparable to the type of machine currently in use in many hospital urological departments. The algorithm must produce valid results for images from the available ultrasound machines, with better results obtained if a modern scanner's images are available.

The first requirement, automated processing, is a significant detour from many outlining algorithms previously reported in the literature. Where user interaction has been used in other's work, the amount required varies from the minimal (such as selecting the prostate's centre point) through to extensive interaction (*e.g.* selecting multiple points on an outline or hand modifying an algorithm's result). The benefits of involving the user include a reduction in the algorithm's processing time (since the search space is reduced to a small area around the specified points), and the ability for the user acceptance of the algorithm result. However, an outline can be quickly and completely specified by a small number of user-defined points joined with a spline curve (as used in section 5.3.4), or in a comparative time by using a light pen to draw the whole outline, without involving any image processing. Therefore, the benefit of an image processing algorithm which requires user initialization of multiple points to find the outline, is questionable. Algorithms which require user specification of a single point inside the prostate are better candidates for the CASMS task, provided some method can be found to automate the initialization process.

4.1.2 Algorithm implementation

Some of the algorithms described in the following sections have been implemented to assess their performance on transrectal ultrasound images. The choice of whether to pursue an algorithm is made predominantly on the presented results in the literature and the degree of user interaction required during the course of

the outline detection. As described in section 4.1.1, an ideal algorithm would produce valid results on many different prostate ultrasound images, would not require user input, and would produce a rapid result.

The algorithms have been implemented on the Matlab^{TM2} numerical computing platform. Matlab has been chosen since it allows fast implementation of the algorithms without the need to code much of the functionality, such as loading, saving and displaying images. The Matlab Image Processing toolbox also provides many functions such as standard filters, morphological operators and convolutions. The genetic algorithms optimization toolbox [102] has been used for implementation of the prostate detection algorithm of Wu [225] (see section 4.4.3). All other software implementation has been designed and written by the author.

One disadvantage of Matlab is that it is an interpreted script language. Interpreted languages compile the program code on execution rather than pre-compiling before execution which occurs in compiled languages such as C or C++. The result is that algorithms written in Matlab typically run more slowly than the equivalent C++ version, unless the algorithm is implemented using standard in-built Matlab functions which are themselves pre-compiled C++ routines. Where possible the in-built Matlab functions have been used to increase algorithm execution speed.

Fast algorithm execution is of paramount importance when the algorithm is used in a real-time system such as CASMS. However, a fast algorithm which fails to find the prostate boundary is of no use at all! During the initial assessment of the existing algorithms for prostate outlining, algorithm execution speed is not considered a factor in determining the success of any particular algorithm.

The implemented algorithms have all been tested using the prostate ultrasound image given in figure 4.6.

4.2 Edge based detection schemes

4.2.1 Edge definition

Edges are areas within an image which exhibit strong intensity contrast. Since edges often occur at locations within the image which correspond with object boundaries, edge detection is extensively used in image processing when the goal

²Matlab, The Mathworks, Inc., Massachusetts, USA.




Figure 4.1: First and second derivatives for a one-dimensional edge.

is to divide the image into its constituent objects. The standard approach to edge detection [82] is to calculate the derivatives of the image, since the gradient of the image changes rapidly across an ideal edge. Figure 4.1 shows a one-dimensional edge profile, and the first and second derivatives of the profile. The position of the edge may be estimated from the maximum value of the first derivative or the zero-crossing of the second derivative. Therefore, a method of approximating the first or second derivative of an image is required. For a discrete one-dimensional profile $f(i)$, the first derivative may be approximated by:

$$\frac{df(i)}{d(i)} = f(i+1) - f(i) \quad (4.1)$$

and the second derivative by:

$$\frac{d^2 f(i)}{d(i)^2} = f(i+1) - 2f(i) + f(i-1) \quad (4.2)$$

These operations may be implemented as convolution operations, by convolving $f(i)$ with $\begin{bmatrix} -1 & 1 \end{bmatrix}$ and $\begin{bmatrix} 1 & -2 & 1 \end{bmatrix}$ for the first and second derivative respectively. The sign of the output is often disregarded, since information regarding the direction of the edge is usually not required.

4.2.2 First derivative approximations

4.2.2.1 Roberts cross operator

In two dimensions, there are different approximations to the first derivative operator. The simplest is the Roberts cross operator [183] which approximates the magnitude of the gradient by :

$$G(f(i, j)) = |f(i, j) - f(i + 1, j + 1)| + |f(i + 1, j) - f(i, j + 1)| \quad (4.3)$$

This can be implemented as a convolution operation:

$$G(f(i, j)) = |G_x| + |G_y| \quad (4.4)$$

where G_x and G_y are the result of convolving $f(i, j)$ with the convolution masks:

$$R_x = \begin{bmatrix} 1 & 0 \\ 0 & -1 \end{bmatrix} \quad R_y = \begin{bmatrix} 0 & -1 \\ 1 & 0 \end{bmatrix} \quad (4.5)$$

The masks have an even number of values in each dimension (*i.e.* 2 by 2) with the result that the differences are computed at the point $(i + \frac{1}{2}, j + \frac{1}{2})$. Therefore, the Roberts operator is an approximation of the gradient at this interpolated point rather than the point (i, j) .

4.2.2.2 Sobel operator

The Sobel operator [24] avoids the problem of interpolating the point between the pixels by using a 3 by 3 neighbourhood centred on the point (i, j) . The operation gives the magnitude of the gradient calculated as:

$$G(f(i, j)) = \left| \sqrt{G_x^2 + G_y^2} \right| \quad (4.6)$$

where the partial derivatives G_x and G_y are calculated as:

$$G_x = (f(i - 1, j + 1) + 2f(i, j + 1) + f(i + 1, j + 1)) - (f(i - 1, j - 1) + 2f(i, j - 1) + f(i + 1, j - 1)) \quad (4.7)$$

$$G_y = (f(i - 1, j - 1) + 2f(i - 1, j) + f(i - 1, j + 1)) - (f(i + 1, j - 1) + 2f(i + 1, j) + f(i + 1, j + 1)) \quad (4.8)$$

and implemented by convolution with the masks:

$$S_x = \begin{bmatrix} -1 & 0 & 1 \\ -2 & 0 & 2 \\ -1 & 0 & 1 \end{bmatrix} \quad S_y = \begin{bmatrix} 1 & 2 & 1 \\ 0 & 0 & 0 \\ -1 & -2 & -1 \end{bmatrix} \quad (4.9)$$

4.2.2.3 Prewitt operator

The Prewitt operator [53] is similar to the Sobel operator except that it does not place any emphasis on the pixels which lie closer to the centre of the masks. G_x and G_y are calculated by convolution with the Prewitt masks:

$$P_x = \begin{bmatrix} -1 & 0 & 1 \\ -1 & 0 & 1 \\ -1 & 0 & 1 \end{bmatrix} \quad P_y = \begin{bmatrix} 1 & 1 & 1 \\ 0 & 0 & 0 \\ -1 & -1 & -1 \end{bmatrix} \quad (4.10)$$

4.2.2.4 Edge maps

The first derivative operators are usually thresholded to form a binary image which shows the positions of likely edges within the image. Such an image is referred to as an *edge map*. Setting the threshold is problematic; too low and many spurious edges are included in the edge map, too high and some correct edge information is lost. One approach used to automate the threshold setting is to calculate a root mean squared estimate of the noise in the image and set the threshold to this value.

4.2.3 Second derivative approximation

The one dimensional first derivative shown in figure 4.1 demonstrates one of the problems of using the first derivative and a threshold to find the edge map for an image. The red dashed line indicates a threshold level for the binary edge map – any points above the line will become edge points in the edge map. Therefore, the edge becomes wide and the true edge position is lost.

The second derivative provides a method for estimating the correct location of the edge by finding the zero crossings – signified by the red circle in the second derivative in figure 4.1. The second derivative may be approximated by convolving the image with a symmetrical second derivative mask. The mask varies according

to a parameter α which controls the shape:

$$\begin{bmatrix} \frac{\alpha}{\alpha+1} & \frac{1-\alpha}{1+\alpha} & \frac{\alpha}{\alpha+1} \\ \frac{1-\alpha}{1+\alpha} & \frac{-4}{1+\alpha} & \frac{1-\alpha}{1+\alpha} \\ \frac{\alpha}{\alpha+1} & \frac{1-\alpha}{1+\alpha} & \frac{\alpha}{\alpha+1} \end{bmatrix} \quad 0 \leq \alpha \leq 1 \quad (4.11)$$

Some commonly used second derivative masks are:

$$\begin{bmatrix} 0 & 1 & 0 \\ 1 & -4 & 1 \\ 0 & 1 & 0 \end{bmatrix} \alpha = 0 \quad \begin{bmatrix} 1 & 4 & 1 \\ 4 & -20 & 4 \\ 1 & 4 & 1 \end{bmatrix} \alpha = 0.2 \quad \begin{bmatrix} 1 & 0 & 1 \\ 0 & -4 & 0 \\ 1 & 0 & 1 \end{bmatrix} \alpha = 1 \quad (4.12)$$

The convolution result contains values which are both positive and negative. This is converted to a binary image by setting to one those values which are greater than or equal to zero, and setting values less than zero to zero. The true edge positions occur at the boundaries of the binary image and may be found by applying one of the first derivative edge operators.

4.2.4 Results of the first and second derivative approximations

Figure 4.2 shows a test image of red blood cells³ which exhibits strong edge components around the individual cells. Figure 4.3 shows the first derivative approximations using the Roberts, Sobel and Prewitt operators. Figure 4.4 shows the corresponding edge maps using an automatically determined threshold based on the root mean estimate of the noise. The Roberts operator is clearly more affected by noise and produces edges which are incomplete. The Sobel and Prewitt operators perform much better, producing edge maps which are virtually free of noise whilst retaining the complete outline of the blood cells.

Figure 4.5 shows the result of applying the second derivative approximation to the blood image, and the resulting binary image. The second derivative approximation is displayed with an adjusted grey scale showing pixels with zero value as grey, positive pixels are lighter and negative pixels darker. These figures show the high sensitivity to noise associated with the second derivative. Although the

³The blood cells image is an example image which is included in the Matlab package.




Figure 4.2: Microscopic image of red blood cells (an example image supplied with the Maltlab programming platform).




Figure 4.3: First derivative approximations a) Roberts, b) Sobel, c) Prewitt operators.




Figure 4.4: Edge maps a) Roberts, b) Sobel, c) Prewitt operators.




Figure 4.5: a) Second derivative approximation and b) binary result.

blood cell edges are visible as the intersections between the black rings and white rings in the image, any attempt to find these intersections using a first derivative edge operator would be swamped by noise.

The binary image from the second derivative highlight shows the position of all the zero crossings within the image regardless of how strong those edges are. This leads to an interesting point - what are the correct edges of the blood cells? The first derivative operators are clearly finding the edges which correspond to the largest gradient change – the intersection of the white and black rings for each blood cell in figure 4.2. In fact, the actual outline is probably the outer perimeter of the white ring. Whether this is important or not depends on the intended use of the edge detection, for example if the intention is to count the number of red blood cells the exact outline is not necessary; if the processing aims to calculate the area of the red blood cell, the correct outline is required.

4.2.4.1 First and second derivative results on ultrasound images

Figure 4.6 shows a transrectal ultrasound image of a prostate from a patient with symptoms of benign prostatic hyperplasia. The image is taken at the mid-section of the prostate gland – the section which gives the largest prostate area in the ultrasound image. This image has been selected for its high quality, the prostate is clearly visible and the image does not contain significant artefacts which are often apparent in ultrasound images (see chapter 2, section 2.6.3). The prostate outline is indicated by the blue arrows present within the image. A discussion of the correct location of the prostate outline is presented in chapter 5, section 5.3.

Figure 4.7 shows the author's result of applying the Sobel operator to the prostate ultrasound image, and the corresponding edge map. It is clear that the Sobel operator finds the strongest edges in the image - the edges of the black circle which corresponds to the position of the ultrasound probe. Figure 4.8 shows the ultrasound image with this black circle filled.⁴ This filled image is the test image for the outlining algorithms described in the rest of this chapter.

Figure 4.9 gives the result of applying the Sobel edge operator to the ultrasound image of figure 4.8. The result demonstrates the difficulty of finding the edge of the prostate – no useful information regarding the position of the prostate edge is given in the images. It is clear that standard image processing techniques are

⁴Filling the black circle is an automated process. The black circle is easily extracted since it is the only portion of the ultrasound image where the pixel values are zero. The region is filled using values interpolated from the pixels adjacent to the probe edges.




Figure 4.6: Ultrasound image of the prostate mid-section of a patient with benign prostatic hyperplasia.




Figure 4.7: a) First derivative approximation using the Sobel operator. b) Edge map result.




Figure 4.8: Prostate ultrasound image with interpolated filled probe region.




Figure 4.9: a) First derivative approximation using the Sobel operator on the filled probe image. b) Edge map result.

not applicable when processing ultrasound images. Ultrasound images are low in contrast, dominated by a grainy texture known as *speckle*. Speckle occurs when the reflected ultrasound waves interfere at the transducer causing construction and destruction of the received echo signal. These reflections result from the presence of small scatterers (small on the scale of the ultrasound wavelength) within the tissue. Larger reflections from tissue region boundaries result in the perceived edges visible in the ultrasound image. Unfortunately, it is not possible to discriminate between the reflections from tissue region boundaries and the reflections from the small scatterers. If speckle is considered as noise within the image⁵, then it is clear that the signal-to-noise ratio of an ultrasound image is very low compared to most other types of image. Speckle causes the many spurious edge segments which are visible in the edge maps of figure 4.9.

4.2.5 Noise smoothing filters

The results in the previous section demonstrate the problems caused by noise present in images, compounded in ultrasound images by the low signal-to-noise ratio which renders the first derivative operators useless for finding edges. The second derivative operator used alone is of no use even on a reasonably clean image such as the blood cell example. This may be anticipated since the second derivative operator does not make any use of the magnitude of an intensity change, the smallest change in grey-level value between two pixels will result in a zero-crossing at that point.

Noise smoothing is a standard preprocessing stage in many image processing applications. The goal of any noise smoothing filter is to smooth areas of low grey level variation without blurring edges - regions containing large intensity changes. Standard smoothing filters include median filtering and spatial averaging. Median filtering removes impulse noise within an image by replacing each pixel with the median pixel value in a neighbourhood centred on the pixel. Spatial averaging replaces each pixel with a mean or weighted mean value of pixels in a neighbourhood around the pixel.

Figure 4.10 shows the blood cell image after applying 9 by 9 pixel median and spatial averaging filters, and the corresponding edge maps after applying the Sobel operator and thresholding the result. In comparison with figure 4.4, it is

⁵Although commonly referred to as *noise*, speckle is inherent in the image formation process – with no speckle there would be no image.

clear that the edge maps after noise smoothing contain less false edges and the edge of the blood cells are also more complete.

Figures 4.11a) and b) shows the author's result of applying the median filter and spatial averaging filter to the prostate ultrasound image of figure 4.8. Figures 4.11c) and d) show the corresponding edge maps resulting from Sobel edge detection. Although there are still many false edge segments present in both edge maps, some parts of the prostate outline begin to be apparent, particularly in the edge map formed after spatial averaging.

A further standard noise smoothing technique is to use a Gaussian kernel to blur the image before applying an edge operator. This is very commonly used with the Laplacian second-derivative operator, and has been used by several authors for finding the prostate in ultrasound images. Details of the Laplacian of Gaussian operator are given in section 4.3.2.


Noise smoothing in ultrasound images (and other images containing speckle) has been a major topic of research and many specialized filters have been developed. Such filters are termed *speckle-reduction* filters and a review of many of them is presented in appendix A. Section 6.6 in chapter 6 describes work undertaken to investigate the statistics of ultrasound images using these speckle-reduction filters.

4.3 Outline detection in ultrasound images

Standard edge detection techniques fail to produce good results directly due to the noisy nature of ultrasound images. Therefore, the remainder of the work described in this chapter reviews research carried out in the field of ultrasound image processing, with an emphasis on techniques developed for medical ultrasound image processing. A common theme of medical ultrasound image processing is the detection of anatomical structures within an image.


4.3.1 Edge based detection schemes

The algorithms described in the following sections attempt to find the prostate outline by finding likely edges within the image. The basic concept for each algorithm is essentially the same. Following smoothing and edge detection, likely edge segments are identified and joined to form a final closed curve. Sections 4.3.2



(a)

(b)



(c)

(d)

Figure 4.10: a) Median filtered blood cell image. b) Spatial averaging applied to the blood cell image. c) Sobel edge result following median filtering. d) Sobel edge result following spatial averaging.




Figure 4.11: a) Median filtered prostate ultrasound image. b) Spatial averaging applied to the ultrasound image. c) Sobel edge result following median filtering. d) Sobel edge result following spatial averaging.

through 4.3.5 describe various methods of finding edge segments in the ultrasound image, and section 4.3.6 details the method used to select prostate outline edge segments and form these segments into the final outline.

4.3.2 Laplacian of Gaussian

Several authors have applied the Laplacian of Gaussian (LOG) operator to prostate ultrasound images. Gomes et al [80] use the LOG operator to detect edges in transurethral ultrasound images of the prostate. The Gaussian filter serves to suppress the high frequency speckle inherent in the image, and the Laplacian operator⁶ highlights the positions of the edges. The Gaussian filter is described mathematically as:

$$G(x, y) = e^{-\frac{x^2}{\sigma_x^2} - \frac{y^2}{\sigma_y^2}} \quad (4.13)$$

and the LOG filter as:

$$LOG(x, y) = \frac{\partial^2 G(x, y)}{\partial x^2} + \frac{\partial^2 G(x, y)}{\partial y^2} \quad (4.14)$$

The degree of smoothing in the LOG filtered image depends on the width of the Gaussian filter, i.e. the values of σ_x^2 and σ_y^2 . Gomes et al report best results using $\sigma_x^2 = \sigma_y^2 = 64$ pixels.

The output of the LOG filter is an image containing both positive and negative pixel values. The zero crossings of this image are the locations of edges within the input ultrasound image. These zero crossings are found by forming a binary image of the output (by setting positive pixels to one and negative pixels to zero), and using a Prewitt filter to detect the binary image edges. Figure 4.12 shows the results of applying the LOG filter and zero-crossing detection to the prostate ultrasound test image of figure 4.8.

Richard et al [177] use a similar LOG operator to detect the prostate outline on transrectal ultrasound images.

One problem that is immediately apparent from the LOG result is that any information regarding the strength of the edges is discarded in the binary operation. Therefore, it is not possible to determine whether an edge segment in the final image is due to a large or small grey-level variation in the input ultrasound image. This problem has been broached by Zhao et al [234], who use the gradient

⁶The Laplacian operator approximates the second derivative of the image.




Figure 4.12: Zero-crossings of the LOG filtered prostate ultrasound image.

strength in the image as well as the locations of the zero-crossing of the second derivative, to detect prostate outlines on transrectal ultrasound images. The gradient strength is calculated by replacing each pixel with the maximum of the absolute gradient values in the four directions of the 3 by 3 neighbourhood of the pixel. Figure 4.13 shows the four gradient directions. The gradient strength image is masked by the binary edge image result from the LOG operator, to yield a final edge-intensity image which contains the positions of edges in the input ultrasound image and their relative strengths. Figure 4.14 shows the edge-intensity result from applying the Zhao algorithm to the ultrasound test image




Figure 4.13: The four gradient directions in the pixel neighbourhood.




Figure 4.14: The edge intensity result from the Zhao algorithm implementation, applied to the prostate ultrasound test image.

of figure 4.8, where darker lines represent higher intensity.

4.3.3 Edge strength and second derivative in gradient direction

Aarnink et al [1, 2] use a similar approach to produce an edge intensity image. The initial image is smoothed using a 21 by 21 pixel uniform blur filter instead of the LOG filter discussed previously. The filtered image is true sampled by factor 3 (i.e. keep every third pixel along each row and column) to reduce the total number of pixels nine-fold. The Laplacian of the compressed image is approximated by using the second derivative in gradient direction (SDGD) method of Haralick [87]. First an image is formed by replacing the value of a pixel with the maximum value over a 17 by 17 neighbourhood around the candidate pixel. A second image is created from the minimum of the same neighbourhood. Subtracting the compressed image from the maximum image gives an approximation to the gradient in the direction of the maximum, whilst subtracting the minimum image from the compressed image forms the gradient in the direction of the minimum. The SDGD results from the subtraction of minimum gradient direction from

maximum gradient direction:

$$SDGD(x, y) = \left(\max_n(x, y) - g(x, y) \right) - \left(g(x, y) - \min_n(x, y) \right) \quad (4.15)$$

where $\max_n$ is the maximum value in the n by n neighbourhood of the pixel of interest, $\min_n$ is the minimum value of the neighbourhood and g is the blurred, compressed image.

The zero crossings of the SDGD are found by looking at the eight neighbours of the candidate, and flagging an edge if the pixel values change sign along any of the four directions of the 3 by 3 neighbourhood. Thus, an edge is detected if:

$$\begin{aligned} & [(P_i = 0) \cap (P_{i-1}P_{i+1} < 0)] \cup \\ & [(P_iP_{i-1} < 0) \cap (\text{abs}(P_i) < \text{abs}(P_{i-1}))] \cup \\ & [(P_iP_{i+1} < 0) \cap (\text{abs}(P_i) \leq \text{abs}(P_{i-1}))] \end{aligned} \quad (4.16)$$

is true for any of the four directions (upper left to lower right, top to bottom, upper right to lower left and left to right), where P_{i-1} , P_i and P_{i+1} are the three pixels in the current direction.

The edge strength (ES) image is formed by taking the minimum value of the maximum and minimum gradient direction images [125]:

$$ES(x, y) = \min \left(\max_n(x, y) - g(x, y), g(x, y) - \min_n(x, y) \right) \quad (4.17)$$

The final edge intensity image is obtained from a pixel by pixel multiplication of the SDGD and ES images. The result of applying the Aarnink algorithm to the test ultrasound image is shown in figure 4.15.

The Aarnink algorithm provides a further processing step to detect likely prostate edges from the edge intensity image. The method is based on the hysteresis thresholding technique of Rosenfeld et al [184]. The edge intensity image is thresholded using a high threshold value to produce an initial binary edge map. An iterative process then adds additional pixels to the edge map if those pixels are adjacent to an existing pixel in the edge map, and above a second, lower threshold in the edge intensity image. Figure 4.16 shows the binary edge map overlaid on the input ultrasound image.

Further work by Aarnink et al [3] aims at improving the edge intensity image formation scheme. Following blurring of the input prostate image, an additional stage is added which calculates a binary map of the standard deviation of pixels




Figure 4.15: The edge intensity image resulting from applying the Aarnink algorithm implementation to the prostate ultrasound test image.

in the image. An image is created by assigning the standard deviation of 41 by 41 pixel neighbourhoods to the centre pixel of the neighbourhood. This is converted to a binary map by thresholding at a threshold determined by the global standard deviation in the blurred image. The map is used during the maximum and minimum neighbourhood operations to choose a neighbourhood size of 25 by 25 pixels or 49 by 49 pixels, where the local standard deviation is greater than or less than the global standard deviation, respectively.

4.3.4 Sticks and anisotropic diffusion filtering

Pathak et al [162] approach the problem of finding edge segments in the image in a different manner. The input image is first filtered using a maximum stick filter⁷ [49, 50]. The stick filter acts to accentuate lines within the image. The result is passed through a smoothing filter which acts anisotropically to flatten regions of relative smoothness, whilst maintaining features of significant grey-level variation. The filter used is analogous to the anisotropic diffusion filter for speckle

⁷A full description of the stick filtering algorithms is given in appendix A, sections A.1.8 and A.1.9.




Figure 4.16: Test image with overlaid edge map result from the Aarnink algorithm. The values for the hysteresis threshold were experimentally determined as 0.5 and 0.1 for the high and low thresholds.

reduction described by Perona and Malik [163] (see appendix A, section A.1.1).

The filtered image is modelled as a smooth function which acts as a weak elastic membrane. The function is determined by minimizing the energy terms of the membrane formulated from the degree of mismatch between the model and true image, penalties for discontinuities and for roughness in the absence of discontinuities. The degree of mismatch is determined as:

$$D = \sum_{xy} (S(x, y) - I(x, y))^2 \quad (4.18)$$

where S is the elastic membrane function and I is the smoothed stick filtered image. The discontinuity penalty is:

$$C = \alpha \sum_{xy} l(x, y) \quad (4.19)$$

where $l(x, y)$ is unity if a discontinuity is present at the position (x, y) or zero otherwise. The roughness factor is given by:

$$R = \beta \sum_{xy} (\nabla S(x, y))^2 (1 - l(x, y)) \quad (4.20)$$

Finding the most suitable membrane function can be viewed as a global optimization problem. The optimization aims to minimise the energy of the membrane such that:

$$\min_{\{S, I\}} E = \min_{\{S\}} \left(D + \min_{\{I\}} (R + C) \right)$$

The membrane function exhibits many local minima. Pathak uses the graduated non-convexity minimization algorithm of Blake and Zisserman [18] to calculate the global minimum.

Pathak's approach has been implemented using the analogous anisotropic diffusion filter rather than the elastic membrane function. Figure 4.17 shows the result of applying the sticks and anisotropic diffusion filters to the input test image.

The optimised membrane function is used as an input to a Canny edge detector [32]. The Canny detector applies Gaussian smoothing to the input and calculates the gradient magnitude and direction of the smoothed image. The next step finds points whose strength is locally maximum in the direction of the gradient. All non-maxima pixels are set to zero – hence this process is called non-maxima suppression. The final step is equivalent to the hysteresis threshold function –




Figure 4.17: The edge intensity result from applying the sticks and anisotropic diffusion filters to the input ultrasound test image.

finding the strong edges and iteratively connecting adjacent weaker edges. The edge map resulting from the Canny edge detector is shown in figure 4.18, overlaid on the original input ultrasound image.

4.3.5 Radial bas-relief

Liu et al [136] present an algorithm for initial processing of the ultrasound image to highlight the prostate edges, called radial bas-relief (RBR). The technique is derived from a darkroom technique used to manipulate photographs to produce images which have a 3D effect – ridges appear in the image lit from one side. In RBR, the input ultrasound image is normalized, inverted and enlarged by a constant factor radially from the image centre. The enlarged, inverted image is added to the normalized original image registered on the image centre:

$$A(x, y) = I(x, y) + 255 - I(kx, ky) \quad (4.21)$$

where $I(x, y)$ is the input image normalized to the range $[0 \cdots 255]$, the xy origin occurs at the image centre, and k^{-1} is the enlargement factor. Bi-linear interpo-




Figure 4.18: Edge map from the Pathak algorithm when applied to the initial ultrasound image input.

lation is used to calculate pixel values in the enlarged image. The pixel values in $A(x, y)$ are set to zero if they overflow the pixel range, then thresholded to produce a binary result:

$$B(x, y) = \begin{cases} 0 & \text{if } A(x, y) > 255 \\ A(x, y) & \text{otherwise} \end{cases} \quad (4.22)$$

$$C(x, y) = \begin{cases} 255 & \text{if } B(x, y) > t \\ 0 & \text{otherwise} \end{cases} \quad (4.23)$$

where $C(x, y)$ is the binary result and t is the binary threshold.

Figure 4.19 shows the result of applying the RBR algorithm to the test ultrasound image. The enlargement factor is 109%, which is the value quoted in Liu's paper.

The value of the threshold t used in Liu's work is set at 2. Further work on the RBR by Wu [224] demonstrates that the use of a threshold is unnecessary. The RBR method works by testing whether a candidate point is brighter than a point n pixels closer to the centre along the same radius. If the test is true, then the candidate pixel is marked as an edge pixel. The value of the offset n depends on




Figure 4.19: Binary result from the RBR algorithm applied to the test prostate ultrasound image.

the distance of the candidate pixel from the image centre, given by:

$$n = (1 - k) \sqrt{x^2 + y^2} \quad (4.24)$$

Wu suggests an improved version of the RBR method with a fixed value for n and using the grey-level difference between the two pixels rather than a binary test. This may be considered as a first-derivative type operation applied along the radii from the image centre. The modified RBR is implemented as a Cartesian to polar co-ordinate transformation, followed by convolution with a linear first-derivative mask. Figure 4.20 demonstrates the difference between the original RBR and modified RBR techniques. Referring to figure 4.20, the modified RBR result is calculated as:

$$D(x, y) = \begin{cases} I(x_o, y_o) - I(x_i, y_i) & \text{if } I(x_o, y_o) > I(x_i, y_i) \\ 0 & \text{otherwise} \end{cases} \quad (4.25)$$

Figure 4.21 shows the result of applying the modified RBR to the ultrasound image of figure 4.6 using a pixel offset n of 5 pixels.

Further processing is necessary to produce an edge map from the RBR result.




Figure 4.20: a) Original RBR. b) Modified RBR with fixed radial offset.




Figure 4.21: Modified RBR applied to the input ultrasound image.



Figure 4.22: Edge intensity result from the modified RBR technique.

In original RBR technique, morphological closing is performed on the binary result to smooth the pixel regions, followed by area thresholding to identify regions which correspond to the edge of the prostate. The area thresholding simply discards regions if the number of pixels within the region is less than a certain threshold. Further processing (not described in Liu's paper) is necessary to produce single pixel width edge segments from the region threshold result.

The approach adopted here is to use the zero crossings result of the LOG and Prewitt filtering algorithm described in section 4.3.2 as a mask for the modified RBR technique. The edge intensity image formed from masking the result shown in figure 4.21 is shown in figure 4.22. It should be noted that the use of the image centre for the RBR technique assumes that the image centre falls within the prostate. In the test ultrasound image, this criterion happens to be satisfied, however it is unlikely to hold true for every prostate ultrasound image. The best position for the RBR centre would be the centre of the prostate.

4.3.6 Edge selection

The goal of the image processing is to find a complete outline of the prostate in the ultrasound image. Some of the algorithms discussed so far have produced edge maps. Edge maps have been produced for the remaining algorithms by

applying the hysteresis thresholding algorithm described in section 4.3.3 to these algorithms' edge intensity results. The edge maps contain edge segments which are likely to correspond with the prostate outline, but also exhibit spurious edge segments due to other structures within the image. The next step in the image processing pipeline is to identify the true outline edge segments.

In much of the literature, selection of the appropriate edge segments is glossed over as a trivial step generally described as "further processing". In reality, robust identification of the correct prostate edges is a complex task. Rules which give a correct result for one image may often produce incorrect answers when applied on a complete data set. As an example, Pathak et al [161, 162] use the following procedure to find the likely edge segments⁸. The user is first asked to select a seed point within the prostate. Starting from the seed point a radial search is conducted to find the intersection between each radius and the closest edge segment. For each quadrant in the image (with the seed point defining the quadrant intersection) the mean and standard deviation of the distance from the radial search points to the seed point is calculated. An outlier filter then removes points which lie more than two standard deviations from the mean. A spline fitted to the remaining points forms the final closed boundary. This regime fails under two circumstances; either when a large line segment is present inside the prostate boundary and the radial search erroneously detects these points or when a significant portion of the prostate outline is missing in one or more quadrants.

The following paragraphs describe a series of steps which have been implemented by the author to attempt to distinguish the correct prostate edges segments from the false edge segments and produce a smooth continuous curve describing the prostate outline. Experimental results show that these steps produce the best results when applied across the data set and using the different algorithms' edge maps. Prior knowledge of the likely prostate characteristics are used as the basis for these heuristic rules.

One feature common to most prostate ultrasound images is that the average grey level within the prostate tends to be darker than the surrounding areas, and the prostate boundary is often ringed by a bright halo. Therefore it is highly likely that a true outline edge segment will occur at a dark to bright grey level transition when viewed from inside the prostate. Given a known point within the prostate, it is possible to remove edge segments which correspond to a bright to dark change moving radially out from the known point. Pathak [161, 162] and Aarnink [1, 3]

⁸Pathak's paper is the only one which gives a full description of the process of forming a complete outline from the edge map.

advocate user selection of a point within the prostate. Although this does not represent a time-consuming step, user-interaction is obviously required, departing from the automated criterion necessary for a monitoring system. However, for the initial testing of the algorithms described in this chapter, the point within the prostate has been set for the test ultrasound image of figure 4.8, and supplied to the algorithms as a parameter. Automated assignment of this point is discussed in chapter 5, section 5.6.

The next step removes edges whose length are less than a threshold. Although this may remove some segments which lie on the prostate boundary, many of the small segments result from small bright patches within the image, which may be considered as noise. Segments are also removed which lie within a small radius of the known point within the prostate.

A radial search from the known point finds points at the intersection of the remaining edge segments and the radii. An initial cubic spline is fitted to these points with certain constraints on the maximum change of radius and the smoothness of the spline. An iterative process reintroduces previously removed edge segments which lie close to the spline and the spline is recalculated from the new points.

The final step ensures that the spline doesn't cross the black circle of the probe. This can be a problem if the edge map does not contain an edge segment at the base of the prostate, around the probe. Given the difficulty of detecting the edges of the prostate at the lobes (where the prostate outline lies close to parallel to the ultrasound beam), it is possible that the spline can balloon out at the base and intersect the black circle. If this occurs, a separate local search is made at the probe periphery to find the most likely edge of the prostate.

Figure 4.23 shows the prostate ultrasound image from figure 4.6 overlaid with four outlines – the results of the LOG (red), Aarnink algorithm (green), Pathak algorithm (blue) and modified RBR technique (yellow).

4.4 Model based detection schemes

Algorithms described in the previous section all attempt to find the prostate outline by first finding likely edge segments, selectively choosing some of these edge segments and finally joining and interpolating between these segments to form a final closed outline. This may be considered a top-down approach, starting




Figure 4.23: Results of the edge based detection schemes. a) Laplacian of Gaussian, b) Aarnink, c) Pathak and d) modified RBR.

from the raw image data, and proceeding through a number of steps to arrive at a final result. The algorithms described in this section adopt a different approach. Each algorithm uses a representation of the prostate shape which varies according to the underlying grey level variation of the image to which the model is applied.

Model based detection schemes have the advantage over edge-based schemes that a complete boundary is intrinsic as part of the detection process. Therefore a model-based detection scheme will always produce a result in an image regardless of whether a valid target exists or not. However, each scheme uses some method of scoring to determine the appropriateness of a particular outline, and opts for the outline with the greatest score. Therefore, it is possible to use this outline score as a final measure of the likelihood that a given result is actually valid.

4.4.1 General model based image processing

Model based detection is a somewhat specialized topic in the field of image processing (as opposed to edge detection). One well-documented general technique is the Hough transform. Although generally described in terms of its use for detecting lines in an image at differing orientations and positions, or circles of different size and location, the Hough transform is applicable to any generalized shape.

The Hough transform uses an analytical description of a shape and assesses the likelihood of that shape appearing at each possible location and orientation within the image. For example, the Hough transform could be used to detect the black circle at the edge of the probe.⁹ This would involve constructing a three-dimensional array of bins; two dimensions for the horizontal and vertical location of the circle centre, and one for the radius. Then a given bin gives a circle of given size and position in the image, and the bin is assigned a score depending on how likely it is that the circle exists in the image (for example by averaging the gradient across the circle boundary). The bin which contains the greatest score determines the correct location and radius of the black circle.

The drawback of the Hough transform is that it only detects a single shape within the image even though that shape need not be described by a simple analytical equation. Therefore it is unsuitable for detecting the prostate outline

⁹The Hough transform was originally used by the author to find the black circle within the image. However the technique was abandoned in favour of a faster, direct method which relies on the grey scale value of the black circle pixels always being zero.

in ultrasound images due to the variation in prostate shape which may be observed in prostate images from different patients.

4.4.2 Snake based techniques

Snakes, or *active contours*, are boundary curves which change shape according to the forces applied to them. The forces include internal forces which are intrinsic in the snake itself, and external forces which are produced by the image data to which the snake is applied. These internal and external forces are defined so that the snake will migrate towards the object boundary within an image, with equilibrium achieved once the snake reaches the boundary. The traditional snake [107] is defined as a parametric curve $\mathbf{x}$, that moves within the image so as to minimise the energy function E of the snake:

$$\mathbf{x}(s) = [x(s), y(s)], s \in [0, 1] \quad (4.26)$$

$$E = \int_0^1 \frac{1}{2} (\alpha |\mathbf{x}'(s)|^2 + \beta |\mathbf{x}''(s)|^2) + E_{ext}(\mathbf{x}(s)) ds \quad (4.27)$$

where α and β are weighting parameters that control the snake's tension and rigidity (the internal energy), and E_{ext} is the energy derived from the image, defined such that it is smaller at features of interest such as edges. Typical external energy functions for a grey scale image $I(x, y)$ include:

$$E_{ext} = -|\nabla I(x, y)|^2 \quad (4.28)$$

$$E_{ext} = -|\nabla (G_\sigma(x, y) * I(x, y))|^2 \quad (4.29)$$

where G_σ is a Gaussian function with standard deviation σ and ∇ is the gradient operator.

According to the Euler equation, a snake which minimizes E must satisfy:

$$\alpha \mathbf{x}''(s) - \beta \mathbf{x}''''(s) - \nabla E_{ext} = 0 \quad (4.30)$$

Equation 4.30 is solved by treating the snake as a function of time as well as s . The partial derivative of $\mathbf{x}$ with respect to time is given by:

$$\mathbf{x}_t(s, t) = \alpha \mathbf{x}''(s, t) - \beta \mathbf{x}''''(s, t) - \nabla E_{ext} \quad (4.31)$$

A numerical solution to equation 4.31 is found by separating the equation into

its constituent parts and iteratively solving the discrete system.

4.4.2.1 Ultrasound outline detection with snakes

Snakes have been used previously to find the outline of organs in medical ultrasound images. Ladak et al [117, 118] describe a semi-automatic technique for detecting the prostate boundary. The user gives four points close to the prostate boundary, namely the top and bottom of the mid-section, and the extremities of the two lobes. The snake is initialized to pass through these four points. The snake used is a discrete dynamic contour as presented by Lobregt et al [138], which is itself a variation on the original snake model developed by Terzopoulos et al [211] and others. The discrete dynamic contour is a series of points connected by straight line segments which deform according to internal and external forces. The internal forces depend on the local curvature of the contour and act to oppose high local curvature. The external forces are calculated based on the local gradient value near the contour and the gradient direction. Deformation of the contour proceeds until the internal and external forces reach a state of equilibrium. If equilibrium results in a contour which does not match the prostate boundary, then further user interaction is required to adjust the incorrect segments.

Levienaise-Obadia and Gee [128] use an active contour (snake) technique to detect the boundary of the gall bladder and the kidney. Although their algorithm has not been used on prostate images, it is included here since kidney ultrasound images and prostate ultrasound images have some similar characteristics¹⁰. The technique is similar to that applied by Ladak et al [117, 118] on prostate images, but provides a much more detailed description of the forces acting on the snake, which is not provided by Ladak. The snake contour used is based on a cubic b-spline function which provides inherent smoothness due to constraints set at the control points [43], therefore only external forces need to be calculated. The initial snake is determined from control points selected by the user, which lie close to the boundary.

The image is convolved with a Gaussian filter to reduce high frequency speckle. An initial search window is determined at a set distance from the contour, and both inside and outside the contour, to look for the most likely true boundary

¹⁰Kidney and prostate ultrasound images are distinct from most ultrasound images where outline detection might be applied, since both contain tissue to tissue boundaries, rather than fluid to tissue boundaries.

location. Two potential functions whose minima represent the likely true boundary position are calculated at a sample of points at equal intervals around the contour. The first function is a measure of the local grey-level gradient along the lines normal to the boundary:

$$P_{\nabla}(x, y) = -|\nabla(G(x, y) * I(x, y)) \cdot \mathbf{n}(s_i)| \quad (4.32)$$

where $I(x, y)$ is the image, $G(x, y)$ is the Gaussian filter, ∇ is the gradient operator, and $\mathbf{n}$ is the unit normal vector at sample points along the snake. The second function attempts to find differences in the local texture features inside and outside the boundary by analyzing the mean and variance of small (7x7) pixel regions: For each region centred at (x, y) , a vector is constructed as:

$$\mathbf{x}(x, y) = [\rho, \sigma^2]^T \quad (4.33)$$

where ρ is the mean of the pixel values in the region, and σ^2 is the variance. The mean vector and co-variance matrix are calculated for the set of regions inside and the set outside the boundary. A candidate region is evaluated as inside if:

$$-d_i + d_o > C \quad (4.34)$$

where,

$$d_i = \frac{1}{2} (\mathbf{x} - \bar{\mathbf{x}}_i) M_i^{-1} (\mathbf{x} - \bar{\mathbf{x}}_i)^T \quad (4.35)$$

$$d_o = \frac{1}{2} (\mathbf{x} - \bar{\mathbf{x}}_o) M_o^{-1} (\mathbf{x} - \bar{\mathbf{x}}_o)^T \quad (4.36)$$

$$C = \frac{1}{2} \ln \left| \frac{M_i}{M_o} \right| \quad (4.37)$$

and $\bar{\mathbf{x}}$ is the mean vector, M is the co-variance matrix, and the indices i and o represent inside and outside the boundary, respectively. The texture potential function is calculated as:

$$P_{\Omega}(x, y) = |C - (d_o - d_i)| \quad (4.38)$$

The weighted sum of two potential functions give a function whose minimum is the most likely local boundary position, where the weights are decided by analyzing previous outlines. For a sequence of images, the user describes the initial contour in the first image. The boundary search is conducted solely on the gradient of

the grey-level along the contour normals (i.e. the weight of the texture potential function is zero), and the user corrects any mis-positioned points of the found boundary. The b-spline control points are updated to minimise the discrepancy with the found boundary using linear least squares [43]. Texture statistics are found from regions on both sides of the boundary, and the texture potentials calculated. The weights are updated according to the final boundary and used in the next image, where the previous boundary is used as the initial contour. If user correction is required in subsequent images, then the potential function weights are updated accordingly.

4.4.2.2 Problems of snake based outline detection

The most obvious drawback of the two snake based techniques described in the previous section, is their reliance on user participation in the initialization of the snake position and possible user adaptation of the final result. For a truly automated outline detection algorithm, this user intervention is not permissible. User intervention in both cases is required since the snake is likely to converge to the wrong result if it is not initialized close to the correct boundary.

Several methods have been proposed to address the problem of contour initialization. Each method aims to extend the range over which the external forces act to draw the snake towards the correct outline. Internal pressure forces [41] can be used to push the snake beyond local minima to reach the correct result. However the pressure force can also cause the snake to balloon out at weak or missing edges of a true outline, as shown in the artificial example of figure 4.24. For prostate ultrasound images, breaks in the outline are common and the pressure required to overcome the local minima is high enough that the snake is unlikely to be “caught” by the correct outline. A final problem with pressurized snakes is that the pressure force can only act in one direction (internal or external pressure) so that the initial outline must be defined to be either completely inside or outside the true boundary. Two multi-resolution methods [127] have been defined which attempt to overcome initialization problems. The first method performs by adapting the snake on a series of images starting from a coarse pixellation (each successive level contains four times the number of pixels of the previous level) and successively using the result of the previous scale as an initialization on a finer pixellation. The second method used a Fourier description [199] of the outline with successively more harmonics. However, neither method is likely to cope with the problem that the prostate boundary is generally not the strongest




Figure 4.24: Snake with internal pressure 'balloons' out past the correct outline. a) Initial snake (green) and potential field (black). b) Intermediate stages (red) and final result (blue).

edge feature within the ultrasound image.

4.4.2.3 Gradient vector flow

Xu et al [226, 227] propose a technique known as *gradient vector flow* (GVF) to produce a better external force field for the snake. They claim that their method results in a snake which has no prior bias towards its movement direction (unlike the balloon snake), may be initialized inside, outside or across the actual outline, and has a large capture range - i.e. is capable of finding the correct outline at considerable distance. The gradient vector flow field is used to replace the E_{ext} term of equation 4.27.

An edge intensity image f is defined such that the magnitude of f is greater at features of interest within the image. Xu gives example functions for f as:

$$f(x, y) = -E_{ext}(x, y) \quad (4.39)$$

where E_{ext} is one of the functions given in equations 4.28 and 4.29. Three general properties of the edge intensity image are important in this context.

- The gradient of the edge intensity ∇f contains vectors which are normal to the edges in the image, and point towards those edges.
- Only vectors which are close to the edges have large magnitudes.

- In homogeneous image regions, the magnitude of ∇f is close to zero.

Property I is desirable since the snake will move towards the edges and converge to a stable result. However, properties II and III mean that the snake must be initialized close to the edges otherwise the snake will not move.

The gradient vector flow method aims to overcome the capture range problem by diffusing the large gradients at the edges out into the homogeneous regions, whilst maintaining the large gradients near the edges. The gradient vector flow field $\mathbf{v}$ is found by minimizing an energy function ε , where:

$$\mathbf{v}(x, y) = (u(x, y), v(x, y)) \quad (4.40)$$

$$\varepsilon = \int \int \mu \left(u_x^2 + u_y^2 + v_x^2 + v_y^2 \right) + |\nabla f|^2 |\mathbf{v} - \nabla f|^2 dx dy \quad (4.41)$$

Thus, in regions close to the edge features, where $|\nabla f|$ is large, the energy function is minimized by setting $\mathbf{v} = \nabla f$, maintaining the desired property I. In homogeneous regions, the second term is approximately zero, and the energy is dependent on the sum of squares of the partial derivatives of the intensity image. Therefore, in these regions the minimal energy condition smooths out the large changes in the gradient, and the vector field is interpolated from the boundaries of the homogeneous region. The parameter μ in equation 4.41 determines the sensitivity towards noise within the image.

Figure 4.25 shows the result of applying the gradient vector flow snake to the ultrasound image of figure 4.8. The edge intensity image f is calculated as:

$$f(x, y) = |\nabla (G_\sigma(x, y) * I(x, y))|^2 \quad (4.42)$$

using a Gaussian filter with a standard deviation of 8. The parameters for the snake internal energy terms are $\alpha = 0.6$ and $\beta = 0.0$, and the noise smoothing parameter $\mu = 0.2$ for the gradient vector flow calculation.

Xu gives several examples of the GVF snake working on synthesized images, and one example using an axial MRI image of the left ventricle of the heart (figure 4.26a)). Further examination of the software provided with these examples shows that the gradient vector flow field in each case is normalized by the magnitude of the vector field, *i.e.* the magnitude of *all* the vectors in the field is unity and only the direction of the vectors is taken into account. Careful initialization of the snake in the heart example ensures that the snake converges to the ventricle boundary, however slight perturbation of the initialization causes






Figure 4.25: Gradient vector flow snake result on the test ultrasound image



(a)



(b)

Figure 4.26: a) Axial MRI heart ventricle image. b) Correct GVF snake outline result.




Figure 4.27: a) Local maxima in the heart GVF field. b) Incorrect outline result caused by snagging. Initial outline (red), intermediate outline (green) and final outline black).

the snake to “snag” on a local maxima. Figure 4.26b) shows a magnified portion of the vector field and the initial, intermediary and final snake outlines, where careful initialization leads to the correct result. Figure 4.27 demonstrates the problem of localized “snagging” of the snake. When the gradient vector flow field is used without normalization, the snake shrinks under its own internal energies. Extensive experimentation with changing the various parameters of the snake and gradient vector flow calculation did not produce a solution to this problem. Email exchanges with the authors also failed to provide a satisfactory explanation/solution to the problem.

4.4.3 Statistical shape modelling

Arambula and Davies [5] present a method for finding prostate boundaries using statistical shape models (SSM) [45] to model the range of prostate shapes, and genetic algorithms (GA) to generate populations of boundaries from the model. The method is used for transurethral ultrasound images of the prostate. The first step in generating the SSM is to align a set of sample prostate boundaries, which have been outlined by an expert. Arambula performs the alignment by selecting a typical shape from the sample boundaries, and manually transforming the subsequent boundaries to match the typical shape as closely as possible. The

manual transformations are rigid - translation, rotation and scaling.

Each shape is represented as a vector of the x and y position of the boundary in equal angular increments around the boundary, relative to a common centre point (45 points are used). The mean outline is calculated from the means of points at each angle from the aligned sample outlines. The covariance matrix of the outline vectors is calculated, and the eigenvectors and eigenvalues of the covariance matrix are found. Thus, for the 45 point pairs (x and y) in the vectors, 90 eigenvectors are obtained which give the total variance of the sample boundaries from the mean boundary. However, only the ten eigenvectors with largest associated eigenvalues are found to give 95% of the total variance. Candidate outlines can be formed from the mean boundary plus a weighted sum of these ten eigenvectors. The weights are limited so that the points lie within three standard deviations of the mean [45], ensuring that all possible prostate shapes may be formed:

$$s = \bar{s} + \sum_{i=1}^{10} b_i \mathbf{p}_i \quad (4.43)$$

where s is the outline vector, $\bar{s}$ is the mean outline, $\mathbf{p}_i$ are the principal vectors and b_i their respective weights.

The weights, scale, rotation and translation information representing the boundary are coded as a binary string, where the range of each parameter determines the number of bits required for that parameter. Arambula uses strings of 53 bits; these are the chromosomes for the GA. An initial population of chromosomes is constructed with random values for each parameter (within the allowed ranges). The GA uses crossover and mutation to generate new chromosomes from the initial population. The particular values for the GA parameters are assigned according to the work by Grefenstette et al [83], which give good performance, and stochastic universal sampling [6] is used to select chromosomes for parenting the next generation.

The GA requires a performance measure (normally referred to as an objective function in GA) so that the suitability of a particular solution may be determined. Arambula uses the following scheme for the objective function. Points at 1mm intervals along the boundary are selected, and the mean grey level evaluated along the line perpendicular to the boundary at each point, for 15 pixels inside, and 15 pixels outside the boundary. Then the suitability d of a particular point is given by:

$$d = 255 - out + in \quad (4.44)$$

$$out = \frac{1}{15} \sum_{n=1}^{15} x_{o,n} \quad (4.45)$$

$$in = \frac{1}{15} \sum_{n=1}^{15} x_{i,n} \quad (4.46)$$

where x_o and x_i are the pixels in the line which are outside and inside the boundary respectively. Therefore, points on a dark to bright transition are scored highly. The objective function is found as the mean of the point measures scaled by the inverse of the mean outside grey level, to make the test more sensitive to the bright halo around the prostate capsule:

$$ObjFun = \frac{1}{m\sqrt{out}} \sum_{n=1}^m d_n \quad (4.47)$$

$$\overline{out} = \frac{1}{m} \sum_{n=1}^m out_n \quad (4.48)$$

This test is used for the objective function of the GA; the lower the result, the better the fit of the boundary to the ultrasound image.

Wu [224, 225] applies similar techniques of boundary modelling and GA for finding the prostate capsule in transrectal ultrasound images. An adaptation of PDM called feature modelling (FM) is used where the boundary points are labelled according to the polar co-ordinate system defined by the boundary centroid. The first point in the labelling is the boundary point whose angle is perpendicular (and to the right) of the line joining the boundary centroid and the probe centre. A set of sample hand segmented images is used to construct the FM, without the manual alignment used in [5], rather, the physical dimensions of the points are used (distances in mm from the centroid).

The vector of each boundary is made from the distance between boundary centroid and probe centre, the angle between the line joining probe centre to boundary centroid and the vertical, and the boundary radii. Chromosomes are made for the nine eigenvectors with the largest associated eigenvalues (representing over 99% of the variance) and the angle between the vertical and the line joining the probe centre to the boundary centroid. The objective function for the GA is the sum of the inner product of the normal boundary direction (pointing inward) and the local grey-level gradient direction, multiplied by the local grey-level gradient magnitude. Boundary fitness is proportional to the objective function result.

Figure 4.28 shows the result of applying Wu's FM and GA process to the test prostate ultrasound image. The FM was calculated from a set of expertly out-




Figure 4.28: Feature modelling and genetic algorithm result on the test ultrasound image.

lined ultrasound images (see chapter 5, section 5.3.4), which includes the test image of figure 4.8. The genetic algorithm uses binary encoding with incremental crowd replacement, a population of 400 outlines iterated for 2500 generations, with a 90% crossover rate and 1% mutation rate. These parameters were chosen according to the best results given in Wu's paper.

4.5 Region based detection schemes

This section describes algorithms which attempt to segment an image based on some similarity criterion of the pixels contained in a region. The techniques are not necessarily used predominantly for boundary detection, however they can be used to find the prostate outline if they can reliably segment a prostate ultrasound image into prostate and non-prostate regions.

4.5.1 General region-based image processing

4.5.1.1 Region growing

Region growing is the simplest of the region-based segmentation techniques. The algorithm starts from a set of initial points (seeds) representing distinct image regions which are grown until the regions cover the complete image. The seeds are usually chosen by the user. At each iteration of the algorithm unassigned pixels which border a region are assessed as to whether they are likely to belong to the region. The assessment is often based on a comparison of the pixel's value to the mean value of the pixels contained in the region. The algorithm stops once no more pixels are added to any region during one iteration.

Figure 4.29 shows the result of applying a simple region growing algorithm to the test ultrasound image of figure 4.6. A single seed point is used which is inside the prostate (the same point as used in the edge processing discussed in section 4.3.6). Therefore, an ideal result would assign all pixels within the prostate to the region of the seed, and leave all other pixels unassigned. The test image is first filtered with a Gaussian kernel (with $\sigma = 8$) to reduce the influence of speckle on the segmentation.

Region growing algorithms are unlikely to perform well on ultrasound images due to the speckle inherent within the image. However, one example of region growing was found in the literature and is discussed in section 4.5.4.

4.5.1.2 Merging

One of the problems of region growing algorithms is that they are highly dependent on the initialization of the seed points. If two or more seed points are defined in a single region, then the result will contain more classifications than necessary. A combined merging/growing algorithm [129] may be used to merge regions which have similar properties. The algorithm starts from a single point in the image (often the top left corner) and assigns this point as a region. The algorithms then scans through each point in the image and attempts to add each candidate point to one of the regions assigned to its neighbours, based on some statistical homogeneity criterion. If the candidate cannot be added to its neighbours, a search is made of all the existing regions to see if the candidate may be added to any regions regardless of the geometrical proximity. If no region satisfies the homogeneity criterion, the candidate pixel is assigned to a new region. In the



Figure 4.29: Region growing result for the ultrasound image of figure 4.8.

case of more than one region satisfying the homogeneity criterion, the pixel is added to the region where the homogeneity is greatest.

Although the merging algorithm may assign unconnected pixels to the same region, this is not necessarily a problem if all the pixels which are part of the prostate fall into the same class. Figure 4.30 shows the result of applying the merging algorithm to the test prostate ultrasound image which is pre-filtered with a Gaussian kernel to reduce the effects of speckle. The homogeneity criterion is calculated as the ratio of the region mean to the region standard deviation, as used by Koo and Park [111] in the homogeneous region growing median filter (see appendix A, section A.1.2).

4.5.1.3 Splitting

Region splitting approaches the segmentation from an opposite viewpoint. The image is first assumed to be homogeneous for the statistical criterion used. If this is not true, the image is split into regions, which are also subsequently split if they are also not homogeneous. A common example of a region splitting technique is *quadtrees decomposition* [9]. The initial square image is subdivided into four squares, and the homogeneity test applied. Each of the sub-regions which fails the homogeneity test is further split into four square regions. The process




Figure 4.30: Merging algorithm applied to the prostate ultrasound test image.

continues until all the regions are homogeneous. Figure 4.31 gives an example region configuration and its associated quadtree. The result of applying quadtree decomposition to the test prostate ultrasound image is given in figure 4.32. The input image is pre-filtered with a Gaussian kernel to reduce the effects of speckle on the region splitting. The homogeneity criterion is satisfied if the ratio of region mean over region standard deviation is less than a threshold.

4.5.1.4 Split and merge

Region splitting often leads to a segmentation where many adjacent but distinct regions are homogeneous. Ideally, these regions should be merged. In the example quadtree decomposition given in figure 4.31, there are two distinct regions (represented by light or dark grey pixels), but these regions are further split into seven and nine regions respectively. The split and merge algorithm [101] attempts to overcome this problem.

At each iteration of the algorithm, both splitting and merging are included. A region is split into four sub-regions if the homogeneity criterion fails. Each adjacent pair of regions is then checked to see if they are homogeneous and are merged if the criterion is satisfied. The algorithm proceeds until no more splitting or merging operations are possible.




Figure 4.31: Quadtree decomposition. a) Image sub-divided into homogeneous regions. b) Associated quadtree structure.




Figure 4.32: Quadtree decomposition applied to the test ultrasound image.




Figure 4.33: Split and merge algorithm result.

Figure 4.33 shows the result of applying the split and merge algorithm to figure 4.8, after pre-filtering with a Gaussian kernel to reduce the influence of speckle. The homogeneity criterion compares the ratio of region mean to region standard deviation against a user-defined threshold.

4.5.2 Texture based segmentation

Grey-level texture statistics have been used by several authors to detect different types of tissue in ultrasound images. Basset et al [8] use texture co-occurrence matrices in an attempt to distinguish carcinoma tissue from benign tissue in ultrasound images of the prostate. However, only Richard and Keen [178] have attempted to segment the prostate using texture as the basis for deciding the prostate and non-prostate regions. Given a successful segmentation, it is trivial to determine the prostate boundary by finding the edge of the segmented region.

Richard and Keen assess four different textural energies present in the ultrasound images. These four texture energies are found by first convolving the input image with four texture masks given by Laws [124]. For each of the resulting images, an edge-preserving smoothing operation is performed. The variance of the four quadrants in a window around each pixel is calculated. The pixel is then replaced by the mean value of pixels in the quadrant which has the lowest variance. These

four images reflect the different textural energies of each pixel.

A vector is formed for each pixel containing the values of the four textural energies. The mean vector and standard deviation are calculated and the Euclidean distance from the mean is calculated for all the vectors and the standard deviation. The first seed point is chosen as the pixel whose vector lies closest to the mean. An iterative process assigns further seed points by calculating the distance between each vector and the current seed points, and checking that the distance to the closest seed point is less than k times the standard deviation distance. If the distance is greater, the pixel is added as a new seed point, and the next iteration started. The process ends when all pixels are within the range of k times the standard deviation distance from any seed. The number of seed points determines the number of different texture classes within the image. Hence the value of k controls the number of classes and is a user-defined parameter.

All pixels are assigned as belonging to the texture class of the nearest seed point. Each seed is then changed to be the mean of the pixels assigned to its class, and the pixels reassigned according to the nearest new seed. The process repeats until no seed moves more than a threshold distance between iterations. An initial probability of each vector belonging to each class is calculated by assuming a Gaussian distribution for each class, and using the mean and variance of the vectors in a class. For a pixel i and a class λ , the initial probability of the pixel belonging to the class is denoted $P_i^0(\lambda)$. The mean probability of each class is found as:

$$\hat{P}(\lambda) = \frac{1}{n} \sum_{i=1}^n P_i^0(\lambda) \quad (4.49)$$

where n is the number of pixels in the image.

The joint probability of two neighbours in a particular orientation (e.g. the candidate pixel i and the pixel immediately to its right i_R) belonging to two particular classes, λ and λ_R is calculated as:

$$\hat{P}_{i,i_R}(\lambda, \lambda_R) = \frac{1}{n} \sum_{i=1}^n P_i^0(\lambda) P_{i_R}^0(\lambda_R) \quad (4.50)$$

Equation 4.50 is evaluated for all possible classes λ and λ' , including $\lambda = \lambda'$, and for all eight neighbour orientations:

$$\hat{P}_{i,i+\theta}(\lambda, \lambda') = \frac{1}{n} \sum_{i=1}^n P_i^0(\lambda) P_{i+\theta}^0(\lambda') \quad \theta \in \Delta \quad (4.51)$$

where $i + \theta$ is a specific neighbour of pixel i and Δ is the set of neighbours of

pixel i .

The compatibility coefficients for each orientation and pair of classes are calculated by dividing the joint probability of the two classes in that orientation by the mean probabilities of the two classes:

$$r_{i,i+\theta}(\lambda, \lambda') = \frac{\hat{P}_{i,j+\theta}(\lambda, \lambda')}{\hat{P}(\lambda) \hat{P}(\lambda')} \quad (4.52)$$

The next process in the algorithm is an iterative relaxation of the probability labellings. For the $j + 1^{th}$ iteration, the likelihood of a pixel belonging to one class given that a neighbour (in a particular orientation) is a member of another class is calculated as:

$$S_{i,i+\theta}^{j+1}(\lambda) = \sum_{\lambda' \in \Delta} \hat{P}_i^j(\lambda) \hat{P}_{i+\theta}^j(\lambda') r_{i,i+\theta}(\lambda, \lambda') \quad (4.53)$$

This is normalized by the sum of the measures for all possible classes of the neighbour:

$$q_{i,i+\theta}^{j+1}(\lambda) = \frac{S_{i,i+\theta}^{j+1}(\lambda)}{\sum_{\lambda' \in \Delta} S_{i,i+\theta}^{j+1}(\lambda')} \quad (4.54)$$

Finally, the probability of each vector belonging to a class is found as the mean of the likelihood of the class from each of the eight neighbours:

$$P_i^{j+1}(\lambda) = \frac{1}{8} \sum_{\theta \in \Delta} q_{i,i+\theta}^{j+1}(\lambda)$$

The process iterates for a set number of turns using the updated probabilities. The final result is created by labelling each pixel according to its most probable class. Figure 4.34 shows the result of applying the texture based segmentation algorithm to the ultrasound image from figure 4.6.

4.5.3 Neural networks

Prater and Richard [166] use feed-forward neural network techniques to give a segmentation of the input image into prostate and non-prostate regions. Three architectures, each with three layers of neurons, are presented using different inputs into the network. Figure 4.35 shows the ultrasound image and ideal segmentation used to train the networks. The first neural network consists of an input layer containing five neurons, a 'hidden' layer of six neurons and a final




Figure 4.34: Texture based segmentation result for the test ultrasound image.



Figure 4.35: Training image and its ideal segmentation.




Figure 4.36: Inputs for a given pixel in the 5-6-2 neural network

output layer containing two complementary neurons. Thus for each pixel, five inputs are made, and the network decides on whether that pixel is part of the prostate or not – indicated by which of the output neurons is activated.

For every pixel in the image the mean value of the 9 (3 by 3) pixel neighbourhood centred on the pixel is calculated. When testing a candidate pixel the first input to the network is this mean value. The remaining four inputs are the maximum value of the means in each direction (north, east, south, west) around the candidate pixel. Figure 4.36 shows the inputs for a given pixel.

The network is trained using a back-propagation technique. The training image is hand segmented to produce a binary image containing prostate and non-prostate regions (see figure 4.35). Small random weights are assigned to the network in the initial stage. Calculations are made for the five input parameters for a candidate pixel in the training image and the resulting output compared to the binary segmentation result. Adjustments are made to the weights when the output conflicts with the binary image. Several thousand input and corresponding output sets are used to train the network. Figure 4.37 shows the result of applying the network to the ultrasound image of figure 4.35, following training of the network on the same image using 1334 pixels as training patterns, taken at fixed intervals across the image.



Figure 4.37: 5-6-2 network result on the ultrasound image of figure 4.35. The network was trained on the same image as the input.

The result for the 5-6-2 network exhibits stripe artefacts and suffers from a misclassified region in the centre of the prostate. To overcome these problems a second network has been used, incorporating substantially more inputs. The input layer contains 132 neurons, the hidden layer has 30 neurons, and the output has two complementary neurons. For each pixel, the mean of its 9 pixel neighbourhood, the means of adjacent 3 by 9 pixel neighbourhoods in the same row as the candidate, the means of adjacent 9x3 neighbourhoods in the same column as the candidate, and the means of 3x3 neighbourhoods forming a square around the candidate pixel are used as inputs. The inputs to the network for a given pixel are shown in figure 4.38. The network is trained using 3000 input/output training patterns using figure 4.35 as the training image. The result of applying the neural network to this image is shown in figure 4.39.

The second network is also tested after training with substantially more input training patterns. The resulting network gives an almost perfect segmentation of the training image. Figure 4.40 shows the result of using 15908 training patterns from the input image, and applying the trained neural network to it's training image.

The problem with the second network is that the large number of neurons involved causes the training and processing stages to be excessive. Analysis of the assigned network weights shows that many of the neurons are superfluous – i.e. they have zero weights and play no part in the decision making process. The final network




Figure 4.38: Inputs for the 132-30-2 neural network.



Figure 4.39: 132-30-2 neural network result.



Figure 4.40: 132-30-2 neural network result after extensive training with 15908 training patterns.

uses wider neighbourhoods along the rows, and taller neighbourhoods along the columns, and fewer 3x3 neighbourhoods in a circle formation around the candidate pixel in order to reduce the number of network inputs. Figure 4.41 shows the inputs for a given pixel. This network has a 61-10-2 neuron architecture, and is trained using 7500 training patterns. The result of applying this network to the training ultrasound image given in figure 4.35 is given in figure 4.42. Once again, this is an almost perfect segmentation result when compared to the training segmentation pattern. However, the network does not perform as well when applied to images upon which it has not been trained. Figure 4.43 shows the results of applying the third network to two new images. In each case the boundary of the correct segmentation is given. Both these images are gathered from the same prostate as the training image; the image on the left of figure 4.43 is an adjacent slice to the training image (1mm away), the image on the right is taken at 10mm from the training image.

Jansen [105] provides further work using the Prater and Richard neural network techniques. This work uses the 61-10-2 neural network architecture to segment transurethral ultrasound images of the prostate. The novel aspect of the work is the application of the neural network to a set of images from different patients rather than a single patient. Initial trials show that a network trained on an image from one patient fails to produce a plausible result on images from other patients. Jansen tries to overcome this difficulty by training the network on a series of images – however the results showed that the network “forgets” prior




Figure 4.41: Inputs for a given pixel for the 61-10-2 neural network.



Figure 4.42: Result of the 61-10-2 neural network trained and tested on the ultrasound image of figure 4.35.



Figure 4.43: Results for the 61-10-2 network applied to non-training images. a) Image adjacent image to the training image. b) 10mm from the training image.

information and only performs correctly when applied to the most recent training image.

4.5.4 Gradient direction based region growing

Crivianu-Gaita et al [47] propose a very simple scheme for segmenting the prostate. The input image is median filtered to reduce speckle and the user specifies a point to mark the centre of the prostate. The algorithm uses a region growing technique to label pixels as forming part of the prostate. Starting from the specified centre, pixels are added according to the labelling of their neighbours. A pixel is added to the region if at least seventy percent of its neighbours are also in the region (the neighbourhood size and shape are not specified), and not added if less than thirty percent of the neighbours are in the region. For pixels where the number of neighbouring region labelled pixels is between the two bounds a further criterion is used. The local gradient is calculated along the line between the candidate pixel and the prostate centre, and compared to the gradient between the candidate and the centre. When the cosine of the angle between the gradients is close to one (i.e. the gradients are close) the point is labelled as part of the prostate – as shown in figure 4.44. The region-growing algorithm iterates until no more pixels are added to the region.

The algorithm assumes that the inside of the prostate can be thought of as a




Figure 4.44: Local gradient compared to gradient with respect to the prostate centre.

grey-level valley completely surrounded by grey-level peaks at the prostate edge. However, the algorithm fails when the peaks (prostate edges) are not completely continuous, allowing the region to grow on outside of the prostate. No result is given for the algorithm since it failed even on the high quality input image of figure 4.6.

4.6 Discussion

The objective of applying many different techniques to a single high quality prostate ultrasound image is to identify techniques which warrant further consideration. The approach adopted in selecting which techniques to test has been largely holistic in nature. Many different techniques have been used which may be considered to be “standard” approaches in the image processing field. However, the unusually high signal-to-noise ratio present in ultrasound images and the presence of artefacts cause many of the standard techniques to fail, even for a high quality (in prostate ultrasound terms) image such as the test image (figure 4.6).

Detection of outlines in ultrasound images has been attempted by many researchers. However, in most of these applications, the object to be detected is either filled with fluid or surrounded by fluid, giving a relatively distinct boundary within the image compared to those of changes due to differing tissue. Prostate ultrasound images do not exhibit such clear distinctions between the prostate and non-prostate regions. Therefore, most of the other techniques which have been assessed are based on previous work aimed at detecting the prostate boundary in ultrasound images.

	Edge based outline detection	Model based outline detection	Region based outline detection
Inputs	Raw image pixel data	Raw image pixel data Boundary model	Raw image pixel data System training
Low-level processing	Noise suppression First derivative Second deriavtive zero crossings Edge intensity	First derivative Noise suppression	Noise suppression Statistics extraction
Intermediate level processing	Thresholding Edge map	Model variation Fitness maximisation	Classification
High level processing	Edge selection Edge linking	Boundary searching	Region labelling Region selection Boundary extraction
Outputs	Boundary description	Boundary description	Boundary description

Figure 4.45: Edge based, model based and region based processing pipelines for extracting prostate outlines from ultrasound images.

The algorithms which have been used to extract the prostate outline fall into three categories; edge based, model based and region based techniques. The image processing pipeline for each of these categories is summarized in figure 4.45. Each process has been split into three levels of processing: low, intermediate and high. Low level processing covers operations which are carried out on pixels in the image, such as noise smoothing and gradient extraction. Processing at the intermediate level essentially covers the aspect of grouping entities together. These entities can be the pixels themselves, such as extracting edge segments from an edge intensity image, or clusters or regions of pixels which may be classified as belonging to a group according to some common criterion. High level processing produces a functional description of the objects within the image. This requires

some knowledge about the objects present within the image, such as the likely shape and position. In the case of the boundary detection algorithms, the high level processing results in a description of the outline of the prostate. The definitions of the three levels are somewhat loose in nature, and there is some overlap between them.

All the algorithms presented in this chapter have been tested on a single ultrasound image of reasonable quality, to varying degrees of success. It is clearly not valid to choose a single algorithm as being the best choice for outlining the prostate based on the result of a single test. However, it is a useful exercise since it provides a good indication of which algorithms should go forward for large scale testing on a dataset of ultrasound images from different patients and of varying quality. This large scale testing forms the basis of the work described in the next chapter. The remaining sections of this chapter discuss the qualitative results of each algorithm, and highlight those which are potentially useful.

4.6.1 Edge based detection schemes

The edge based detection schemes follow a data-driven, bottom-up approach. That is, the techniques start from the pixels in the ultrasound image, extract segments by grouping pixels based on some combination of the first and second derivatives of the image, and finally use prior knowledge of the likely shape and position of the object of interest (the prostate) to form a final, closed boundary. The functions which perform each of these stages require various parameters to be set – correct setting of the parameters is essential for a good boundary result. The parameter setting for each algorithm has been tuned to give the best result for the test ultrasound image of figure 4.6. These parameter values will be used in the large scale trial - *i.e.* individual tuning for each image in the trial is not acceptable.

4.6.1.1 Laplacian of Gaussian

The Laplacian of Gaussian operator described in section 4.3.2 performs well at the dual tasks of smoothing out the speckle in the image and detecting the edges. The edges lie on closed curves which do not touch or intersect, therefore the prostate edge lies on more than a single contour. This causes difficulty in selecting the appropriate edges which correspond to the boundary, and the places where

those edges deviate from the boundary. Therefore the approaches suggested by Gomes [80] and Richard [177] are not practicable.

Incorporating the edge strength information as described by Zhao [234] gives an improved result since the edges of the prostate are highlighted in comparison to many of the other contours at the zero-crossings of the Laplacian of Gaussian. The hysteresis thresholding and knowledge based edge removal processes lead to an edge map which contains a large proportion of the prostate outline. This leads to a final outline which is close to the true boundary, and is the best of all the results from the edge-based schemes.

The Laplacian of Gaussian based scheme is deemed successful in the initial assessment and proceeds as a candidate for the the large scale trial.

4.6.1.2 Edge strength and second derivative in gradient direction

The result from the edge strength and second derivative in gradient direction algorithm is also close to the correct boundary. The result is slightly enlarged in the lower left portion of the boundary and is overall curve is less smooth than the results of the other edge based detection schemes. However, the algorithm warrants further analysis in the large scale trial.

Use of the integrated edge maps variation [3], based on the neighbourhood standard deviation, gives no improvement to the result for the test image, since the local standard deviation only exceeds the global standard deviation in areas of the image which are not part of the prostate. Manually selecting the threshold so that more of the prostate is filtered with a smaller filter produced no discernible improvement in the final result.

4.6.1.3 Sticks and anisotropic diffusion filtering

The sticks and anisotropic diffusion result suffers from a slightly underestimated position of the outline in the lower left quadrant. This is due to the spline interpolation – the algorithm failed to find points corresponding to the edge in this region. However, the result is close to the true boundary and the technique proceeds as a candidate for the full-scale trial.

4.6.1.4 Radial-bas relief

The radial-bas relief technique gives an under-estimate of the outline in the lower regions of the prostate, and is the poorest of edge-based detection schemes. However, the result is still reasonably close to the correct boundary, and its errors verge on the safer condition of underestimating the size of the prostate, rather than exceeding the true boundary. Therefore, the radial-bas relief based scheme is included in the full-scale trial.

4.6.2 Model based detection schemes

The basis of the model based detection schemes is a description of the shape of the object to be detected. The major difference between each of the schemes identified in section 4.4, is the “tightness” of the model – *i.e.* the amount of shape variation allowed in the model. For the Hough transform, the allowed variation is small, corresponding to a single shape which may be altered through rigid transformations of scale and position. Therefore, the Hough transform is unsuited to the task of detecting prostate outlines where the correct outline varies between images.

4.6.2.1 Snake-based techniques

The snake model gives the loosest description of the allowable shape, constraining the shape to be smooth, but offset against the attractive forces calculated from the underlying image features. The snake-based detection schemes described in section 4.4.2, all fail to produce a valid result on the test image, due to the difficulties of assigning an initial snake which is close to the true boundary and the failure of the snakes to converge on the correct (weak) boundary. The failure to produce a result on a test image deliberately chosen for its relatively clear boundary means that the snake based techniques are deemed to have failed, and will not be pursued in the large scale testing.

4.6.2.2 Statistical shape modelling techniques

The basis of the SSM is that the model should maintain the characteristics of the represented objects, but allow deformation to include the variation which is present in the objects. Cootes et al [45] describe two fundamental characteristics

which the models should exhibit. Firstly, a model must be general – *i.e.* it must be capable of generating any example of the object which the model represents. Secondly, the model must be specific – it should only generate legal examples of the objects.

Initial results using the FM technique are promising.

4.6.3 Region based detection schemes

4.6.3.1 Texture based region segmentation

The texture based region segmentation algorithm produces very poor results. The regions cross the prostate outline, therefore there is no way to use this technique to determine the correct outline of the prostate. Given the failure of the algorithm on a high quality image such as the test image, it is highly unlikely that the algorithm will produce better results on other images. The texture-based region segmentation algorithm will not be included in further trials.

4.6.3.2 Neural network based segmentation

The neural network techniques proposed by Prater and Richard give good results when applied to the image upon which the network has been trained. However, the results quickly degrade when the network is applied to images upon which it has not been trained, even for images which are similar to the training image by virtue of being taken from the same prostate. The work of Jansen showed that the networks tend to learn how to segment a particular image, but do not learn the general rules which would allow different images to be correctly processed.

The appeal of neural networks for the image processing is based upon their similarity to the processes which must be occurring in the human visual cortex to allow a human expert to identify the boundary of the prostate in an ultrasound image. These processes are not clearly understood, but the general idea behind neural networks is that they may “learn” the processes through training. However, whilst the precise factors which a human uses to decide on the boundary position are not known, it is highly unlikely that they include deciding the fate of a single pixel based on a few values of region means surrounding that pixel.

Rather than assess each pixel individually, it is preferable for the whole image to be considered as an input to the network. This would give a far closer analogy

to the information used by a human to determine the position of the outline. However, for a 512 by 512 pixel image such as the test ultrasound image, the network would require 262144 input neurons, 262144 output neurons and a substantial number of intermediate hidden layer neurons. Despite the improvements in computing power which have occurred since the Prater and Richard work (where segmentation on their 61-10-2 network took 46 minutes on a high powered workstation), such a network is still beyond current capabilities. Dedicated neural network hardware does exist, but the top of the range boards are limited to 10000 neurons, falling far short of the requirements. The number of neurons required could be reduced by sub-sampling the image, but the requirements would still exceed current capabilities.

It is the belief of this author that neural network techniques will become feasible as computing power continues to improve at a fast rate. The analogy of the neural network to the human brain gives the clearest indication of future success – it is relatively simple for a human subject to become adept at outlining the prostate in previously unseen images, given a few example images and their correct segmentation. However, for the present study, the neural network technique is deemed unsuitable for the task of finding the prostate outline and will not be included in further trials.

4.6.3.3 Gradient direction based region growing

The gradient direction based region growing algorithm relies on a continuous, strong edge around the prostate. However, the region growing “leaks” past the prostate edge even on the high quality prostate ultrasound image used for testing. Therefore the method is not suitable as a candidate for further trials.

4.7 Conclusion

Many image processing techniques which are potential candidates for finding the prostate outline in ultrasound images have been described in this chapter. The algorithms cover a broad range from standard techniques found in any introductory image processing text [88, 168] through to algorithms specific to ultrasound image processing, and methods explicitly for prostate boundary detection in ultrasound. Potential algorithms have been implemented and tested on a single prostate ultrasound image judged to be of high quality in terms of the clarity of

Scheme	Successful	Failed
Edge based detection	Laplacian of Gaussian Sticks & ADF Edge strength & SDGD Radial-bas relief	Sobel operator Prewitt operator Roberts operator Second derivative
Model based detection	Statistical shape model	Region growing Region merging Region splitting Split & merge <i>Hough transform</i> Snakes
Region based		Texture based <i>Neural network</i> Gradient direction region growing

Table 4.1: Initial outlining results. Algorithms on the left show promise and are included in the large-scale trial. The italicised algorithms were not implemented.

the prostate outline. Two algorithms were judged to be infeasible for finding the prostate outline without implementation¹¹.

Table 4.1 gives a summary of the results of the initial testing of the algorithms. Algorithms in the centre column show promise in the task of finding the outline of the prostate in ultrasound images, and are candidates for further testing on a set of ultrasound images from different patients, and exhibiting variations in quality. This forms the basis of the work described in the next chapter.

¹¹The Hough transform was not considered since it cannot cope with the shape variation present in the prostate. The neural network approach was dismissed on the basis of the presented results by Prater and Richard [166], and the work of Jansen [105].

Chapter 5

Assessment of outlining algorithms

5.1 Introduction

Chapter 4 introduced many of the algorithms previously described in the literature which have been used to find the outline of the prostate in ultrasound images. Many of these algorithms were implemented and assessed on the basis of applying the algorithm to a single prostate ultrasound image. The test image was chosen as an example of the “best” quality image likely to be encountered in typical routine transrectal ultrasound examinations.

The outline results using this test image demonstrate that several of the algorithms are worthy of further investigation. Clearly, applying an algorithm to a single image is not a rigorous test of its general suitability for outlining, particularly when the chosen image is taken from the high end of the range of image quality. This is a criticism of most of the reviewed papers in chapter 4 with the exception of the work by Pathak et al [161, 162] and Wu [224].

This chapter describes the work undertaken to evaluate the success of the promising algorithms (chapter 4, section 4.7) applied to a library of images which covers the range of quality observed in routine ultrasound examinations. Acquisition of the data set is described in section 5.2. Section 5.3 discusses the problem of deciding what the correct outline result for a given image should be, and assesses various potential methods for attaining the correct result. A discussion of various methods of comparing the ideal result with the algorithms’ results is given in section 5.4. Statistical analyses for assessing the results across the range of images are discussed in section 5.5. Finally, results of the statistical analysis for the five promising algorithms are given and discussed in section 5.6.

5.2 Image data

The prostate varies in shape and size from one person to another. In addition, the tissues which constitute the prostate and its surrounds vary between patients. Calcifications are known to occur within the prostate resulting in hyper-echoic regions on an ultrasound image. Cysts can produce regions of hypo-echoicity. These variations mean that the appearance of prostate ultrasound images can vary considerably with different patients. An automated outlining algorithm must be able to produce a reasonable result regardless of the variations found in the input ultrasound images. Testing the algorithms requires an image data set for use as inputs to the algorithms.

Sixty two ultrasound images have been collected during routine ultrasound examinations on sixty two different patients presenting symptoms of benign prostatic hyperplasia. The scanner used for the ultrasound examinations was a Kretz 420 Combison¹ machine. The scans were acquired using a 7.5 MHz transrectal ultrasound probe containing a mechanically spun transducer providing a 350° degree polar image around the probe. The ultrasound probe was set to acquire images in transverse mode. The video signal from the ultrasound scanner was connected to a PC with an on-board Matrox Meteor² frame grabber card. Each image was captured at the mid-gland portion of the prostate – the probe position at which to capture the image being decided by choosing the image with the largest cross-sectional area of the prostate. The scale and gains of the images were set by the ultrasound examiner during each examination – hence they are different in the collected scans.

The images were not selected according to perceived image quality, size of gland *etc.* The image set is therefore not biased in any way other than that which arises from the fact that the ‘capture image’ (or equivalent) button was pressed at a particular moment and the settings on the ultrasound machine were selected by an expert to achieve an image that conveyed the most or clearest information to a human observer. The automated outlining techniques discussed here are meant to work without the luxury of this initial quality filter.

The images in the data set are typical of the range in quality observed during routine prostate ultrasound examinations. Figure 5.1 shows three example images from the data set, qualitatively chosen as examples of low, medium and high quality in terms of the ease of identification of the prostate boundary.

¹KretzTechnik AG, Zipf, Austria.

²Matrox Electronic Systems Limited, Quebec, Canada.




Figure 5.1: Examples from the ultrasound image data set. a) Low quality. b) Medium quality. c) High quality.

5.3 Reference outlines

In order to assess the performance of any particular outlining algorithm, it is necessary to have a benchmark against which the algorithm's result may be compared. The obvious benchmark is to compare the algorithm's result with the correct outline, which is termed the *reference outline* (RO). However, for a given ultrasound image, it is not possible to obtain a definitive answer to the question "what is the correct prostate outline?". The following sections discuss various methods which have been/might be tried in order to find the best answer to this question.

5.3.1 Direct measurement of the outline

In an ideal situation, the prostate would be available following the ultrasound examination such that the prostate could be cut open at the plane of the ultrasound image and the outline measured directly. However, following the majority of prostate ultrasound examinations the prostate is not available. This author does have experience of one experiment to perform a direct measurement on a prostate from a patient undergoing radical prostatectomy³. The ultrasound examination was conducted in the operating theatre immediately prior to the prostate removal. The mid-gland region of the prostate was identified on the ultrasound scan, at which point the ultrasound probe was mechanically clamped to prevent further probe movement. The ultrasound image was then captured using a PC with an on-board frame grabber card connected to the ultrasound machine's video output.

A needle was inserted into the prostate through the perineum until the needle's tip just became visible in the ultrasound scan. A small quantity of indigo blue dye was injected into the prostate to mark the position at which the ultrasound image was acquired. The needle and ultrasound probe were removed and the operation was completed.

The excised prostate was placed in a jar of formaldehyde and sent for histological analysis (the analysis taking place directly following the operation). One millimetre sections of the prostate were dissected by hand. The sections showing traces of the dye were cut by microtome, histologically stained and mounted.

³Radical prostatectomy is an open surgery technique where the whole prostate, including the capsule, is removed. It is most commonly performed on patients with prostate cancer rather than benign prostatic hyperplasia.




Figure 5.2: A typical transurethral ultrasound image of the prostate (the prostate outline is shown in red and the urethra in green).

The intention of the experiment was to measure the prostate outline on the histological slides. This outline would then be compared to the information contained in the corresponding ultrasound image. However, a number of concerns were encountered which prevent the histological slide outline from being used as a valid reference outline for the ultrasound image.

Despite the precautions taken to obtain the ultrasound image immediately prior to the prostate's removal, there were clear indications that the prostate shape and size changed between ultrasound scanning and histological slicing. The presence of the ultrasound probe within the rectum caused distortion of the prostate – such distortion is clearly visible in the majority of transrectal ultrasound images as a concave region in the vicinity of the probe. These concave regions are not observed in transurethral ultrasound images – a typical transurethral ultrasound image is shown in figure 5.2. Further gross movement and local deformation was caused by the insertion of a needle into the prostate.⁴

⁴The problem of prostate movement and deformation following needle insertion is well documented as a problem for brachytherapy – a treatment for prostate cancer where needles are used to implant radiological seeds within the prostate.

Within the body, the prostate was constrained by a surrounding matrix of muscle, fat and ligaments. Once the prostate was removed from the body these constraints were no longer in place with the result that the prostate changed shape. The fixation of the prostate tissue using formaldehyde caused dehydration of the tissue, resulting in significant shrinkage of the prostate tissue.

Since there were no structures visible in the ultrasound image which could be registered to anatomical features, orientating the scalpel to cut slices parallel to the image plane was not practicable beyond rough estimation. During the slicing, the prostate was observed to deform under the action of the scalpel. This resulted in slices whose surfaces were not planar. The marker dye was found to have permeated through the prostate tissue so that had the slicing been fortunate enough to coincide with the ultrasound image plane, it would still not have been possible to tell. Therefore, it was not possible to have any confidence that the cut surfaces correspond to the ultrasound image.

These problems demonstrate that direct measurement of the outline is not a feasible method for obtaining the reference outline.

5.3.2 Cross-modality registration

Considerable research has been undertaken with the aim of registering images from a region of the body, acquired using different imaging modalities. Most of the work is aimed at matching computed tomography (CT) images with magnetic resonance (MR) images, with ultrasound / MR registration less common. The general principle involved in the registration techniques are to find a transformation matrix which matches corresponding features between the two images. The corresponding features are usually anatomical landmarks, such as protrusions on bones, ventricles in the brain *etc.* Once the transformation matrix is calculated, it may be applied to all the pixels in one image, warping that image to match the other.

The prostate is more clearly visible in images acquired using different modalities to ultrasound. Figure 5.3 shows an axial MR image with a clear distinction between the prostate and its surrounds. Figure 5.4 shows an axial CT image where the prostate is also more easily discernible than in transrectal ultrasound images. It should be noted, however, that the clear outline of the prostate in an MR or CT image may not necessarily be the same as the outline which is apparent on an ultrasound image. The three modalities (MR, CT and ultrasound) each



Figure 5.3: An example axial MR image. The prostate outline is clearly visible.

acquire images which vary according to different properties of the tissue (MR depends on the number of protons in the tissue, CT changes according to the attenuation of the x-rays, ultrasound varies depending on the reflections from the tissue).

The prostate is most clearly visible in CT images. If it is assumed that the prostate boundary is the same on the different modalities, then it is clear that the registration should be made between ultrasound and CT images. There are many practical reasons why obtaining CT images of the prostate would be difficult. Not least is the expense of CT in comparison to an ultrasound examination. However, even if the factors of cost, ethical approval *etc.* are not considered then there are fundamental reasons why comparing ultrasound scans to CT images is not feasible.

Prostate ultrasound images do not exhibit the necessary anatomical landmarks which would enable direct registration to the CT image. This could possibly be overcome by inserting fiducial markers into the prostate which would show up on both the CT and ultrasound images.. The deformation of the prostate caused by the transrectal probe during ultrasound imaging, could be replicated by using an ultrasound probe shaped phantom whilst acquiring the CT image information.




Figure 5.4: CT image of the prostate.

However, there is additional movement and deformation of the prostate which could not be compensated for. This movement results from the change in body position of the patient between the two imaging modalities. During CT imaging, the patient is positioned within the bore of the machine with legs out straight and together. Patients undergoing prostate ultrasound examinations lie on their side with legs drawn up. Studies involving patients undergoing conformal radiotherapy have shown movement of the prostate of 10 mm depending on the position of the legs [92].

5.3.3 Prostate phantoms

One possible solution to the problem of lack of knowledge of the correct outline is to use a phantom, since a phantom allows images to be acquired with full knowledge of the underlying shape of the phantom prostate. However, such a study will only have relevance if the phantom behaviour when probed, cut, deformed or ultrasound imaged is reasonably similar to that in-vivo. A commercial prostate ultrasound phantom is available, marketed as a training aid for surgeons to practice prostate brachytherapy techniques⁵. This phantom may be imaged using a standard transrectal ultrasound probe and has an elastic membrane through which

⁵Prostate brachytherapy phantom, Nuclear Associates ltd, NY, USA.




Figure 5.5: a) Brachytherapy phantom. b) Resulting ultrasound image.

needles may be inserted into the prostate. The phantom and an ultrasound image of the phantom are shown in figure 5.5.

The ultrasound image in figure 5.5 was acquired with the phantom in air, using a transrectal ultrasound probe smeared with ultrasound gel to give a good acoustic coupling between the probe and the phantom. The ring artefacts below the probe are due to the lack of contact between the probe and phantom, and the bright areas at the sides of the images are due to the phantom/air interface. These could be reduced or removed by imaging the phantom inside a water bath. However, the important elements of the image are the prostate and its immediate surrounds, and these are strikingly different from the prostate observed in clinical ultrasound images. The phantom prostate image is uniform in grey level and considerably brighter than the surrounding matrix. Therefore, images from this phantom are of little use for testing the outlining algorithms.

Given the lack of an appropriate commercial phantom, the author investigated the possibility of developing a phantom which provides clinically realistic ultrasound images. A search of the literature reveals a variety of materials which have been used in ultrasound phantoms to mimic the echogenic properties of tissue. There are two types of application; phantoms for testing ultrasound machine performance and specification, and phantoms used as teaching and training aids. The first type attempts to create phantoms where physical ultrasound properties (such as attenuation coefficients and sound velocity) match those of real tissue. The second category generally consists of phantoms for practicing surgi-

cal procedures which are conducted under ultrasonic guidance, such as biopsy. The important requirement for these phantoms is that the surgical targets (simulated cysts and masses) are clearly visible; producing clinically realistic images is secondary. However, some of the techniques produce images which provide reasonable approximations to the clinical images.

Cos [46] describes the construction of a transrectal prostate ultrasound phantom from a Foley balloon and catheter, and a collapsible wine container, but the presented results are far removed from the images obtained in in-vivo scanning even though the purpose of the phantom is to provide training in prostate ultrasound anatomy rather than surgical procedure. Biopsy phantoms generally consist of a gelatin or agar gel with added scatterers. Various materials have been used as scatterers such as flour [152], cornstarch [185], silicium-carbide powder [75], graphite powder [160], polyethylene granules [30] and psyllium hydrophilic mucilloid fibre [29].

A trial by the author, using each of these materials, showed that the most suitable material for the scattering agent is the psyllium fibre used by Bude and Adler [29], assessed qualitatively in terms of the visual appearance of the speckle on the ultrasound image. The particle size of flour and cornstarch is too small for realistic images, and the silicium-carbide and graphite powders make the mixture too dark to distinguish the different colours of the prostate and surrounding regions. Polyethylene granules provide good images, but it is difficult to prevent them from floating on the surface.

After much experimentation adjusting the quantities of psyllium fibre in the prostate and non-prostate regions, the best prostate phantom found produced images of improved realism over the commercially available phantom. Figure 5.6 shows the phantom suspended in a custom-built phantom box. The wires suspended across the box are visible on the ultrasound image and coincide with a slicing template which may be used post-imaging to slice the phantom at the same location as the image acquisition. A sample ultrasound image of the phantom and an optical image of the associated slice are given in figure 5.7. It is clear the ultrasound phantom produces images which are more realistic than those from the commercially available phantom. However, there are still significant differences between the phantom images and clinical images. For example, the speckle patterns observed in the phantom image are highly regular since the phantom material is uniform (within the prostate or non-prostate regions), and the boundaries between the regions are much clearer than those observed in clin-




Figure 5.6: Ultrasound phantom box and prostate.




Figure 5.7: a) Ultrasound image of the phantom. b) Optical slice at the same location with overlaid hand draw outline.

ical prostate ultrasound images. Therefore, use of the phantoms for testing the outlining algorithms was abandoned.

5.3.4 Ultrasound outlining by “experts”

From the discussion in sections 5.3.1 through 5.3.3, it would appear that it is not possible to obtain the true outline of the prostate by either direct measurement or matching with a different imaging modality, neither can a phantom be produced which produces realistic images. In the absence of a direct measurement, the best estimate of the outline can be taken to be that drawn by a human observer with considerable knowledge of prostate anatomy, surgery and ultrasound. Such an observer is termed “expert”.

Four experienced urologists were asked to perform outlining on a series of ultrasound images. Since hand outlining is time-consuming, a sample of twenty one images was randomly selected from the complete set of images. The selection was performed once, *i.e.* each urologist was shown the same set of images. Each image was displayed on a computer and a urologist marked salient points to identify his opinion of the outline. Parametric cubic spline interpolation was used to connect these points. Points could be added, moved or removed until the urologist was satisfied that the displayed outline was satisfactory. Each of the urologists worked independently, therefore they were not biased by the outlining opinions of their counterparts.

Each of the urologists have extensive experience and therefore may be considered to be experts in the art of prostate outlining. However, the prostate outlines as drawn by the experts are not always the same. In some images there is a close match between the boundary opinions of the experts. Others show significant discrepancies between the position of the outline. Figure 5.8 shows two ultrasound images with overlaid expert drawn outlines. The outlines in the ultrasound image on the left of figure 5.8 show close agreement between the experts, while those on the right show clear discrepancies. The choice of which outline to use as the *de facto* reference, depends on the intended use of the outline. For example, if the outline were used to decide the resection boundaries for an automated TURP system such as CASMS, it would be prudent to err on the side of safety and combine the innermost outline portions to create the reference.

For the benchmark in the assessment of the algorithms, the reference outline is taken to be the average of the outlines from the group of experts. The average




Figure 5.8: Transrectal ultrasound images with expert delineated prostate boundaries. a) Close agreement between the experts' outlines. b) Marked disagreements between the experts' opinions.

outline can be calculated from corresponding points on the outlines. Since there are no anatomical landmarks on the prostate outline which may be used for correspondence, some means of assigning artificial corresponding points is required.

The method for finding the corresponding points is based on work reported by Chalana and Kim [38]. From a set of N outlines, a random point P_1 on a randomly selected outline is chosen. This outline is divided into n^6 segments of equal length starting from point P_1 . These n segments have n endpoints which define the outline. For each of the $N - 1$ remaining outlines, the closest point on the outline to the point P_1 on the first outline is found, and denoted P_1 for that outline. These outlines are also divided into n equal segments, which define n endpoints describing each outline. Thus each outline is described by points $(P_1, P_2, \dots, P_n)$, and the points P_i of all the outlines are corresponding points.

An initial approximation of the average outline is found by taking the mean x and y values of each set of corresponding points. These mean corresponding points are used in a cubic spline interpolation to produce the initial average outline. Chalana and Kim [38] update the initial approximation in the following manner.

⁶ n is a user defined parameter.




Figure 5.9: Breakdown in the Chalana and Kim method of finding the average outline. The input curves are the red, green, blue and yellow lines, and the black curve is the average at each iteration. a) Initial average with two points (red and blue) and their normals highlighted. b) Migration of the blue point due to the shallow angle between the normal and the yellow curve. c) After three iterations the normal at the red point of the mean curve no longer intersects all the input outlines.

At each of the mean corresponding points, the normal to the average outline is calculated. The intersections between these normals and the original outlines are then taken to be a new set of corresponding points. The means of the new corresponding points give a better approximation of the average outline, and the process is repeated. The final average outline is found when none of the mean points changes by more than a threshold between iterations.

The problem with this method is that the outlines must be concave, otherwise it is possible that no intersection exists between the normal to the average outline and one (or more) input outlines. An example where the method breaks down is shown in figure 5.9. The black curve in figure 5.9 is the mean of the four input curves (red, blue green and yellow). As the algorithm progresses, the intersection between the normal to the mean curve at the blue asterisk and the yellow curve causes the migration of the blue asterisk. In the third iteration, this movement results in a situation where the normal to the mean curve at the red asterisk does not intersect the yellow curve.

This breakdown in the Chalana and Kim method has been solved by taking the new corresponding points as the point on each outline which is closest to the point on the average outline, repeated for every point on the average outline. A




Figure 5.10: Finding the average of two outlines. a) The input (green and blue) curves and initial (black) mean outline. The initial correspondence is highlighted by the red markers. b) The final mean outline and corresponding points after algorithm convergence. Note the movement of the mean curve on the right hand side of a) and b).

better approximation to the average outline is then found by taking the mean of the new corresponding points. The process is repeated until the average outline changes by less than threshold. It is important to assess changes over the whole outline between iterations since the mean corresponding points may change while the outline remains the same – *i.e.* the points cycle around the outline without the shape changing. This new method is robust for any shape of outline and any degree of overlap between the two outlines. The average outline of two outlines is shown in figure 5.10, where a) shows the initial average, and b) the final average outline.

5.4 Outline comparison

In the previous section various methods of determining the ideal reference outline for a given images were discussed. Given a reference outline and a candidate outline result from one of the outlining algorithms, the question which needs to be answered is how well do the two outlines match? This section discusses various possible measures which could be compared to answer this question.

5.4.1 Derived parameters

Knoll et al [109], Pathak et al [162] and Aarnink et al [2] calculate the area contained within an automatically determined prostate outline and assess performance by comparing the area with the area contained within an expert's manually drawn outline. In each case, the purpose of the algorithm is to determine the volume of the prostate (by summing the areas of each scan and multiplying by the inter-scan distance), since prostate volume is an important clinical factor when deciding treatment methodologies.

In other areas of medical imaging it is common to derive different clinical parameters from the outlining results. Chalana et al [37] derive head circumference and abdominal diameter from outlining ultrasonic images of the foetus. Geiser et al [76] use the perimeter of heart outlines as a basis for evaluating their outlining algorithm.

The choice of which parameter to use is normally based on the clinical relevance of a particular measure. However, the problem with using a derived clinical parameter is that very different outlines may produce the same result. In applications where the desired goal is simply to calculate the clinical parameter, then comparing the parameter calculated from an automatically determined outline to the parameter calculated for the reference outline is valid. In terms of the current investigation for surgical monitoring systems, it is the shape and position differences between the candidate and reference outlines which is important. Therefore, comparison using only a derived parameter is not justified.

5.4.2 Error metrics

Chalana and Kim [38] present three methods for determining a metric suitable for assessing the similarity of two outlines. The first two methods calculate the metric using a set of corresponding points between the curves. The method for obtaining corresponding points is discussed in section 5.3.4. Given outlines A , B and C , an error metric e must exhibit the following three properties:

$$\begin{aligned}
 I) \quad & e(A, A) = 0 \text{ and } e(A, B) \geq 0 \\
 II) \quad & e(A, B) = e(B, A) \\
 III) \quad & e(A, C) \leq e(A, B) + e(B, C)
 \end{aligned} \tag{5.1}$$




Figure 5.11: Hausdorff distance (thick black line) for two pairs of outlines

5.4.2.1 Hausdorff distance

Consider two outlines A and B , each described by a set of m corresponding points such that $A = \{a_1, a_2, \dots, a_m\}$ and $B = \{b_1, b_2, \dots, b_m\}$, where a and b are pairs of (x, y) coordinates. The *distance to closest point* (DCP) for a point a_i to the curve B is described as:

$$DCP(a_i, B) = \min_j \|b_j - a_i\| \quad (5.2)$$

where $\|\cdot\|$ is the Euclidean distance between two points.

The *Hausdorff distance* [104] is defined as the maximum value of the DCP for all points on both contours, i.e.:

$$e(A, B) = \max \left(\max_i \{DCP(a_i, B)\}, \max_j \{DCP(b_j, A)\} \right) \quad (5.3)$$

The Hausdorff distance obviously obeys the first two requirements of a metric, and Chalana and Kim [38] show that the third property is also valid. However, the Hausdorff distance is a measure of the largest discrepancy between two outlines, therefore it does not provide a satisfactory measure for assessing the similarity of two outlines. Figure 5.11 shows two pairs of outlines which have the same Hausdorff distance, but the outlines are very different.

5.4.2.2 Distance between outlines

Chalana and Kim [38] describe a second metric for analyzing outlines based on the root mean squared distance (RMS) between the contours, however, the method

uses corresponding points which are found in a radial manner. First, a common centroid is specified for both contours. The radial corresponding points are defined by extending rays from the centroid at equal angular increments - the intersection of the rays with the outlines give the points. The metric is defined as the RMS distance between the corresponding points, and has been used to analyze ventricle outlines in ultrasound images [64, 74].

The problem with this metric is that it is not applicable in general, since it would fail on outlines where some parts of the outline are occluded by other parts, when viewed from the centroid - *i.e.* where outlines have concave regions with a high degree of curvature. In the set of expert drawn prostate outlines, most shapes were suitable for this method, but a few failed. Therefore, the RMS distance is not suitable for use as a metric in the current study.

The average distance between curves - the mean of the distance between corresponding points - is a useful measure for comparing outlines. For example, the curves in figure 5.11a) have a lower mean distance than the curves in figure 5.11b), despite having the same Hausdorff distance measure. However, average distance is not strictly an error metric, since it does not obey the triangle property (equation 5.1 III), but will be used to compare outlines in the statistical analysis (section 5.5.2).

5.4.2.3 Area between outlines

The third metric is a measure of the area between the outlines. Each outline is converted to a binary image where white pixels constitute the outline and the area enclosed by the outline, and black pixels form the remainder. By performing an XOR operation on the two binary images, the result has white pixels corresponding to the area between the outlines. Chalana and Kim [38] state "the average number of non-zero pixels (white) in the resulting image defines a metric for the two boundaries". It is unclear what the authors mean by average since the XOR operation produces a single image, and the sum of white pixels would not be a valid metric. However a valid metric may be obtained by normalizing by the area of one of the contours (dividing the sum of white pixels by the sum of white pixels in one of the outline binary images).

Prater and Richard [166] use percentage of correctly classified pixels (%CCP) as a parameter for determining the success of their prostate segmentation algorithms. The algorithm produces a binary image of prostate and non-prostate pixels, and




Figure 5.12: Area between corresponding points

is compared against a manually derived image. This may be converted to a valid metric by using the percentage of incorrectly classified pixels (%ICP) i.e.

$$\begin{aligned} \%ICP &= 1 - \%CCP \\ &= 1 - \frac{\sum prostate^+ + \sum nonprostate^+}{\sum pixels} \end{aligned} \quad (5.4)$$

where $\sum prostate^+$ is all the pixels correctly identified as part of the prostate, $\sum nonprostate^+$ is the pixels correctly identified as not part of the prostate, and $\sum pixels$ is the total number of pixels contained within the image.

The drawback of these methods is that they are computationally expensive. An alternative is to approximate the area between the outlines using the corresponding points [25]. Two possibilities exist for overlapping areas between adjacent pairs of corresponding points, either a quadrilateral or two triangles, as shown in figure 5.12.

Consider four points P_1, P_2, P_3, P_4 where P_1, P_2 are adjacent points on one outline, and P_3, P_4 are the corresponding points on the second outline, and each point P_i has a pair of Cartesian coordinates (x_i, y_i) . The area of the quadrilateral formed by the four points is given by:

$$A_{quad} = \frac{1}{2} (x_1y_1 + x_2y_3 + x_3y_4 + x_4y_1 - x_1y_4 - x_2y_1 - x_3y_2 - x_4y_3) \quad (5.5)$$

For the triangles case, the intersection of the outlines must be approximated by finding the intersection of the straight line segments joining P_1 to P_3 and P_2 to P_4 . If the intersection is denoted P_5 , then the area of the two triangles is calculated as:

$$A_{\Delta} = \frac{1}{2} (\begin{aligned} &x_1y_5 + x_4y_1 - x_1y_4 - x_4y_5 - x_5y_1 + \\ &x_2y_5 + x_3y_2 + x_5y_3 - x_2y_3 - x_3y_5 - x_5y_2 \end{aligned}) \quad (5.6)$$




Figure 5.13: Outline area approximation using the trapezium rule

The total intersection area is normalized by dividing by the area of one of the outlines, to give a valid metric. The outline area may be approximated by treating the outline as a polygon constructed of straight line segments joining the corresponding points. The area is found by subtracting the area below the lower half of the polygon from the area under the upper half of the polygon, as shown in figure 5.13. These areas are found using the trapezium rule.

5.4.2.4 Error metric selection

The Hausdorff distance assesses the worst discrepancy between two outlines, therefore it can be used to measure the maximum error between the outline result from an algorithm and the reference outline, giving a valuable measure of algorithm performance. However, use of the Hausdorff distance alone is flawed as demonstrated by the outlines in figure 5.11, since different outlines may produce the same result. Therefore, an additional assessment is required using either the average distance between the outlines, or the area between the outlines. For the application of assessing prostate outlining algorithms for use in the CASMS system, the average distance provides a more suitable measure, because it is not affected by the overall size of the prostate. This is important because the final use of the outlines – to prevent the resectoscope from puncturing the capsule – is governed by the absolute distance to the outline, not a relative measure in terms of prostate size.

5.5 Statistical analysis

The objective of assessing the different outlining algorithms is to find the algorithm which performs best over the set of images. Therefore the question posed is: Can an expert observer be replaced with an automated outlining algorithm? The following sections discuss various statistical techniques which have been previously used in an attempt to answer such a question. Section 5.5.1 describes a repeated measurements study which compares each algorithm's measurement of a clinical parameter with the measurement made by the experts. Although, a clinical parameter based assessment is not ideal, the study is included since it takes into account the possibility of variation in one expert's opinion when a prostate image is outlined multiple times by the same expert. Two additional statistical analyses are presented which determine algorithm performance either by comparing error metrics calculated between the algorithm outlines and expert outlines (joint agreement study in section 5.5.2), or by assessing the ability of each algorithm to produce outlines which lie within the range of the experts' results.

5.5.1 Repeated measurements study

The first statistical test is a repeated measurements study first described by Eliasziw et al [69], and used by Tong et al [212] to assess the prostate volume determination performance of expert observers. In their study a group of o observers make m measurements (on their outlines), on n subjects (prostate ultrasound images). In the current study, the observers are the experts (urologists) and the automated outlining algorithms. Prostate area is chosen, rather than volume, since only a single image from each prostate is used.

The study results are formulated into a table for each observer as shown in table 5.1, with the result X_{ijk} indicating the j^{th} observation of the i^{th} subject by the k^{th} observer, where $i = 1, \dots, n$, $j = 1, \dots, m$ and $k = 1, \dots, o$.

The composition of an observation is broken down into the true value of the observation with certain errors added:

$$X_{ijk} = \mu + p_i + o_j + po_{ij} + e_{ijk} \quad (5.7)$$

where μ is the overall mean value of the prostate area, p_i is the effect of prostate i on the measurement, o_j is the effect of observer j on the measurement, po_{ij} is

subject	trial					sum	mean
	1	...	j	...	m		
1	X_{11k}	...	X_{1jk}	...	X_{1mk}	$X_{1..k}$	$\bar{X}_{1..k}$
...	...	...	...	...	...	...	...
i	X_{i1k}	...	X_{ijk}	...	X_{imk}	$X_{i..k}$	$\bar{X}_{i..k}$
...	...	...	...	...	...	...	...
n	X_{n1k}	...	$X_{nj k}$	...	X_{nmk}	$X_{n..k}$	$\bar{X}_{n..k}$
sum	$X_{.1k}$	...	$X_{.jk}$	...	$X_{.mk}$	$X_{..k}$	
mean	$\bar{X}_{.1k}$	...	$\bar{X}_{.jk}$	...	$\bar{X}_{.mk}$		$\bar{X}_{..k}$

Table 5.1: Observations of observer on a set of subjects, conducted n times.

effect of the interaction of observer j and prostate i , and e_{ijk} is the intra-observer random error in X_{ijk} .

The study allows the inter-rater and intra-rater reliabilities to be determined concurrently. Eliasziw et al provide a lucid example detailing the procedures for estimating the quantities in equation 5.7 from a completed table of values such as table 5.1. However, there are a number of difficulties which preclude the use of a repeated measurement study as an initial assessment of the outlining algorithms.

A subtle problem is that choosing the correct sample size is not straightforward. An appropriate number of observers and subjects must be used for a valid study, without over-burdening the expert observers and spending too long acquiring the data set. However the appropriate numbers which can be used to estimate reliability are dependent on the reliability. This problem could be overcome by performing an initial pilot test to assess the likely reliability and choose numbers for the full study based on these.

The use of a derived parameter such as prostate area is undesirable for the reasons discussed in section 5.4.1. However, the use of an error metric breaks the statistical basis of the study since the overall mean of the error metric tends toward zero. The statistical model is also violated if the automated outlining algorithm is completely deterministic, i.e. the algorithm always gives the same result when repeatedly applied to the same input image, since the terms p_i and e_{ijk} will be zero. This is the case for the edge base detection algorithms, therefore a repeated measurements study is only applicable to the feature modelling and genetic algorithms approach by Wu [224, 225].

5.5.2 Joint agreement study

Chalana and Kim [38] describe a different test known as a *joint agreement* study. The study aims to determine whether an individual observer agrees with the set of observers less than any observer agrees with the set, and was first presented by Williams [222]. Consider $n + 1$ observers, numbered from 0 to n , making measurements on N subjects. The average level of agreement between observer 0 and the set of observers $1, \dots, n$ is given by:

$$P_0 = \frac{1}{n} \sum_{j=1}^n P_{0,j} \quad (5.8)$$

and the average level of agreement between the set of observers is:

$$P_n = \frac{2}{n(n-1)} \sum_j \sum_{j': j' \neq j} P_{j,j'} \quad (5.9)$$

where $P_{jj'}$ is the proportion of agreements for a pair of observers j and j' ($j, j' \in [1, \dots, n]$).

For the nominal data in Williams' analysis, the ratio $\frac{P_0}{P_n}$ is an index which compares observer 0 agrees with the remaining group of n observers. Chalana and Kim define a modified Williams index for the numeric multivariate data used in the current analysis:

$$I' = \frac{\frac{1}{n} \sum_{j=1}^n P_{0,j}}{\frac{2}{n(n-1)} \sum_i \sum_{j': j' \neq j} P_{j,j'}} \quad (5.10)$$

The problem is reduced to how to calculate the agreement between two observers' outlines. Chalana and Kim [38] define the average *disagreement* D of two observers as the average of the chosen error metric calculated for the N subjects, i.e.:

$$D_{j,j'} = \frac{1}{N} \sum e(X_{ij}, X_{ij'}) \quad (5.11)$$

where X_{ij} is the i^{th} observation of observer j .

Chalana and Kim next define the agreement of the two observers as the reciprocal of the average disagreement:

$$P_{j,j'} = \frac{1}{D_{j,j'}} \quad (5.12)$$

However, the asymptotic nature of this function creates a certain bias in the value

of the agreement. Therefore, the agreement has been redefined as:

$$P_{j,j'} = \exp\left(\frac{-D_{j,j'}}{K}\right) \quad (5.13)$$

where K is a positive real number.

The confidence interval for the index is estimated using a jackknife non-parametric sampling technique as described by Bland and Attman [19]. The results for one of the subjects (images) is omitted and the modified Williams index is calculated for the remaining observers. The calculation is repeated by leaving out the results for each subject in turn. If the modified Williams index calculated by leaving out the results for the i^{th} subject is denoted by $I'_{(i)}$, then the standard error in the modified Williams index may be estimated as:

$$se = \left\{ \frac{1}{N-1} \sum_{i=1}^N [I'_{(i)} - \bar{I}]^2 \right\}^{\frac{1}{2}} \quad (5.14)$$

where the mean modified Williams Index is given by:

$$\bar{I} = \frac{1}{N} \sum_{i=1}^N I'_{(i)} \quad (5.15)$$

The 95% confidence interval may then be calculated as:

$$CI_{95} = \bar{I} \pm z_{0.95} se \quad (5.16)$$

where $z_{0.95} = 1.96$ is the 95th percentile of the standard normal distribution.

5.5.3 Percentage statistic

Chalana and Kim [38] propose a second technique which tests the number of successful results produced by a given algorithm. Success is judged according to whether or not the result of the algorithm lies within the range of results produced by the observers. For derived parameters such as area, the test simply assesses whether the result from the outlining algorithm is greater than the minimum value and less than the maximum value of the experts' results.

For the error metrics, the success may be approximately judged by comparing the maximum algorithm to observer metric to the maximum observer to observer metric. For a given metric e , the success for one result of the outlining algorithm

depends on whether:

$$\max_i \{e(A, X_i)\} \leq \max_{i,j} \{e(X_i, X_j)\} \quad (5.17)$$

where A is the algorithm result for one image, X_i are the results from the experts for the same image.

The assumed hypothesis is that the observers (algorithm and experts) produce boundaries which are samples from the same distribution. Therefore, for $n + 1$ observers, the probability that one observer's result lies outside range of the other observers' boundaries is $1/(n + 1)$. Similarly, the probability that one observer's result is inside the range is given by $n/(n + 1)$. Therefore, using the outline result for the four human expert observers. the expected percentage is 80%.

5.6 Results

The successful algorithms identified in the previous chapter were those based on work by Aarnink [1] (section 4.3.3), Pathak [162] (section 4.3.4), Liu [136] (section 4.3.5) and Zhao [234] (section 4.3.2) from the edge based detection schemes, and the feature modelling approach adopted by Wu [225] (section 4.4.3). Each of these algorithms was applied to the 62 images in the dataset. For the edge based methods, several parameters are required such as the high and low levels for the hysteresis thresholding (section 4.3.3), the prostate centre, the radius of the circle from the prostate centre within which no edges are permitted and the minimum allowed segment length.

For the prostate centre parameter, it is not appropriate to ask the user to pick the centre point as this violates the automated algorithm criterion, The centre point must be determined in some automated fashion. However, this prostate centre point has an important role in all the edge based methods, since it determines which of the edge segments (dark to bright transitions and at a set distance from the centre) are kept in the construction of the closed outline. It is also the centre point of the radial-bas relief process used in Liu's algorithm. If the centre point lies too close to or outside the boundary then all the edge based techniques will fail. The adopted approach to find the centre has been to calculate the position⁷ of the centroid of the expert delineated outlines relative to the probe

⁷The physical position is measured rather than the pixel position, *i.e.* the image zoom factor is taken into account.




Figure 5.14: Results for the ultrasound image with close expert agreement. a) Reference outline. b) Edge based detection results (Aarnink – red, Pathak – green, Liu – blue, Zhao – yellow). c) Best (green) and worst (red) feature modelling result.

position for the set of prostate ultrasound images. The mean relative position was calculated. The prostate centre is then taken as the mean physical position relative to probe for each of the images. Several attempts were made to optimize the other parameters to give the best results judged qualitatively over the whole image set.

The feature modelling approach of Wu does not require any parameters once the feature model is formed. The construction of the feature model is discussed in detail in the next chapter, section 6.2. In contrast to the edge based detection schemes, Wu's algorithm produces different results when applied repeatedly to the same image due to the random nature of the genetic algorithm optimization. Therefore, Wu's algorithm was applied ten times to each image. For a single image, the ten results were compared to the reference outline⁸ and ranked according to the average distance between the corresponding points. Both the sets of best results and worst results are included in the analysis.

The reference outlines, edge based detection algorithms results, and feature modelling results for the prostate ultrasound images of figure 5.8 are shown in figures 5.14 and 5.15.

⁸The reference outline is either the average of the experts' outlines in the case of an image from the multiple experts set, or the single expert result for the remaining images.




Figure 5.15: Results for the ultrasound image with poor expert agreement. a) Reference outline. b) Edge based detection results (Aarnink – red, Pathak – green, Liu – blue, Zhao – yellow). c) Best (green) and worst (red) feature modelling result.

5.6.1 Multiple experts results

Table 5.2 shows the Hausdorff distances found by comparing each of the algorithms to the four expert outlines and by comparing each of the expert outlines to the other experts' results, averaged over the 21 images in the multiple expert data set. The calculated values for the Williams index and percentage statistic are also given, along with their respective 95% confidence intervals (CI).

The results demonstrate the importance of comparing the algorithm outlines to the opinion of more than one expert. If the values for the computer to observer difference are taken in isolation, then even the lowest value of 4.63 mm for the set of best Wu results can be construed as poor. However, when the variation in the experts results is also considered, the algorithms' results are far more promising. The slight variation in the inter-observer agreement is due to the random starting point for the corresponding points algorithm.

The results in table 5.2 show that the mean computer observer differences for the four edge based detection schemes are very similar, as are the corresponding Williams indices. This is expected since the four algorithms themselves only differ in their methods of finding the initial edge segments which are likely to be part of the prostate. Given that one of the algorithms finds a strong, suitable edge segment at a particular location in an image, it is likely that the other edge based algorithms will also find an edge segment at the same location, regardless

	COD[mm] / (σ)	IOD[mm] / (σ)	WI / [CI]	P(%) / [CI]
Aarnink	5.95 (0.14)	3.75 (0.19)	0.80 [0.82, 0.78]	23.8 [22.8, 24.8]
Pathak	6.48 (0.34)	3.75 (0.16)	0.76 [0.74, 0.78]	19.0 [18.3, 19.8]
Liu	5.97 (0.41)	3.72 (0.12)	0.80 [0.78, 0.82]	14.3 [13.8, 14.8]
Zhao	5.99 (0.20)	3.74 (0.15)	0.80 [0.79, 0.81]	23.8 [24.8, 22.8]
Wu _{worst}	13.97 (0.31)	3.73 (0.13)	0.36 [0.35, 0.37]	0 [0, 0]
Wu _{best}	4.63 (0.38)	3.75 (0.19)	0.92 [0.90, 0.93]	52.4 [49.8, 55.0]

Table 5.2: Comparing the algorithms' outlines to the experts' outlines using the Hausdorff distance as an error metric. COD is the mean computer to observer difference, IOD is the mean inter-observer difference, σ is the standard deviation, WI is the Williams index, P is the percentage statistic and CI are the respective confidence intervals.

of whether the true prostate edge lies at that location.

The set of best results using Wu's algorithms show the closest agreement to the expert results. This is reflected in their Williams index which is closer to the optimal value than the index for any other algorithm. Unity for the Williams index implies that no distinction can be made between the outlining algorithm's result, and the set of observers' results, i.e. the algorithm can replace a human observer without any change in the statistical comparisons. However, the percentage statistic is more exacting since it tests the frequency with which the algorithm result lies within the range of the experts' results, rather than how well the algorithm results agrees with the set of experts results. Although this seems to be a contradiction, the hypothetical outlines shown in figure 5.16 demonstrate that the Williams index and percentage statistic do measure different aspects of agreement.

Figure 5.16 shows outlines from three experts (black and blue circles) and one algorithm (red circle). The Hausdorff distances between each of the experts' results and the algorithm result are also shown. Clearly the algorithm's result agrees with the set of outlines as much as the blue expert result agrees with the set, therefore it is consistent with the hypothesis that the algorithm can replace




Figure 5.16: Hypothetical outlines demonstrating the stringency of the percentage statistic.

a member of the group. However, the percentage statistic is not satisfied; the green line representing the largest Hausdorff distance for the algorithm outline does not lie in the range of the other experts.

The most striking result from table 5.2 is the how badly the set of worst results from Wu's algorithm compare to the experts results – the Williams index is much lower than for any other algorithm's results, and the percentage statistic is zero, meaning that the worst results from Wu's algorithm are never within the range of the experts' outlines. Therefore, although capable of producing very good results, the feature modelling approach is too variable to be used reliably as an automated outlining technique.

Table 5.3 shows a similar set of statistics to table 5.2 where the average distance is used as the performance measure rather than the Hausdorff distance. Similar trends are observable in the presented results, with the best set of feature modelling outlines lying closer to the experts opinions than the result of any other algorithm. The Williams index for the best feature modelling set is almost unity, therefore fulfilling the hypothesis that the algorithm is a reliable member of the expert group. However, the percent statistic is still some way short of the expected value of 80%.

Once again, the worst set of feature modelling results, falls far below the standard of the other algorithms' results.

	COD[mm] / (σ)	IOD[mm] / (σ)	WI / [CI]	P(%) / [CI]
Aarnink	2.21 (0.10)	1.22 (0.13)	0.90 [0.89, 0.92]	33.3 [31.8, 34.9]
Pathak	2.53 (0.19)	1.23 (0.14)	0.88 [0.87, 0.89]	28.6 [27.3, 29.9]
Liu	2.20 (0.20)	1.22 (0.13)	0.91 [0.89, 0.92]	28.6 [27.3, 29.9]
Zhao	2.27 (0.12)	1.23 (0.14)	0.90 [0.89, 0.91]	14.3 [13.8, 14.8]
Wu _{worst}	9.14 (0.15)	1.22 (0.13)	0.45 [0.44, 0.46]	0 [0, 0]
Wu _{best}	1.31 (0.17)	1.24 (0.15)	0.99 [0.99, 1.00]	47.6 [45.2, 50.0]

Table 5.3: Comparing the algorithms' outlines to the experts' outlines using the average distance measure. COD is the mean computer to observer difference in mm, IOD is the mean inter-observer difference in mm, σ is the standard deviation, WI is the Williams' index, P is the percentage statistic and CI are the respective confidence intervals.

5.6.2 Reference outline comparison

Figure 5.17 shows the average boundary distance for the 62 results of each outlining algorithm from the 62 reference outlines. The red line in each box marks the median value for the data, the box limits are placed at the lower and upper quartile boundaries, and the whiskers extend across the data (1.5*inter-quartile range). Data outliers are marked with a red plus.

The plot shows that the edge based outline detection algorithms perform similarly to each other, with the modified RBR method of Liu showing slightly better results than the other algorithms. The set of best results from the feature modelling algorithm by Wu shows a significant improvement, with the lowest median result, the smallest interquartile range, and the minimum outlier values. However, the average distance data for the worst feature modelling set shows an order of magnitude degradation over any of the other algorithms.

The results show that whilst Wu's algorithm has the potential to perform very well, it is not well enough confined and can also produce very poor results. Therefore, the statistical analysis shows that the RBR method is the best overall performing algorithm from the literature when factors such as the repeatability and robustness are taken into account. However, the potential of the feature mod-




Figure 5.17: Box plots of the average boundary distances from the reference outlines.

elling algorithm to provide the best results warrants further investigation. The analysis of various aspects of the feature modelling algorithm forms the basis of the next chapter.

5.7 Conclusion

This chapter provided a rigorous assessment of the promising algorithms from the initial outlining trials of chapter 4. Acquisition of an image data set which covers the range of image quality found in routine clinical ultrasound examinations has been described. Various possible techniques for assigning a correct outlining result to each of these images were investigated, with repeated outlining by human expert being chosen as the most suitable method.

A modified method for assigning artificial correspondency between outlines has been presented, which overcomes the fallibility of the previous technique. The use of these corresponding points to calculate comparison metrics between outlines has been demonstrated. Several statistical methods were discussed which use the error metrics to analyze the performance of each algorithm over the whole image set, in comparison to the experts' evaluations.

The results of the statistical analysis have shown that the feature modelling approach is capable of performing significantly better than any of the other algorithms. However, the feature modelling technique is not completely deterministic, and has the potential to perform significantly worse than the other methods. This leads on to the next chapter which analyses the feature modelling technique and the associated pitfalls which degrade performance, and investigates methods for improving the result.

Chapter 6

Optimizing the prostate outlining

6.1 Introduction

The results from the previous chapter identified statistical shape modelling combined with genetic algorithms as the best approach from the literature for finding the prostate outline, such as the technique proposed by Wu [225]. However, the technique does not achieve a correct result when applied to every image in the test dataset. The rest of the image processing work presented in this thesis aims to improve upon these results. There are three areas of the genetic algorithms technique which warrant further investigation to provide improved results.

The genetic algorithm is a global optimization technique whose aim is to maintain a population of different boundaries and create new boundaries derived from promising members of the population. The final output of the genetic algorithm is a boundary which corresponds to the best fit with the underlying image. There are many different genetic algorithm schemes available where factors such as the population size, encoding scheme, mutation and crossover rates *etc.* may be varied, which will affect the outcome of the algorithm and the speed of execution. The work of Wu [224] explored the effects of varying the genetic algorithm parameters, but the results were assessed based upon the maximum value of the objective function for a given image and the repeatability of the objective function score when the same image is processed multiple times. A better assessment would score the schemes according to the quality of the result in comparison to the experts' results in the same manner as the statistical analysis described in chapter 5. Other techniques besides genetic algorithms are possible candidates for performing the global optimization. These include simulated annealing methods

(described in section 6.3.1 for automatically aligning shapes) and various simplex methods. However, the method of obtaining the global optimization is secondary to the problems of finding the best objective function and ensuring the validity of candidate outlines.

The objective function used by Wu uses some measure based upon the first derivative of the image to decide what is a good result and what is not. In essence it is very similar to the edge-based detection schemes discussed in chapter 4, section 4.3.1 – using the strength of edges within the image to determine the position of the final outline. However, the strongest edges present in a given image do not necessarily correspond with the prostate outline which is characterized by its lack of a distinct edge. Therefore, measures other than the first derivative should be investigated to find an objective function which scores highly for the correct outline in the majority of the images. The work described in the second half of this chapter, starting on page 211 describes a statistical analysis whose aim is to determine the best objective function.

The work described in the first part of this chapter investigates the problem of modelling the shape variation of the prostate. Section 6.2 discusses the model used by Wu’s algorithms and highlights the deficiencies of the method. Section 6.3 describes a method of automatically aligning shapes using affine transformations as a pre-cursor to the point distribution modelling described by Cootes et al [45]. The background to the point distribution model is given in section 6.4, and several models are presented using different strategies for assigning correspondency between the aligned shapes. Section 6.5 presents the work undertaken to optimize the correspondency of the shapes, and gives example results using the optimal statistical shape model.

6.2 Feature modelling

Wu uses an adaptation of the point distribution model of Cootes and Taylor, called *feature modelling*. The essential difference between the two modelling processes is that Wu omits the alignment stage in the model formation. Wu argues that by maintaining the physical dimensions of the model training shapes, the model is able to provide a better representation when used for “featureless” shapes such as the prostate.




Figure 6.1: Corresponding points of two outlines where mismatch occurs due to lateral displacement.




Figure 6.2: Expert outline which cannot be represented by polar points.

6.2.1 Feature model formation

The initial stage of the feature modelling process is an automated labelling of corresponding points on each prostate outline. First, the positions of the ultrasound probe centre and centroid of the outline are found. The corresponding points are then found at equal angular increments around the outline, with the initial point situated at 90 degrees (to the right) of the line joining the probe centre to outline centroid. Figure 6.1 shows the labelling for two prostate outlines.

The automated labelling introduces the first problem of feature modelling. Since the corresponding points are taken at angular intervals, only outlines which are continuous in polar space may be represented (where the origin of the polar coordinate system is the outline centroid). Figure 6.2 shows an example outline from the set of expert drawn outlines which cannot be represented in polar coordinates.

A second, less obvious problem with the automated labelling scheme is that it is based on the assumption that the prostate lies directly above the rectum, and that any rotation of the outline centroid around the probe centre, is caused by rotation of the ultrasound probe within the prostate. This is not necessarily true; the prostate can move around inside the body therefore a lateral displacement of the prostate with respect to the probe centre is possible. Figure 6.1 shows the correspondency of two outlines which are identical in shape but have different positions with respect to the probe centre. Ideally, the corresponding points of the two outlines would be identical, but this does not occur in practice.

The radius of the black probe circle is used to calculate the zoom setting of the ultrasound machine during the image acquisition. This allows conversion of the pixel position of the labelled points to positions in physical dimensions (mm). However, this means that the algorithm is tied to one particular ultrasound machine model rather than allowing images of prostates captured using different scanners to be analyzed using the same feature model.

Once all the outlines have been labelled, the feature modelling process follows the same procedure as the point distribution modelling, which is described in full in section 6.4. A vector representing each outline is constructed from the distance between the probe centre and outline centroid, and the radii of each of the landmark points. The vectors for each outline in the training set define a region in n^1 dimensional space which is the “allowable shape domain” – *i.e.* all valid prostate outlines should lie within this domain. The allowable shape domain is assumed to be ellipsoidal in n dimensions – *i.e.* the distribution along each axis is approximately Gaussian. The centre of the allowable shape domain is calculated from the mean of the feature vectors in the training set. Following construction of the covariance matrix of all the vectors, a principle components analysis provides the eigenvectors and eigenvalues which may be used to approximate the allowable shape domain.

6.2.2 Feature modelling discussion

Figure 6.3 shows the labelled point distribution of the expert outlines using feature modelling, and the mean outline (the centre of the ellipsoidal space) aligned so that the distance from the probe centre is the same in each case. The most obvious feature from figure 6.3 is that there is significant variation in the radius

¹The n dimensions are $n-1$ outline radii and the distance from the probe centre.




Figure 6.3: Corresponding points and mean outline for the feature modelling prostate model.

of the points for any given angular position. Figure 6.4 shows the variation in shape achievable by independently altering the weights applied to each of the first four principal components.

At first sight the results seem promising – the shape variations contained in figure 6.4 are nearly all valid prostate shapes, with the exception of the second mode extremes where the lateral lobes exhibit uncharacteristic concavities, and the minimum weighted first component (figure 6.4a) black curve) which has a small sharp corner at the base of the prostate.

The shapes in figure 6.4 are created by adjusting only one weight at a time with all the other weights fixed at zero. Figure 6.5 shows the effect on the prostate shape of adjusting more than one weight at a time. The curves in figure 6.5a) are formed by adjusting the odd numbered mode weights from the minimum value to the maximum value, while conversely adjusting the even numbered weights from the maximum to the minimum value. All the weights are adjusted from minimum to maximum value to produce the curves in figure 6.5b). These two sets of adjusted weights show some of the extremes of the permissible space occupied by the model. It is clear that these curves are not legitimate examples of prostate shapes – the outlines exhibit sharp corners and even double back on themselves, despite being “valid” in terms of the model definition.




Figure 6.4: Effect of varying the weights on the feature modelling prostate model. The black curves are obtained using the minimum value for the weights; the yellow curves use the maximum. The variance of each component lessens from a) through d).




Figure 6.5: "Un-prostate-like", valid shapes from the feature modelling prostate model. a) The ten weights are varied so that the odd numbered weights (1,3,5,...) vary from the minimum to maximum value and the even numbered weights (2,4,6,...) vary from maximum to minimum value. b) All the weights vary from minimum to maximum value.

6.3 Automated shape alignment

The problem highlighted in the previous section is the "looseness" of the model formed using the feature modelling approach. The initial hypothesis is that the large variation in size between the outlines in the training set results in an allowable boundary region which is too large, and contains many regions where candidate outlines are invalid. The first approach adopted to overcome this problem has been to revert to the standard point distribution model approach of Cootes and Taylor. This approach has been used previously by Arambula [5] to model the shape of the prostate in transrectal ultrasound images.

The first stage of the point distribution modelling is to align the training shapes before labelling. Arambula [5] used a manual approach to match outlines from transrectal prostate ultrasound images. Initially, one shape in the set was chosen as a typical prostate shape; subsequently each of the remaining outlines was transformed to fit the example shape through a series of translations, rotations and scalings. The transformation parameters were adjusted manually until the fit was deemed suitable enough by visual inspection.

Manual alignment of the training set is time-consuming and tedious, and is also

highly dependent on the choice of the typical prostate shape and the user's judgement of optimal fitness. Therefore a method of automated shape alignment has been developed.

The initial stage of the automated alignment randomly selects one outline from the set as the base against which the other outlines are matched. The adapted corresponding points technique presented in section 5.3.4 is used to assign an initial correspondence to each outline and the base outline. For each outline, simulated annealing optimization is used to adjust the pose parameters (translation, rotation and scaling) of the candidate outline so as to minimise the sum of squares distance between the corresponding points.

6.3.1 Simulated annealing

Simulated annealing [108] is a global optimization technique which performs particularly well in problems where the global minimum is hidden among many local sub-optimal minima, and is therefore ideally suited to the problem of prostate shape matching. The technique takes its name from the analogy with the thermodynamic process of formation of crystals in freezing liquids, or the annealing of slowly cooled metals. In a liquid at high temperature, the molecules have high kinetic energy and are able to move around freely. As the temperature drops, the molecules lose energy and are able to move less. If the cooling process is sufficiently slow, then the molecules can form a structured, ordered crystalline lattice, which is uniform throughout the medium. This structured state is the lowest possible energy form of the crystal.

There are four components to the simulated annealing technique. The first component is a description of the possible configurations of the system; for the current problem, the possible configurations are controlled by adjusting the pose parameters of the outline transformation. Five parameters are employed in this case: translation in the x and y dimensions, scaling in the x and y dimensions, and one rotation parameter. The second element of the simulated annealing process is a generator which causes random changes in the configurations of the system. As each change is made, an objective function is used to assess the effect of the change; the objective function for the current problem calculates the sum of the squared distance between the corresponding points of the base contour and the transformed candidate contour. The objective function forms the third component of the simulated annealing process. The final, fourth component is a

control parameter which is analogous to the temperature of a cooling liquid or metal. At high temperatures, the control parameter allows large random changes to occur even if these changes result in a higher cost when the objective function is processed, therefore the system has the ability to escape local minima. As the temperature drops, the ability to escape is reduced and the algorithm makes changes which result in a lowering of the system energy, i.e. the objective function gives a lower score than previous states. The simulated annealing process concludes once the temperature reaches zero. The control parameter also specifies the rate at which “cooling” occurs - i.e. the number of random changes which may occur before a temperature drop occurs, and how large the temperature drop is.

Once each shape has been matched to the base outline, the mean of all the corresponding points from the aligned outlines is calculated. This mean outline becomes the new base against which the complete set of outlines is matched. The process of aligning the shapes, calculating the mean of the aligned shapes and recalculating the pose transformations is repeated until the mean shape changes by less than a threshold. The mean shape change is calculated using the sum of squares distance between the corresponding points on the mean shapes between iterations.

The simulated annealing algorithm was implemented on the Matlab programming platform using an adaptation of the method from [172]. Figure 6.6 shows two outlines used as an input to test the efficacy of the simulated annealing outline matching method. The second outline (the ellipse) was created by applying an affine transformation to the first (circular) outline. Each curve was represented by fifty corresponding points² and initial correspondences are shown by the coloured points. The outlines provide a rigorous test of the algorithm since there are many configurations with a low energy state (matching the circular outlines but without matching the indentations), but the algorithm managed to find the global minimum where the indentations are aligned.

6.3.2 Expert outlines alignment

In the aligning of the expert outlines, each shape was represented by fifty corresponding points. The algorithm converged after a few iterations using a distance threshold corresponding to less than one pixel displacement for any corresponding point pair, between the mean base outline points from one iteration to the

²Sixteen corresponding points are shown in the figure to preserve clarity.




Figure 6.6: Test input outlines for the simulated annealing outline matching algorithm.

next. The algorithm was run with different random initial base outlines with no perceptible change in the final mean base outline. The alignment took approximately one hour to complete on a dual PIII 1GHz processor PC running Windows 2000. However, the alignment needs only to be performed once for each point distribution model so the time requirement is not a critical factor.

The aligned outlines are shown in figure 6.7, which shows that the outline shape is tightly bounded to an approximately elliptical shape with an indentation caused by the probe, which is also the region of most outline variation.

6.4 Point distribution model for the aligned expert outlines

The point distribution model relies on the degree of “correspondency” of the points which are considered to be equivalent between the shapes in the training set. For many shapes, the choice of some corresponding points is obvious - for example corners of objects, or junctions between lines on the object’s boundary are obvious landmarks. In such cases, the simplest method of establishing correspondency is for a human observer to annotate each of the example shapes with a set of corresponding points. However, more corresponding points are required for the modelling than would be provided by the landmarks alone, therefore additional




Figure 6.7: Outline results from the simulated annealing automated alignment.

points are added, often at equally spaced locations between the landmark points. Figure 6.8 shows the landmark and intermediate points for a pair of rectangles.

Hand annotation of the shapes is tedious and time consuming. More importantly in the case of the prostate outlines, no suitable landmarks exist making the choice of corresponding points arbitrary and difficult. Therefore, some method of automatically applying correspondence is required.

6.4.1 Selecting corresponding points

Two methods of finding the corresponding points have been previously discussed in this thesis. The first method assigns correspondence to points at equal angular elevation to a common centre of the shapes – this is the polar method used by Arambula and Davies [5] and Wu et al [225]. Whereas this method was unable to accurately represent all the training shapes in the feature modelling method described in section 6.2, the transformed, aligned expert outlines are continuous in the $r\theta$ polar representation using the centroid of the mean shape as the origin. However, this does not hold for the general case of modelling arbitrary objects, and it is certainly possible to have other prostate shapes which would still not be representable in polar form after the alignment phase.

The second method of assigning correspondence uses the adaptation of the approach described by Chalana and Kim [38], described in section 5.3.4. A third, obvious method is to divide each contour so that the distance between adjacent




Figure 6.8: Manually annotated landmark points and automatically placed intermediate corresponding points.

points is the same for a given contour. Results for the point distribution models created with these three methods of establishing correspondency are given in the following sections.

6.4.2 Model formation

Once the corresponding points are established for each outline in the training set, the point distribution model is calculated as follows. Let there be N_s shape vectors $\mathbf{s}_i$ defined as:

$$\mathbf{s}_i = [s_1, s_2, s_3, \dots] \quad (6.1)$$

In the general Cartesian case, there are n corresponding points, each with an x and a y coordinate, therefore:

$$\mathbf{s}_i = [x_1, y_1, x_2, y_2, \dots, x_{n-1}, y_{n-1}, x_n, y_n] \quad (6.2)$$

These vectors may be represented as points in a hyper-ellipsoid in $2n$ -dimensional space. The point distribution model (PDM) aims to model the distribution of these points and provide a parameterisation of the model which may be used to produce new example shapes which are similar to those in the training set. First,

the mean of all the vectors is found as:

$$\bar{\mathbf{s}} = \frac{1}{N_s} \sum_{i=1}^{N_s} \mathbf{s}_i \quad (6.3)$$

Next, the covariance of the data is calculated:

$$\mathbf{C} = \frac{1}{N_s - 1} \sum_{i=1}^{N_s} (\mathbf{s}_i - \bar{\mathbf{s}}) (\mathbf{s}_i - \bar{\mathbf{s}})^T \quad (6.4)$$

Using principal components analysis, the main axes of the point distribution in n -dimensional space are calculated. This equates to finding the eigenvectors ($\mathbf{b}_i$) and associated eigenvalues (λ_i) of the matrix $\mathbf{C}$, then ranking them according to the eigenvalues (largest first). The eigenvalues give the variance of the data around each of the eigenvector axes, therefore only the principal eigenvectors are kept – such that the model retains a high percentage of the total variance of the data (in this case 98% is used). Thus, the first t ranked eigenvectors are retained according to:

$$\sum_{i=1}^t \lambda_i \geq k \lambda_T \quad (6.5)$$

where $\lambda_T = \sum_{i=1}^n \lambda_i$ is the total variance of the data and $k = 0.98$.

New example shapes can be created starting from the mean of all the outlines and applying differing weights upon each of the principal component axes. Mathematically:

$$\mathbf{s}_{candidate} = \bar{\mathbf{s}} + \sum_{i=1}^t \mathbf{b}_i w_i \quad (6.6)$$

Following the assumption of Gaussian distributions for the data along each of the principal axes, the weights are allowed to vary within limits set by the eigenvalues, such that:

$$-3\sqrt{\lambda_i} \leq w_i \leq 3\sqrt{\lambda_i} \quad (6.7)$$

6.4.3 Polar corresponding points

Figure 6.9 shows the corresponding points and mean outline from the aligned outlines where the points are taken at equal angular elevations referenced to the centroid of the mean outline. There is significantly less variation in the point spread at any given angle in comparison to the feature modelling points set shown in figure 6.3. Figure 6.10 shows the effect of independently varying the weights of each of the first four principal components for the polar PDM prostate




Figure 6.9: Polar corresponding points from the aligned expert outlines.

model. The variation in the curves shown in figure 6.10 is markedly less than the curves from the feature modelling approach (figure 6.4). Since the curves were aligned prior to the model formation, the first principal component does not exhibit the large size change apparent in the most significant feature modelling component. However, the outline created with the maximum value in the first weight (the yellow curve in figure 6.10a) exhibits concavities at the side lobes which are uncharacteristic of prostate shapes.

Figure 6.11 shows the effects of varying more than just a single weight at a time. The outlines shown in figure 6.11a) result from varying the odd numbered weights (1st component, 3rd component, *etc.*) from the minimum to the maximum value, and the even numbered weights from the maximum to minimum value. In figure 6.11b), each weight is varied from the minimum to maximum value. These outlines represent some of the extremities of the hyper-ellipsoid space occupied by the model and display sharp indentations which are “un-prostate like”. However, the curves do not overlap themselves as observed in the feature modelling extremities shown in figure 6.5.

6.4.4 Equally spaced corresponding points

Figure 6.12 shows the corresponding points and mean outline from the aligned outlines where the points are taken at regular intervals along each contour. The starting point for each contour is at zero degrees with respect to the contour centroid, as can be seen by the close clustering of the points on the right hand




Figure 6.10: Effect of varying the weights on the polar PDM prostate model. The variance of each component lessens from a) through d).




Figure 6.11: Example curves from the extremities of the polar corresponding points PDM model.

side of figure 6.12. Since the contours are different lengths, the clustering of each corresponding point becomes more spread when moving away from the starting point (towards the left hand side of the figure).

Figure 6.13 on page 199, shows the effect of independently varying the weights of each of the first four principal components for the PDM prostate model constructed using equally spaced corresponding points. All the curves given in figure 6.13 are valid examples of prostate shapes Figure 6.14 (on page 200) shows the effects of varying more than just a single weight at a time using the same methods of adjusting the weights as used previously with the feature modelling model (figure 6.5, page 188) and the polar corresponding points model (figure 6.11). The resulting curves are much improved in terms of their “closeness” to what would be expected of a prostate like shape – i.e. they do not exhibit the sharp corners presence in the previous models.

6.5 Optimizing the corresponding point positions

The PDM models described in the previous section using pre-aligned contours give improved results over the feature modelling method described in section 6.2. It is also clear that the positions of the corresponding points on the aligned contours play an important role in the success of the model – as demonstrated by the




Figure 6.12: Equally spaced corresponding points from the aligned expert outlines.

differences between the extremities of the polar and equally-spaced corresponding point models. However, at present there is no quantitative measure available for how good a particular model is, rather the comparison have been based upon qualitative visual inspections of the sample contours resulting from each model.

In order to optimize the positions of the corresponding points, a quantitative measure is required to assess how optimal the resulting PDM model is. Previous work in this area has tended towards assigning correspondency to equally spaced points along each contour [12] (as used in section 6.4.4), or to shape features present on the outlines [12, 106, 207, 215], such as areas of similar curvature. However, none of this work shows that the models are in any sense optimal.

Several authors have used the variance of the PDM model to provide a quantitative assessment of the optimality of a model. Hill and Taylor [93] use the total variance (the sum of the eigenvalues of the co-variance matrix from equation 6.4) as an objective function in the optimization problem and adjust the corresponding points so as to minimise the total variance. Rangajaran et al [176] also optimize to reduce the total variance, but include the alignment parameters by which pairs of contours are matched to one another. Kotcheff and Taylor [112] use genetic algorithms to find the optimal parameters for the alignment pose transformations and corresponding points positions. The objective function is calculated from the determinant of the covariance matrix of the model and therefore favours models with fewer, more significant modes of variation.

Table 6.1 gives the values of variance for the four most significant components




Figure 6.13: Effect of varying the weights on the equally spaced corresponding points PDM prostate model.




Figure 6.14: Example curves from the extremities of the PDM constructed from equally spaced corresponding points.

of the three models presented thus far. The total variance for each model, and the relative percentage of each component are also shown, along with the number of components required to represent 98% of the total variance, according to equation 6.5. The feature modelling approach requires the lowest number of components (10) in order to achieve 98% of the total variance, but the dominance of the first component is evident – 70% of the total variance. This is corroborated by the visual effects of the first component as shown in figure 6.4a) on page 187, *i.e.* the first component is dominated by the change in size of the prostate shapes rather than any change in shape. The total variance for the feature modelling approach is also significantly greater than the total variance for the two aligned PDM models, which matches the significantly poorer results from the feature modelling approach which were shown previously in figure 6.5 on page 188.

The variances of each component of the two aligned PDM models are similar, with the equidistant points model showing lower variance in all but the first component. The total variance of the equidistant points PDM is also significantly lower than the total variance of the polar points PDM. Since the equidistant points PDM model provides the best results visually (compare figure 6.10 with 6.13, and figure 6.11 with 6.14), the results seem to corroborate the use of the total variance as an objective function for the optimization problem, as used by previous researchers. However, the results of the following section show that a lower total variance does not necessarily lead to a better model qualitatively.

Component i	Variance λ_i		
	FM	Aligned Polar	Aligned Equidistant
1	22279	2026	2321
2	4040	721	464
3	1494	623	409
4	1304	362	230
Total Variance λ_T	31837	4733	4171

Component i	Percentage Variance		
	FM %	Aligned Polar %	Aligned Equidistant %
1	70.0	42.8	55.6
2	12.7	15.2	11.1
3	4.7	13.2	9.8
4	4.1	7.7	5.5
Number of modes t	10	15	15

Table 6.1: Variances of the first four principal components for the feature modelling points model, aligned polar corresponding points PDM and aligned equidistant corresponding points PDM models. The total variance (λ_T) for each model is given along with the percentage variance of each of the four principal components. The number of modes (t) is the number of principal components required so that 98% of the total variance of the corresponding points models are contained in the respective models.




Figure 6.15: Modified Chalana and Kim corresponding points from the aligned expert outlines.

6.5.1 Modified Chalana and Kim corresponding points

The third method of establishing corresponding points on the set of aligned expert outlines uses the modified Chalana and Kim method described in section 5.3.4. The resulting corresponding points and mean outline are shown in figure 6.15. Table 6.2 gives the variances and percentage variance of the first four principal components of the model, as well as the total variance and number of modes to produce 98% of the total variance. The total variance for the model is much less than the totals for the polar and equidistant corresponding points models presented in the preceding section.

Figure 6.16 shows the effect of independently varying the weights of each of the first four principal components. Each of the outlines in figure 6.16 is an acceptable prostate shape. However, the model does not perform as well when more than one weight is adjusted at a time. Figure 6.17 shows the outlines resulting from changing the weights according to the same regime as used previously – alternately changing the weights from minimum to maximum and maximum to minimum in figure 6.17a), and adjusting all weights between the minimum and maximum in figure 6.17b). The resulting outlines are worse than those for the equidistant points model (in figure 6.14), exhibiting sharp corners and kinks.




Figure 6.16: Effect of varying the weights on the modified Chalana and Kim corresponding points PDM prostate model. The variance of each component lessens from a) through d).

Component i	Variance λ_i	Percentage Variance
1	965	41
2	369	16
3	238	10
4	160	7
	Total Variance λ_T	2342
	Number of modes t	22

Table 6.2: Variances and percentage variances of the first four principal components for the modified Chalana and Kim corresponding points model. The total variance (λ_T) and number of modes (t) to represent 98% of the total variance are also given.




Figure 6.17: Outlines from the extremities of the PDM constructed using the modified Chalana and Kim corresponding points method.

6.5.2 Model information

The results from section 6.5.1 show that optimizing the points correspondences on the basis of minimum total model variance does not lead to a model which reflects the twin goals of 1) generality – the model is able to produce any of the shapes in the training set, and 2) specificity – only legal examples of prostate shapes should be produced. This is a somewhat surprising result since total variance has been used as a scoring parameter in many of the optimizing regimes found in the literature. However, the use of the total variance alone neglects the influence of the number of modes of variation present within the model – despite having a far lower total variance than the preceding models, the modified Chalana and Kim corresponding points model also has many more modes. Therefore, the model must be optimized by some means which takes both the variance of each mode and the total number of modes into account.

6.5.3 Minimum description length

The work of Davies et al [58, 59, 60] uses the concept of the *model information* as a parameter for optimization. They argue that the best model balances the cost of the model complexity against the ability of the model to fit the data, using the principle that the simplest description provides the best general model. The measure of this “simplicity” is made in terms of the *minimum description length* principle of Rissanen [181, 182]. The minimum description length parameter is based upon the concept of the cost of trying to transmit the model information across a network, where the data stream must include not only the data values, but also the model parameter values.

The following is a summarized version of the minimum description length derivation given in [60]. To calculate the description length for a given set of corresponding points, the PDM is constructed in exactly the same manner as described in section 6.4, except that the model is not approximated to a high percentage of the total variance. By noting that only the $N_s - 1$ eigenvalues are not identically zero, equation 6.6 may be adapted to include the complete model as:

$$\mathbf{s}_i = \bar{\mathbf{s}} + \sum_{j=1}^{N_s-1} \mathbf{b}_j w_j \quad (6.8)$$

The eigenvectors $\mathbf{b}$ form the principal axes of the data points in n -dimensional

space. The data points, projected onto these new axes become:

$$\mathbf{y}_j = \mathbf{s} \cdot \tilde{\mathbf{b}}_j$$

where

$$\tilde{\mathbf{b}}_j = \frac{\mathbf{b}_j}{\|\mathbf{b}_j\|}$$

are the normal vectors. The range of values of all the original points $\mathbf{s} = \{s_{i\alpha}\}$ is

$$-\frac{r}{2} \leq s_{i\alpha} \leq \frac{r}{2} \quad i = 1 \text{ to } N_s, \alpha = 1 \text{ to } 2n$$

which gives the range R for the projected values $\mathbf{y}_j = \{y_{j\beta}\}$ as:

$$R = r\sqrt{n}, \text{ therefore } -\frac{R}{2} \leq y_{j\beta} \leq \frac{R}{2}$$

The projected data must be quantized into discrete levels for transmission. The quantization parameter Δ is chosen so that the quantization accuracy is of the order of one pixel in the original data. Along each of the principal axes, the data is modelled as a Gaussian distribution of variance σ and mean μ . The mean of the data is calculated and quantized to the same accuracy, Δ , to give the quantized mean $\hat{\mu}_j$. The variance of the quantized data is then calculated as:

$$\sigma^2 = \frac{1}{N_s} \sum_{j=1}^{N_s} (\hat{y}_j - \hat{\mu}_j)^2$$

with an upper limit of $\sigma_{max} = \frac{R}{2}$, and a lower limit of $\sigma_{min} = 2\Delta$. The variance is approximated to a quantization level δ .

The length of code required to transmit the model parameters is given by:

$$L_{params} = 1 + \ln\left(\frac{R}{\Delta}\right) + \ln\left(\frac{\sigma_{max} - \sigma_{min}}{\delta}\right) + |\ln \delta|$$

The description length of code required to transmit the model parameters and the data depends on the variance σ of the data. For $\sigma \geq \sigma_{min}$:

$$DL = L_{params} - N_s \ln \delta + \frac{N_s}{2} \ln(2\pi\sigma^2) + \frac{N_s}{2} + \frac{N_s \delta^2}{12\sigma^2}$$

For the case where $\sigma < \sigma_{min}$, but the data still occupies more than one bin, i.e.

$\max(\hat{y}_j) - \min(\hat{y}_j) > \Delta$, the description length is:

$$DL = L_{params} + \frac{N_s}{2} \ln(2\pi\sigma_{min}^2) - \frac{N_s\delta^2}{24\sigma_{min}^2} + \frac{N_s\sigma^2}{2\sigma_{min}^2} \left(1 + \frac{\delta^2}{4\sigma_{min}^2}\right)$$

Finally, for the case where all the data lies in one bin, only the mean need be transmitted, with a description length:

$$DL = \ln\left(\frac{R}{\Delta}\right)$$

6.5.4 Outline parameterisation

The previous section described the procedure for calculating a score for a particular set of corresponding points. The procedure may be used as the objective function for an optimization regime which aims to find the best set of corresponding points for a set of curves. All that remains is to define a method by which the optimization regime can adjust the position of the corresponding points. Davies et al [58] use a set of parameterizations functions $\phi_i(t)$ which control the corresponding point positions of each curve. The parameterisation functions are monotonically increasing functions of t with a range of $0 \leq \phi(t) \leq 1$, where t is the normalized distance along the contour. Figure 6.18 shows an example parameterisation function, and the resulting points on a circular example contour. The points are sampled uniformly along the x-axis, and the corresponding y-axis values are taken as the distance from the starting point of each of the points on the curve. Therefore, the positions of the corresponding points on the contours are controlled by adjusting the parameterisation functions $\phi_i(t)$. The parameterisation functions can be described by a cumulative distribution function of the form³:

$$\phi(t) = \epsilon + \int_0^t \rho(x) dx$$

where $\rho(x)$ is approximated by a sum of Gaussian functions:

$$\rho(x) = c \left[1 + \sum_i \frac{A_i}{\sigma_i \sqrt{2\pi}} \exp\left(-\frac{1}{2\sigma_i^2} (x - a_i)^2\right) \right]$$

where

$$c^{-1} = 1 + \sum_i \frac{A_i}{2} \operatorname{erf}\left(\frac{1 - a_i}{\sigma_i \sqrt{2}}\right) + \sum_i \frac{A_i}{2} \operatorname{erf}\left(\frac{a_i}{\sigma_i \sqrt{2}}\right)$$

³The method of forming the parameterisation functions is described fully in [58]




Figure 6.18: a) Points of correspondence on a circle determined by b) the parameterisation function $\phi(t)$. (Taken from [58]).

and

$$\operatorname{erf}(x) = \frac{2}{\sqrt{\pi}} \int_0^x e^{-r^2} dr$$

The parameterisation is controlled by the heights A_i , widths σ_i and positions a_i of the i Gaussian curves and ϵ , the offset from the origin of each of the outlines. For a given number of Gaussian functions, a direct optimization of all the parameters leads to a search space which is too large for efficient convergence. Therefore, the parameters are found in a multi-resolution approach as defined in [58].

The initial stage uses a single Gaussian of width $\sigma_1 = \frac{1}{4}$ at the position $a_1 = \frac{1}{2}$, and optimizes its height A_1 and the offset ϵ . For the second level, two additional Gaussian curves are added at positions $a_2 = \frac{1}{4}$ and $a_3 = \frac{3}{4}$, each with a width of $\sigma_{2,3} = \frac{1}{8}$, and their heights are optimized. At each subsequent level k , an additional 2^{k-1} Gaussian curves are added, at intervals of $\frac{1}{2^k}$, and with widths $\left(\frac{1}{4}\right) \left(\frac{1}{2}\right)^k$, and the optimization algorithm adjusts the heights to give the best result for the objective function.

6.5.5 Optimized corresponding points

The parameterisation scheme and minimum description length were used as the basis for optimizing the corresponding point positions on the aligned expert




Figure 6.19: Optimized corresponding points from the aligned expert outlines.

outlines using the simulated annealing optimization algorithm described in section 6.3.1. The algorithm was run for four iterations, giving a total of 15 Gaussian functions to describe each parameterisation function. The algorithm was implemented on the Matlab programming platform and processing time was approximately 72 hours on a dual 1GHz PC running Windows 2000. Although this processing time is extremely lengthy, it only needs to be performed once for a single set of outlines, so execution speed is not a limiting factor.

Figure 6.19 shows the optimized corresponding points and mean outline from the aligned outlines. Figure 6.10 shows the effect of independently varying the weights of each of the first four principal components for the optimized model. The outlines in figure 6.20 are all acceptable prostate shapes. Figure 6.21 shows the effects of varying more than just a single weight at a time, using the same method of adjusting the weights as with the previous models. The worst outlines in figure 6.21 are rather un-prostate-like, since they exhibit concave regions in areas other than around the probe, however, the outlines are smooth without the corners and kinks observed in some of the previous models. These outlines are qualitatively the best of the results when compared to the previous models.

Table 6.3 lists the total variance and objective function score for all the models. The results clearly show that the values of the minimum length description for each model is much more closely aligned to the qualitative visual results, than the corresponding total variance value.




Figure 6.20: Effect of varying the weights on the optimized PDM prostate model. The variance of each component lessens from a) through d).

Model	Total Variance λ_T	Description Length
Feature Modelling	31837	36408
Aligned Polar	4733	35377
Aligned Equidistant	4171	32944
Aligned Modified Chalana and Kim	2342	33774
Aligned Optimized	3874	32906

Table 6.3: Total variances and description length objective function scores for the models. The description length results mirror the qualitative results seen in the example outlines from each model.




Figure 6.21: Outlines from the extremities of the PDM constructed using minimum description length optimized corresponding points.

6.6 Statistical analysis

The work of the preceding sections of this chapter has attempted to determine an optimal statistical shape model for the prostate. The purpose of the model is to provide candidate shapes for use in a genetic algorithm which determines the prostate outline in an ultrasound image. At each iteration of the genetic algorithm the population of shapes are assessed so that outlines which give a good fit to the underlying ultrasound image are proposed as candidates for parenting the next generation of shapes. The work presented in the remainder of this chapter describes the approach adopted to best analyze the fitness of any candidate shape to the ultrasound image.

Any technique for finding the prostate within an ultrasound image must ultimately obtain its information from the grey-levels of pixels within the image. Many techniques for finding the prostate were discussed in chapter 4. Most of these (with the exception of the artificial neural net-based methods) assume that there is one or more characteristics of the image in the locality of the prostate boundary that allows the boundary to be identified. However, the suitability of the characteristic that has been chosen as the boundary identifier (often some combination of the first and second derivatives of the low-pass-filtered image density gradient perpendicular to the outline) can only be assessed as part of the overall outlining algorithm.




Figure 6.22: Histogram result for an ideal statistical measure.

In the rest of this chapter, the properties of the ultrasound image in the region of the boundary are examined using statistical means to evaluate whether these different assumed characteristics of prostate boundaries do actually correlate to the known boundary on a set of images. A number of different candidate indicators are evaluated by directly comparing how well they correlate to the known correct boundary rather than by their efficacy as part of some outlining algorithm.

The approach adopted by previous researchers to find the prostate in ultrasound images is top-down in nature. The input image is filtered and assessed in terms of a particular indicator with the expectation that the indicator will give a high score for the prostate boundary. In many cases, the indicator used combines the zero-crossing of the second-derivative (the position of edges in the image) with the first-derivative (the strength of the edges). The algorithms then attempt to select and combine the sections which score highly for the indicator.

The basis for the statistical analysis described in this chapter questions this assumption that the first-derivative/second-derivative indicator is the best measure for selecting prostate boundaries. The work follows a bottom-up approach; *i.e.* starting from the known expert-drawn boundaries, assess many indicators in the image in both the boundary and non-boundary regions. The search aims to find the indicator which best separates the boundary and non-boundary regions correlated across the set of images. The prostate/non-prostate distributions for an ideal indicator are shown in figure 6.22.

The statistical analysis follows a Monte-Carlo type approach. Given an input image and a statistical test, the analysis computes a histogram of values of the statistical test at points in the image where the boundary is known not to occur – the null hypothesis. The same statistical test is then applied to points on

the expert drawn outline, and the histogram for these results calculated. If the histograms are similar to those shown in figure 6.22 then a simple thresholding (the dotted red line) may be used to test how well a candidate outline matches the prostate boundary.

There are two components which may be adjusted for any given outline, namely the image and the particular statistic measured. Varying the image involves applying different filters to the image with the aim of emphasizing particular image features. The set of filters which have been used to produce input images for the statistical analysis are described in section 6.7. The statistical measures are described in section 6.8.

6.7 Image pre-processing

6.7.1 Speckle reduction

The indicators used for the statistical analysis will be affected by the presence of speckle within the image. Twelve different speckle-reduction algorithms from the literature have been implemented and applied to the ultrasound images. These speckle reduction algorithms attempt to reduce speckle by smoothing the image in areas where the pixel values do not vary much, whilst maintaining the differences in areas where the grey-levels change considerably. The speckle-reduction algorithms are described in appendix A which also gives results of applying the algorithms to a typical image.

All the algorithms have input parameters which affect the degree of smoothing applied to the input image. Each algorithm has been applied four times to each image, with the parameters varied so the results range from light smoothing to heavy smoothing (the degree of smoothing is qualitatively judged). The speckle reduction algorithms are listed in table 6.4.

6.7.2 Edge-enhancement

During the search for speckle-reduction algorithms, three algorithms were found which are classed as speckle reduction algorithms, but perform in a very different manner to the region smoothing algorithms described in section 6.7. These algorithms operate by accentuating edges within the image and suppressing re-

Speckle reduction filters (section/reference)	Parameters
Anisotropic diffusion (A.1.1/[163])	$\lambda = 0.25$ $n = 100$ $K = 2, 4, 8, 16$
Anisotropic diffusion (A.1.1/[202])	$\lambda = 0.25$ $n = 100$ $K = 4, 8, 12, 16$
Homogeneous region growing median (A.1.2/[111])	$x = 7$ $y = 7$ $H_T = 0.01, 0.03, 0.05, 0.1$
Median (A.1.3/[167])	$n = 5, 9, 15, 23$
Separable median (A.1.4/[132])	$n = 5, 9, 15, 23$
Variable length median (A.1.5/[140])	$C = 0.2$ $K = 2, 4, 7, 11$
Double window modified trimmed median (A.1.6/[125])	$m = 3, 5, 9, 13$ $n = 7, 9, 15, 23$
Adaptive weighted median (A.1.7/[140])	$w = 99$ $c = 20$ $K = 2, 4, 7, 11$
Gaussian (4.3.2/[168])	$n = 12, 24, 36, 48$ $\sigma = 4, 8, 12, 16$
Directional median (A.1.8/[48])	$\text{thickness} = 3$ $n = 5, 9, 15, 23$
Maximum sticks (A.1.9/[49])	$\text{thickness} = 3$ $n = 5, 9, 15, 23$
Decision-directed sticks (A.1.9/[50])	$\text{thickness} = 3$ $n = 5, 9, 15, 23$

Table 6.4: Speckle reduction algorithms and parameter values

Edge enhancement algorithms (section/reference)	Parameters
Order statistics (A.2.1/[164])	$m = 5, 9, 15, 23$ $k = 1, 2, 4, 6$
Rotating kernel transformation (A.2.2/[126])	$\text{thickness} = 3, 3, 5, 5$ $n = 15, 23, 15, 23$
Modified ratio of averages (A.2.3/[232])	$n = 5, 9, 15, 25$
Sobel (4.2.2.2/[168])	—
Laplacian (4.2.3/[168])	$\alpha = 0.2$

Table 6.5: Edge enhancement algorithms and parameter values

gions of relative smoothness, therefore they are classed here as edge enhancement algorithms rather than speckle-reduction algorithms.

Each of the images from the pre-processing stage (including the initial input images with no pre-processing) is used as an input for the edge-enhancement algorithms. Five edge-enhancement algorithms have been implemented, listed in table 6.5. Three of the algorithms have parameters which determine the size of the edges which are enhanced. These algorithms have been applied with varying parameter values so that they progressively ignore smaller edges in the input image (again the parameter selection is qualitative in nature, based on visual results as presented in appendix A).

6.7.3 Statistical image input dataset

The sixty-two ultrasound images from the data set described in section 5.2, chapter 5, have been used as initial inputs for the statistical analysis. Each of these im-

ages has also been filtered with the twelve speckle reduction filters (section 6.7.1) and five edge enhancement algorithms (section 6.7.2) using the various parameter values given in table 6.4 and table 6.5 to produce a further set of 3720 images. Finally, each of the results of the speckle reduction algorithms has been further filtered for edge-enhancement, to create an additional set of 40320 images. Therefore, there are a total of 44100 input images for the statistical test.

The total size of the dataset is large and some of the combinations of filters produce outputs in which all the useful ultrasound information is destroyed, *e.g.* some results only contain the edge information for the strongest edge in the input – the periphery of the ultrasound probe. Therefore, the filters should be tested on a single image to remove those combinations which destroy all the useful information. However, the software to calculate the dataset, and to perform the statistical analysis (described in the proceeding sections) was written so that it runs completely without user intervention, thus no interim checking of the results was performed, and the full set of images was processed.

6.8 Statistical measures

6.8.1 Point, line and patch selection

The aim of the statistical analysis is to find an indicator which best distinguishes the boundary from non-boundary regions. This section describes the approach used for selecting the points at which to calculate the indicators used in the statistical analysis.

One hundred points on each expert drawn boundary are selected at equal distances along the length of the boundary. These points form the set at which the boundary indicators will be assessed. However, it is unlikely that the expert drawn boundary will fall exactly onto the pixels with the best significant indicators in the boundary proximity. Therefore, the analysis is extended to look for indicators along line segments and patches at the boundary. At each selected point on the boundary, the normal to the boundary is found (the normal direction is always chosen to lie in the external region of the boundary). The line segments are selected from the normals extending ten pixels either side of the boundary (thus each line segment is 21 pixels in length). The ten pixel length is chosen to compensate for an error on the expert delineated outline of up to ± 5 mm in the normal direction, and the statistical analysis will still find the most




Figure 6.23: Bi-linear interpolation

appropriate value. The patches are created from 21 by 10 pixel patches centred on the boundary point of interest and oriented so that they lie along the normal to the boundary. The justification for using patches is to use the complete boundary rather than sections of it. However, the patches involve significantly more calculation.

The value of the pixels in the lines and patches are calculated using bi-linear interpolation. Referring to figure 6.23, the value of an interpolated pixel W is calculated as:

$$W = \alpha\beta G + \alpha(1 - \beta)R + \beta(1 - \alpha)B + (1 - \alpha)(1 - \alpha)Y \quad (6.9)$$

where G , R , B and Y are the grey-level values of the four pixels. A sample point (A), line (B) and patch (C) for an outline are shown in figure 6.24. The green/red line (D_i/D_e) and patch (E_i/E_e) in figure 6.24 are used for additional statistical measures where the values are calculated for the internal and external line/patch separately.

To calculate the statistical indicator under the null hypothesis, sets of 100 points were randomly chosen. These points were taken from the areas of the image which did not form part of the expert drawn boundaries, or the non-ultrasound areas (the black areas of the image at the probe and the extremities of the beam). Since there are multiple expert boundaries for a single image, the out of bounds region for the null hypothesis points includes all of the area bounded by the expert boundaries. Figure 6.25 shows an ultrasound image overlaid with expert boundaries and the regions of the image where the points are randomly chosen.





Figure 6.24: Point, line and patch regions for the statistical analysis.



(a)

Figure 6.25: a) Expert boundaries on a typical ultrasound image. b) Resulting null hypothesis regions (white areas).

Random orientations are also calculated at each of the random points so that lines and patches may be generated. This orientation is also used to define the pseudo internal/external parts for the random lines and regions.

6.8.2 Indicators

The indicators differ depending on whether a point, line or region is chosen. For points, the actual grey-level value of the indicator at the point is the only statistical measure.

When the line segments and patches are used, six different metrics are used. These are the maximum, minimum, mean, median, variance and max/min difference of the pixels within the lines or patches. These same tests have been applied separately for the internal/external split lines and patches. In addition, eight metrics have been used for the internal/external split lines and patches, which are the differences in internal and external value for each metric (max, min, mean, *etc.*) and the maximum internal minus minimum external and minimum internal minus maximum external value. The 45 metrics are summarized in table 6.6.

6.9 Analyzing the indicators

The statistical analysis of the outline indicators was conducted in three stages. The first stage assessed whether the outline indicators gave results which were significantly different for the expert drawn outlines compared to the indicators for the sets of random points.

6.9.1 Kolmogorov-Smirnov test

The first part of the statistical analysis attempts to answer the question “does the statistical measure give significantly different results for the region in the image which contains the prostate boundary in comparison to the regions of the image without the boundary?”. The statistical measure (for a given image, filter and statistic) provides two datasets, the first for the points on the experts’ boundaries, and the second for the non boundary areas. Therefore, the previous question may be simplified to “are the two datasets drawn from the same distribution function, or from different distribution functions?”. The null hypothesis in this case is that the two datasets come from the same distribution. Therefore, the objective is

Points A
grey-level value
Lines B, D _i , D _e and patches C, E _j , E _e
maximum grey level
minimum grey level
mean grey level
median grey level
max minus min
Lines D and patches E
max internal - max external
max internal - min external
max external - min internal
min internal - min external
mean internal - mean external
median internal - median external
(max-min internal) - (max-min external)

Table 6.6: Statistical measures. The points, lines and patches (A–E) are described in figure 6.24.

to ascertain whether, to a given level of significance, the null hypothesis may be rejected.

The Kolmogorov-Smirnov [36] (KS) test is ideally suited to assessing the null hypothesis in the current application since it makes no assumptions about the distribution of the datasets (it is non-parametric and distribution free). This is in contrast to most statistical tests (*e.g.* the Student's t-test) which do require the data sets to fit some assumed distribution. The KS test calculates an empirical cumulative distribution for each of the samples and assesses the likelihood of the null hypothesis being true based on the maximum deviation between the two distributions.

The level of significance is an important factor since it determines the degree of confidence that may be placed in the KS test. Common values are 5% or 1%, meaning that (for the 5% level) one out of every twenty tests where the null hypothesis is ascertained as true will be a false conclusion. However, reducing the significance level comes at the cost of rejecting the null hypothesis in many more cases where it might be true. A further aspect of the significance level is that the results from the test are cumulative. For the current application, there are 100 sets of random results for each set of expert results, therefore to reject the null hypothesis to a level of 5% each individual test must be assessed at a level of 0.05%.

The KS test was used as a first pass test to assess each combination of filter and indicator in the set of 21 expert outlined images. The test was applied to compare each of the four expert outlines and the reference outline for each image with the random results, to an overall significance level of 5%. Images where the null hypothesis could not be rejected for all the random points and outline points were not carried forward to the later stages of analysis. The rejected images comprise those where the combination of filters (speckle reduction and image enhancement) caused the destruction of the edge information at the prostate boundary.

6.9.2 Random points histogram

The second part of the statistical analysis aims to determine those combinations of filter and indicator which give the greatest difference between the random points and the points from the expert and reference outlines. For each combination, the mean indicator value was calculated for the sets of random points and the outline points. The histogram of the random point indicator means was constructed,




Figure 6.26: a) Original input prostate image. b) HRGMF filtered image – the best result of the random point analysis.

and the positions of the mean outline indicators compared to the histogram. The basis of the test is to search for the filter / indicator pairs for which the expert and reference results all fall into the first region or all fall into the last region of the histogram, *i.e.* the real outlines can be distinguished from most of the random points by a simple threshold.

During the histogram analysis, many combinations of filters and indicators were found where the expert and reference outline indicator means fell outside the range of values in the random mean indicator histogram. Therefore the combinations which were put forward for the third stage in the statistical analysis were those for which all the expert and reference outlines fell to one side (always the same side) of the histogram for all the 21 images - *i.e.* the expert outlines and reference outlines could always be distinguished using a simple threshold with the same value on all the images.

Figure 6.26 shows one image from the dataset which has been filtered using the homogeneous region growing median filter (HRGMF) (see appendix A, section A.1.2) with $H_T = 0.05$. This filter, combined with the maximum minus minimum line indicator, gave the best overall result – the largest interval between the maximum random mean and the minimum expert mean.

6.9.3 Random outlines histogram

The best result from the random point statistical assessment would suggest that the combination of the HRGMF and maximum minus minimum line indicator is the best method to use as a fitness function in the genetic algorithm outline detection scheme. However, although the test provides a clear distinction between actual outline points and random points, it does not necessarily follow that the best result comes from the ideal test for the fitness function, particularly because the random points are never chosen from the region around the actual boundary.

The promising results from the random point analysis (those where the expert and reference outlines were always outside the random histogram) were further processed to assess their performance at distinguishing between candidate outlines and the expert outlines. This third test analyses the results in an identical fashion to the the selection method of the genetic algorithm, *i.e.* for each image a set of 100 random outlines was generated using the optimized outline model of section 6.5.5. The remainder of the test is identical to that used for the random point analysis, except that the histogram of random means was generated using sets of 200 points on each of the 100 random outlines. No restrictions were made on the positions of the random outlines, other than ensuring that they remained wholly within the data regions of the images (*i.e.* the regions outside the probe circle and image extremities; those areas where the pixel value is not zero).

The result of the random outlines histogram analysis shows that the best overall result for the 21 images in the multiple expert outlined set is the combination of the Gaussian filter ($\sigma = 4$) with the “maximum-minus-minimum” line indicator. The HRGMF filter (best result from the random points histogram analysis) also performs to a similar degree of effectiveness at distinguishing the expert outlines from the random outlines. However, the Gaussian filter has been chosen as the best result since its execution speed is far higher than the HRGMF filter. This combination is referred to as the best fitness function in the remainder of the chapter.

Figure 6.27 shows the best and worst results when comparing the best fitness function values for the expert, reference and random outlines for the 21 images in the dataset. The graphs show that even in the worst result, the reference outline (red line) is always outside the histograms of random results (although it is very close in the worst result). The best result shows a clear distinction between all the expert drawn outlines and the random outlines. However, this does not hold for the worst result, where only one expert’s outline is greater than all the random




Figure 6.27: Histograms of random outlines and expert and reference outlines. a) Worst result of the HRGMF/max minus min indicator. b) Best result.

outlines, but these expert boundaries still fall within the top 15% of the random outline results.

The best and worst results from applying the fitness function of Wu's algorithm [224, 225] to the reference, expert and random outlines are shown in figure 6.28. The best results show that Wu's fitness function can score the expert and reference outlines higher than the random outlines. However, the worst result shows that the fitness function can also score the expert and reference outlines so that they fall on both sides of the random histogram, as well as within the histogram. This suggests that the assumption of Wu's fitness function – that the best outline highlights a dark to bright transition when looking from inside the prostate – is false.

The best fitness function was also tested on the complete set of 62 images by comparing the reference outline to a set of 100 random outlines generated for each image. The 62 images were also assessed using the fitness function from the Wu algorithm (with a small modification so that the absolute value of the test was used in each case). Figure 6.29 shows the number of random outlines for which the fitness score was greater than the reference outline fitness score, for both the Wu algorithm and the best result from the analysis.

Figure 6.30 shows the input prostate ultrasound image and the reference outline for which the best fitness function failed. In this particular image much of the




Figure 6.28: Comparing the Wu fitness function for the random outlines and expert and reference outlines. a) Worst result. b) Best result.




Figure 6.29: Comparison of the Wu and best fitness function for all the 62 images in the dataset.




Figure 6.30: Input image for which the best fitness function does not distinguish the reference outline (green) from the random outlines.

boundary is not evident, particularly in the left hand lobe where no ultrasound information is present. The bottom region of the reference outline also shows no clear distinction between the prostate and non-prostate regions. There is also a bright region on the right hand side of the image which is sufficiently “prostate-boundary-like” to make random outlines in this region have a higher fitness score than the reference outline.

6.10 Results

The optimized prostate model and fitness function were implemented into the genetic algorithm and applied to the 62 images in the data set. Each image was assessed 10 times due to the random element involved in the construction of the outline populations in the genetic algorithm.

The best and worst results of the optimized algorithm applied to two different images are shown in figure 6.32 on page 228 (compared to the reference outline using the average distance metric of section 5.4.2.2, chapter 5). The input images were those showing good and poor expert outline agreement; *cf.* the results of

	COD[mm] / (σ)	IOD[mm] / (σ)	WI / [CI]	P(%) / [CI]
Wu _{worst}	13.97 (0.31)	3.73 (0.13)	0.36 [0.35, 0.37]	0 [0, 0]
Wu _{best}	4.63 (0.38)	3.75 (0.19)	0.92 [0.90, 0.93]	52.4 [49.8, 55.0]
Optimized _{worst}	5.55 (0.38)	3.75 (0.17)	0.84 [0.80, 0.86]	61.9 [58.9, 64.9]
Optimized _{best}	4.30 (0.35)	3.75 (0.15)	0.95 [0.93, 0.97]	71.4 [68.2, 74.7]

Table 6.7: Comparing the optimized algorithm's outlines to the experts' outlines using the Hausdorff distance as an error metric. COD is the mean computer to observer difference, IOD is the mean inter-observer difference, σ is the standard deviation, WI is the Williams index, P is the percentage statistic and CI are the respective confidence intervals.

figures 5.14 and 5.15 on pages 175 and 176, respectively.

6.10.1 Multiple expert results

Table 6.8 gives the Hausdorff distances comparing the best and worst outlines from the optimized algorithm with the expert outlines, and the corresponding Hausdorff distances for the best and worst results of Wu's feature modelling algorithm. The Hausdorff distance results show that the best outlines from the optimized algorithm are significantly better than the best results of Wu's algorithm. This is confirmed in the William's indices and percentage statistics which are closer to the optimal values of 1 and 80 respectively, than the corresponding Wu results. However, the major finding of table 6.8 is that the worst results of the optimized algorithm are much better than the worst results of Wu's algorithm. This demonstrates that the optimized algorithm is far more consistent than the feature modelling algorithm of Wu, since it uses a tighter, more compact prostate model, and a more efficient fitness function.

Table 6.8 shows a similar improvement in the optimized algorithm's results versus Wu's algorithm's result when using the average distance to the expert outlines as a metric. The best of the optimized results are slightly better than the best of Wu's algorithm, with the mean computer observer difference only slightly greater than the inter observer difference, This is reflected in the William's index which is unity for the best of the optimized results.

	COD[mm] / (σ)	IOD[mm] / (σ)	WI / [CI]	P(%) / [CI]
Wu _{worst}	9.14 (0.15)	1.22 (0.13)	0.45 [0.44, 0.46]	0 [0, 0]
Wu _{best}	1.31 (0.17)	1.24 (0.15)	0.99 [0.99, 1.00]	47.6 [45.2, 50.0]
Optimized _{worst}	1.68 (0.20)	1.22 (0.14)	0.95 [0.94, 0.97]	66.67 [63.5, 70.0]
Optimized _{best}	1.28 (0.11)	1.24 (0.15)	1.00 [0.99, 1.01]	66.67 [63.5, 70.0]

Table 6.8: Comparing the optimized algorithms' outlines to the experts' outlines using the average distance measure. COD is the mean computer to observer difference in mm, IOD is the mean inter-observer difference in mm, σ is the standard deviation, WI is the Williams' index, P is the percentage statistic and CI are the respective confidence intervals.

6.10.2 Reference outline comparison

Figure 6.31 shows the average boundary distance for the 62 results of each outlining algorithm from the 62 reference outlines. The plot shows that the best results of the optimized algorithm give the best overall results, with the lowest median value (red line), smallest inter-quartile range (blue box) and minimum outlier values (red pluses). This is expected since the best results of the optimized algorithm gave the best results in the Hausdorff and average distance test presented in section 6.10.1.

The worst results of the optimized algorithm give the second best set of results in figure 6.31, performing slightly better than the best results of Wu's algorithm. This is a somewhat surprising result initially, since the best result from Wu's algorithm outperformed the worst optimized results in the tests of section 6.10.1. However, the reference outline comparison was conducted over the whole image dataset of 62 images rather than the subset of 21 images for which multiple expert boundaries were obtained. For the remaining 41 images, the worst results of the optimized algorithm are sufficiently better than the best of Wu's algorithm, such that the overall result is slightly better for the optimized algorithm.




Figure 6.31: Box plots of the average boundary distances from the reference outlines.




Figure 6.32: Optimized algorithm outlines (green – best, red – worst). a) Results for image with close expert agreement (*cf.* figure 5.14 on page 175). b) Results for image with poor expert agreement (*cf.* figure 5.15 on page 176).

6.10.3 Algorithm execution speed

The final part of the assessment of the optimized algorithm was to measure the speed of execution of the algorithm. All the development work for the image processing was conducted using the Matlab platform to program each algorithm. However, despite providing an easy framework for algorithm evolution, Matlab is notoriously slow in execution, particularly for non-linear algorithms where pixel based processing is required.

The optimized outlining algorithm was implemented in C++ using the Metrowerks Codewarrior platform to compile the code. The genetic algorithm code was taken from a freely available library developed by Matthew Wall [216], and the image processing and spline drawing routines from the Magic software source code [142].

The C++ version of the optimized algorithm was tested using boundary populations of 400 outlines, iterated for 2500 generations before termination. The complete process of pre-filtering the image, and running the genetic algorithm takes approximately one minute to execute on a dual 1GHz processor PIII computer (using only a single processor to execute the code).

6.11 Discussion

6.11.1 Shape model optimization

The work of the first part of this chapter developed the concept of the point distribution model and investigated the effect of the assignment of corresponding outline points in the model. Various methods for deciding the correspondence were implemented and assessed in terms of the visual appearance of a discrete set results from each model.

Two quantitative criteria were investigated as possible measures of the optimality of a particular model, namely, the model variance and the description length. The results show that use of the model variance alone is flawed, since there is no correlation between the variance value for and qualitative assessment of each model. A better quantitative parameter could be found by using some combination of the number of modes and the variances of each mode, however, it is difficult to formulate a rigorous argument for why any particular combination would be optimal.

The description length principle does give a better correlation between the visual appearance of the outlines and the quantitative derived parameter, so that the optimized model with the lowest description length also produces the most prostate-like outlines, even at the extremes of the model. However, there are two closely related problems which prevent a definitive statement that the optimized model of section 6.5.5 is indeed the best possible model for the set of prostate shapes. The first problem is that although minimizing the description length seems like a plausible method for approaching the best model, there is no way to directly prove that the model with the lowest description length is the best model. The work of Davies et al [58, 59, 60] followed the principle of Occam's razor (i.e. the simplest model gives the best performance), which is superficially satisfactory, but does not provide a rigorous proof.

The second area of difficulty is that it is not possible to define exactly what is the optimal model. Two aspects of an ideal model have been previously defined – the generality of the model to describe all valid prostate shapes and the specificity of the model to only describe valid prostate shapes. Clearly, the optimized model from section 6.5.5 fails to achieve these objectives, since it allows prostate shapes which would never occur in reality. However, it is certainly possible and probably very likely that an ideal model cannot exist for the particular outlines used in the model creation.

6.11.2 Objective function optimization

The second part of the chapter investigated the underlying statistics in both outline and non-outline regions of the image to try and identify the best criteria for deciding whether an outline is present in a particular position in a candidate image. Many different indicators and filters were included in the analysis to provide a thorough investigation of the properties which determine the presence of an outline. The statistical analysis provides a robust method of determining which of the filter/indicator combinations is most suited to discerning the prostate outline in any given image.

Once again, it is not possible to declare that the best combination for assessing the prostate outlines is the one presented in section 6.9.3, only that this is the best combination selected from those which were included in the analysis. Other filters have been presented in the literature; likewise other statistical indicators could be investigated, and these could give improved results over the current optimum.

For example, no attempt has been made in the current investigation to assess the frequencies (using the Fourier transform) or textural energies associated with the prostate outline. However, such work falls outside the scope of this thesis and forms one possible direction for future work in the field.

One aspect of transrectal ultrasound images which has not been included in the analysis is the polar nature of the image formation process. The ultrasound beam is emitted radially from the probe, and reflected signals are also received radially. In accordance with Snell's Law, the strength of the reflected signal is dependent upon the angle of incidence between the beam and the reflecting surface. Therefore, in areas of the ultrasound image where the prostate outline lies perpendicular to the ultrasound beam, a strong received echo is expected. Conversely, regions where the ultrasound signal and outline are parallel are unlikely to show a strong edge feature. This is apparent in most transrectal prostate ultrasound images, where the top and bottom sections of the outline are comparatively strong, while the side lobes are more poorly discerned. The inclusion of the angle between the ultrasound beam and the outline should be included into the statistical analysis, since it is very likely to improve the overall performance of the outlining algorithm. This forms one major area for further work in the automated outlining of prostates in transrectal ultrasound images, but does not lie within the scope of the current research.

6.11.3 Algorithm speed

The speed of the final optimized algorithm implemented in C++ is not fast enough to achieve the overall goal of monitoring the position and shape of the prostate as it is affected by the movement of the resectoscope during TURP. The current time of one minute to process each image is significantly slower than the ideal, for which the maximum allowable time is of the order of hundreds of milliseconds per image, to achieve a rapid response to changes in the prostate. However, several factors could be investigated which would allow the processing time to be reduced.

The first aspect which could be optimized is the hardware upon which the algorithm is run. Computer processor speeds are constantly increasing, with a common rule of thumb being that processor speeds double every 18 months. The test computer for the algorithm was a dual 1GHz PIII processor machine; at the time of writing top of the range workstations contain 2.8GHz processors, which

could give an instant reduction to approximately 20 seconds per image. Further performance gains could be achieved by adapting the genetic algorithm code so that the task of assessing each outline in the population at any generation is performed in parallel by more than one processor, either in a multi-processor machine, or a more distributed architecture across a network of machines. Some performance gains are likely to result from hand tuning of the assembler code so that fast multimedia features of the processor (such as vector processing) are exploited, redundant code is removed and memory requirements are reduced. Finally, various parts of the algorithm could be implemented in hardware (either dedicated for a particular purpose, or using existing features such as the processing capabilities of the built-in graphics card).

6.11.4 Additional outlining applications

The final objective for the outlining algorithm results is to provide information about the prostate position and shape during surgery; this objective cannot currently be achieved using any of the outlining algorithms presented in this thesis. However, the optimized PDM and genetic algorithm outlining method presented in this chapter gives significantly improved results over previous outlining algorithms presented in the literature. These benefits include both improved correlation with the reference outlines for a set of image, and an improvement in algorithm speed over previously reported execution times (although much of the increase can be attributed to the improvement in computer hardware).

Despite the failure for prostate monitoring, there are many other areas in which the optimized algorithm can provide benefit. Pathak et al [162] define three areas where prostate outlines obtained from transrectal ultrasound images provide key information in the clinical decision making process. The first area is in diagnosis where prostate specific antigen density [34, 173] is a factor in assessing the likelihood of cancer, and prostate size is a crucial parameter in the analysis of benign prostate disease [99]. Both these applications require an accurate assessment of prostate volume rather than the actual outlines, however, various studies have shown that obtaining the prostate volume by outlining a series of scans provides a better estimate than cruder methods based on measuring the height, width and depth of the gland and assuming an ellipsoid shape [212]. The second area where outlining is routinely performed is during the treatment of prostatic disease. Therapies such as brachytherapy [84, 85], high intensity focused ultrasound (HIFU) [94, 141] and cryotherapy [33, 134] all require accurate knowledge of

the prostate outlines so that treatment can be confined to the prostate, and surrounding tissues are not unduly affected. Other therapies such as TURP [35, 208] and transurethral microwave therapy [51, 175] require knowledge of the prostate volume; calculation of the volume from the outlines would again provide more accurate information. The final area for which prostate volume is important is in post-operative treatment assessment, and in the monitoring of non-invasive therapies such as hormonal treatment [220, 217].

6.12 Conclusion

The work presented in this chapter has been conducted with the aim of improving upon the best method for outlining the prostate algorithm from chapter 5 – the feature modelling and genetic algorithm method presented by Wu et al [224, 225]. Two areas were identified which were candidates for optimization – namely, the method of modelling the variation in shape of the prostate, and the method of testing the “fitness” of a candidate outline with the underlying image.

Several methods for generating prostate shape models were investigated, with the best result found using an implementation of the optimization method from Davies et al [58, 59, 60]. This gave a significantly better model than the feature modelling approach, since the model is more compact, and does not exhibit the sharp corners or crossings apparent in the former method.

The criteria for assessing the fitness of each outline was investigated using a novel approach, based upon identifying the features present in the images at the locations of the expert drawn outlines, rather than by assuming that any particular feature best signifies the outline. The approach adopted was to include many potential outline indicators and perform a rigorous statistical analysis to assess the potential of each indicator. The results of the analysis show that the best result performs significantly better than the fitness function used by Wu.

The optimized algorithm was applied to the prostate image dataset and the resulting outlines compared to those drawn by the experts. Whilst still not performing to the level of consistency of the human experts, the optimized algorithm shows a significant step forward towards this goal, especially in the congruity of outlines from the repeated application of the algorithm on a single image.

The original goal of the image processing of prostate images to find the prostate outline was aimed at monitoring the movement and deformation of the prostate

which is apparent during TURP operations. The final section of work presented in this chapter assessed the execution speed of the optimized algorithm. The result of approximately one minute per image shows an improvement over previous researchers' algorithms, however, it is still much too long for the goal of real-time processing to be achieved.

Chapter 7

CASMS TURP monitoring system

7.1 Introduction

The work of the previous chapters has been aimed at finding the outline of the prostate in ultrasound images automatically, without user interaction. The purpose of finding the outline is to use this information to track the position of the prostate, and monitor changes in the prostate shape, during the course of the TURP operation. The outlining work of this thesis has made substantial improvements in the reliability and accuracy of automated outlining over previous techniques detailed in the literature. However, the execution speed of the new outlining algorithm is still not sufficiently swift for the timely tracking of the prostate in a real TURP operation.

The problem of the execution speed precludes the development of a fully functional CASMS system for monitoring TURP procedure in the OT. Despite this obstacle, initial advances towards reaching the CASMS goal have been made. Section 7.2 gives an overview of the CASMS system and describes the additional equipment for such a system over and above that required for the CASTS TURP system. Section 7.3 presents the initial trials of the new equipment in the OT and highlights additional problems for CASMS, other than the image processing.

The remainder of the chapter describes the implementation and testing of a laboratory-based CASMS system in which the problems of the OT need not be considered. The purpose of the laboratory system is to demonstrate that despite the difficulties, ultrasound imaging of the prostate and subsequent processing is a viable method for tracking deformation and movement. The laboratory based CASMS system also provides the additional tracking required to improve the ac-

curacy of the CASTS trainer system discussed in chapter 3. In this respect, the training and monitoring system are essentially the same. However, functionally the emphasis of the trainer is to provide a positionally accurate visual display to the trainee and a true record of performance, while the main objective of the monitoring system is to provide warnings of proximity to resection limits. Therefore, the laboratory based trainer/monitor system is referred to as CASTMS (Computer Aided Surgical Training and Monitoring System) throughout the chapter.

Section 7.4 describes the components and software for CASTMS, including the steps taken to circumvent the problems associated with an OT based CASMS system (section 7.4.1). The image processing required to detect the phantom's outline in ultrasound images is given in section 7.4.2. These outlines must then be collated to form the complete 3D model of the resectable region in the phantom. It should be noted that the algorithm to perform this task is also applicable to formation of a resectable region model from outlines of real prostate ultrasound images; where possible all aspects of CASTMS are directly transferable to the OT based CASMS. The formation of the model is presented in section 7.4.3.

The software for CASTMS is described in section 7.4.4. The software is markedly different from the original CASTS system although the graphical user interface remains practically identical. The main changes are that the architecture has been implemented with concurrent threading for true multitasking (CASTS used a single execution thread) since the computational demands are far greater, and a different programming framework has been used (the Microsoft Foundation Classes used in CASTS are too restrictive and overly bulky for rapid execution).

Section 7.5 presents the testing of the system and the results of a resection trial on a prostate phantom. The results are discussed in section 7.6 which highlights the improvements afforded by the use of ultrasound for prostate motion tracking. Finally, section 7.7 provides a conclusion to the chapter.

7.2 System components

Many of the components used in the CASTS TURP system are also necessary for the CASMS system. The key difference is that while the CASTS system uses and models a prostate phantom, the CASMS system must track changes to the unique prostate within each patient. However, there are other, smaller differences which are required for the in-vivo system. An overview of the CASMS monitoring system for TURP is shown in figure 7.1.




Figure 7.1: Components of the CASMS TURP system.

The surgical equipment used for the real operation is identical to that used in the training system – a desirable feature of the training system since the trainees can familiarize themselves with the equipment prior to surgery. However, the instrumentation attached to track the surgical tools must be adapted to suit the operating theatre environment – the largest difference being that the components must be sterilizable.

7.2.1 Tracking the resectoscope

The resectoscope ring used in the training system was constructed of a material which could survive the sterilization process in an autoclave. However, after several sterilizing cycles, the ring was found to shrink slightly causing difficulty in mounting the ring onto the resectoscope. The strain relief on the connecting wire was also insufficient and did not stand up to the rigours of use, leading to problems with IRED malfunctions. An additional problem with using the ring was that the positional accuracy deteriorated along the perpendicular axis to the ring plane due to the co-planar geometrical arrangement of the IRED trackers.

The original ring design was modified and the new “endotracker” is shown in figure 7.2. The updated design contains additional IRED trackers which are no




Figure 7.2: New IRED tracking component.

longer co-planar to provide improved positional accuracy. The endotracker is constructed from aluminium which improves durability and contains a screw-fit strain relief to prevent damage to the internal wiring. An additional advantage of the aluminium material choice is that the overall weight of the endotracker is reduced so the change to the “feel” of the resectoscope is reduced.

In the CASTS system, motion of the resectoscope’s trigger was monitored using a potentiometer and cog arrangement directly screwed to the body of the resectoscope. This required direct modification of the resectoscope with threaded holes for the clamping screws. This method of trigger monitoring is unsuitable for the OT because the potentiometer does not survive the autoclave sterilization. The position of the potentiometer on the resectoscope and its proximity to the site of operation preclude the use of drapes to form a sterile barrier.

A new design was chosen using the IRED tracking system to provide information on the trigger position in addition to the gross position of the resectoscope. The trigger tracker contains four IRED markers, two on its top surface and two on the underside allowing position tracking regardless of the rotation of the resectoscope. Relatively few markers are required in the design since the tracking system only needs positional information of a single IRED at any given moment, to provide the trigger position relative to the resectoscope body. The trigger tracker fits onto the resectoscope body over the electrosurgical lead terminal and is locked in place by the electrosurgical lead. The tracker is shown in figure 7.3.

7.2.2 Ultrasound probe platform

Ultrasound was discussed in chapter 3 as a means for tracking the position of the prostate phantom in the CASTS system, and in chapter 2 for tracking the real prostate within a patient. For both applications, ultrasound images are required



Figure 7.3: Optical tracking IRED mount for monitoring the trigger motion.

which cover the whole of the resection region within the prostate. Some modern ultrasound scanners are equipped with 3D scanning probes which allow volumetric ultrasound information to be acquired whilst the probe remains stationary. However, for the author's PhD work, the only scanner available was an older ultrasound system¹ which provides 2D ultrasound images. This ultrasound machine is similar those in use in most hospital urological departments - the modern 3D systems are expensive and their benefit has yet to be proved in the urological arena. The potential benefit of 3D ultrasound to the CASMS system is discussed in chapter 8, section 8.5.1.

Obtaining 3D ultrasound information is achievable using a 2D ultrasound machine by moving the ultrasound probe along a line perpendicular to the image plane whilst monitoring the probe position. The concept is not new – previous work in the literature includes methods for obtaining 3D ultrasound using a hand held probe [165], or using a mechanical system for moving the probe in a controlled motion [166]. For the CASMS application, motion by hand is not feasible since the ultrasound information must be continuously updated, and the operative site is too crowded for an additional person to be able to access the probe easily (the surgeon has too many other concerns to expect them to move the probe!).

A motorized platform has been designed which holds the ultrasound probe and allows precise translation of the probe to give registered 3D ultrasound information. The platform can be clamped to a standard photographic tripod for gross positioning, whilst finer positioning may be achieved through manual adjustment of one translational (in the axial plane of the patient) and one rotational (in the patient's coronal plane) axis. The motorized axis of the platform moves the ultrasound probe in a reciprocating motion within the rectum, so that axial ultrasound images through the prostate can be acquired. The power train for the

¹B&K Medical type 1849 ultrasound machine.



Figure 7.4: Ultrasound probe platform.

platform consists of a lead screw which is driven by a stepper motor through a toothed pulley and belt connection. The stepper motor controller is connected to the parallel port of the CASMS PC; each step of the motor is achieved by pulsing a pin on the parallel port, while the state of a second pin determines the direction of motion. The ultrasound probe platform is shown in figure 7.4.

7.3 Initial operating theatre trials

During the course of development of the equipment for CASMS, various trials were conducted in the OT, both during real TURP operations and whilst the theatre was not in use. The purpose of the trials was to obtain feedback from the surgeons as to the suitability of each piece of equipment, principally the impact on the normal routine of the operation.

A large proportion of the OT is occupied by equipment and people during a typical TURP procedure. The patient lies supine on the operating table with his legs clamped open in stirrups. The surgeon is seated at the end of the operating table between the patient's legs. A scrub nurse stands to the side of the surgeon beside a trolley containing the various surgical equipment required during the course of the operation. The equipment required for the operation (electrosurgical current generator, endoscopic displays, endoscopy camera base unit and light source) resides on a trolley which is positioned on one side of the operating table. The head end of the table is occupied by the anaesthetist and the equipment required

monitor and maintain anaesthesia. Several additional nursing staff are usually present within the theatre during the course of the procedure.

The OT environment is already crowded during a conventional TURP, therefore one of the requirements for CASMS is that the equipment be as unobtrusive as possible. Access to the patient is difficult – although one side of the operating table is clear of equipment, it has to remain so to allow access to the patient by the nursing staff.

7.3.1 Optical tracking system camera

The “Flashpoint” tracking system uses an array of three infra-red cameras which view IREDs on the tracked object and allow calculation of its orientation and position. The camera array has a limited working volume within which tools can be tracked, and also requires a direct line of sight between each camera and the tracked tool. This imposes a stringent limit on the placement of the camera system in relation to the position of the tracked tools – the camera must have a view of the site of the operation.

Initial attempts were investigated using a large camera tripod holding a long boom, with the camera system bolted to the end of the boom, allowing the camera to be sited above the patient and in front of the surgeon – the only position at which a clear view of the operating site is attainable. However, the tripod requires a large footprint to maintain stability and severely restricts access to the patient from the open side of the operating table. The surgical equipment trolley on the other side of the operating table prevents siting of the tripod on that side. Therefore, the use of a tripod to hold the camera is not a viable solution.

The adopted solution was to mount the camera on a mechanical arm fixed to the ceiling of the OT. The mechanical arm is a pre-existing fixture, usually used to hold a video camera and allow filming of an operation. The arm is spring loaded to counterbalance the video camera weight (12kg) therefore an additional weight attachment was made for the Flashpoint camera system (which weighs 3kg) to balance the arm. The arm allows passive position of the Flashpoint camera system above the operation site without additional attachments, and provides a guide way for the trailing cable between the camera and the Flashpoint base unit. The Flashpoint camera, mechanical arm and counterweight are shown in figure 7.5.

One problem introduced from using the mechanical arm is that the Flashpoint



Figure 7.5: Flashpoint camera on the OT mechanical arm.



Figure 7.6: The mock-up tools for tracking the resectoscope and trigger motions in use during a TURP procedure.

camera cannot be locked in position. Therefore, the dynamic reference frame (see section 3.2.3.1, chapter 3) of the Flashpoint system must be used so that motion of the camera is not mistaken for motion of the resectoscope. The DRF must be within the working volume of the camera (alongside the tracked tools) thus introducing an additional object close to the operation site. However, even with a locked system (such as when the camera was held on the tripod and boom arrangement), the use of the DRF is a prudent precaution to detect any motion of the camera.

7.3.2 Tracked instruments

Mock-ups of the attachments for tracking the resectoscope and resectoscope trigger motion have been trialled in the OT to assess their impact on the TURP procedure. Figure 7.6 shows the mock tracking tools in use during a TURP. The resectoscope endotracker was found to give minimal interference to the surgeon's

movements. In the trial pictured in figure 7.6, the trigger tracker also did not cause problems. However, additional trials showed that the trigger tracker could prevent the full range of lateral motions during resection when the patient's legs are not widely separated. Most patients undergoing TURP are elderly, and many have hip problems which prevent the legs from opening to a large angle. At present, a successful method for trigger tracking in all cases is still an issue to be resolved.

7.3.3 Ultrasound platform

Designs for the ultrasound platform were discussed with the surgeons and a mock-up of the platform was used to assess the likely impact of the probe platform on the surgery. During TURP, the surgeon is usually seated on a stool placed between the patient's legs. The only possible location for the probe platform is in front of this stool, between the surgeon and the patient. The platform's position requires the surgeon to be seated further from the patient than during standard TURP, however, the surgeons felt that this would not be overly inconvenient.

A more pressing concern is that the platform restricts the downward movement of the resectoscope, preventing resection in the upper sections of the prostate. Most of the resection performed during a TURP occurs in the lateral lobes and base of the prostate and therefore the presence of the probe platform would not prevent TURP surgery, but would hinder a complete resection in some cases.

The mock probe platform is shown in figure 7.7.

7.3.4 Ultrasound

Four TURP operations were conducted with concurrent TRUS imaging of the prostate, with the ultrasound probe positioned within the patient's rectum and clamped rigidly to the operating table. Initial concerns that the presence of the ultrasound probe within the rectum would prevent full motion of the resectoscope were unfounded. The resectoscope causes some acoustic shadowing within the ultrasound image, but not enough to fully occlude the prostate boundary distal to the resectoscope, as shown in figure 7.8.

The major finding from the ultrasound trials is that the electrosurgical current causes severe disruption to the ultrasound imaging system during cutting and coagulation of the prostate tissue. Figure 7.9a) shows the image artefact which




Figure 7.7: Mock-up of the ultrasound probe platform positioned in relation to the surgeon during TURP.




Figure 7.8: Acoustic shadowing due to the presence of the resectoscope.




Figure 7.9: Artefacts due to the electrosurgical current during a) cutting and b) coagulation.

occurs during cutting, while figure b) shows the corresponding artefact resulting from coagulation. It is clear that the electrosurgical current during cutting causes severe interference with the ultrasound signal and effectively destroys all useful image information. During coagulation, the ultrasound image can be partially observed, however the image quality is degraded substantially.

7.3.5 OT trial outcomes

Following initial trials of the equipment in the operating theatre, it was decided that too many problems exist to enable the goal of a full-scale OT monitoring system for TURP to be achieved at present. These problems are listed in table 7.1, which is subdivided into major and minor difficulties. Table 7.1 also includes a list of potential problems which would also need to be investigated should the project be continued in the OT scenario.

The loss of ultrasound information during both cutting and coagulation of tissue means that it is not possible to obtain knowledge of the prostate boundary during these periods. This is precisely the time when knowledge of the prostate boundary is most crucial, since the likelihood of damage to the critical structures within the prostate (verumontanum and puncture of the capsule) is most likely during the cutting strokes! However, even if the cutting and coagulation did not cause artefacts on the ultrasound image, the time-delays required for processing the ultrasound images and moving the ultrasound probe mean that real-time knowledge

Major issues
Cutting current destroys ultrasound image
Coagulation current partially destroys ultrasound image
Latency of ultrasound probe platform
Latency of ultrasound image processing
Minor issues
Trigger tracker interference with full resectoscope motion
Probe platform interference with full resectoscope motion
Less comfortable surgeon posture
Potential problems
Loss of line of sight between optical tracker and tools

Table 7.1: Current difficulties for establishing a fully functional CASMS system for TURP.

of the prostate shape and position would still not be possible. These time-delays are presently too large to be solved simply by using new, faster processors.

Given the major difficulties described in the previous paragraph, no further work was conducted to reduce the minor problems described in table 7.1. The interferences to resectoscope motion could be partially reduced by redesign of the components, particularly for the trigger tracker component. In addition the potential problems of loss of sight of the IRED trackers was not investigated. These issues should be dealt with once solutions to the major problems are found.

A more detailed description of possible methods of overcoming these problem to enable a fully working OT based CASMS system for TURP is given in the next chapter, section 8.5.1. The remainder of the work in this chapter is based around a dual purpose system which extends the CASTS system and also gives a laboratory based proof-of-concept for using ultrasound as a means to track the prostate in the CASMS system.




Figure 7.10: Laboratory based CASTS/CASMS system.

7.4 Lab based TURP monitoring/training system

The purpose of the laboratory based system is two-fold in nature, namely the extension of the CASTS trainer so that registration is maintained between the computer model and the prostate phantom, and secondly, to show that ultrasound can be used in CASMS to track the prostate once the issues of timeliness and interference are resolved. Figure 7.10 shows the modified laboratory based combined CASTMS system with the patient replaced by the ultrasoundable prostate phantom described in chapter 3, section 3.4.2.

7.4.1 Overcoming the major OT problems

The problems of electrosurgical current interference on the ultrasound image and the time delays from image processing and ultrasound probe motion are not immediately solved when moving from the OT to the laboratory.

7.4.1.1 Electrical interference

Resection of the original prostate phantom² used in CASTS, required current to the cutting element to heat the loop and burn through the phantom material. However, the gelatin phantom can be resected using a wire loop electrode without heating, *i.e.* the electrosurgical current does not need to be activated for cutting to occur. Therefore, the solution to electrical interference in the ultrasound image is to use the loop element as a cutting blade without heating. Resection without heat has the additional advantage of cutting without melting the gelatin, so the resected chips do not tend to stick to the remaining material.

The adverse effect of not using the electrosurgical current for cutting the phantom is that the cuts are not necessarily detected by the laboratory CASTMS PC. The electrosurgical current is activated by pressing a foot switch; a micro-switch on the foot switch signals this activation to the PC (as described in section 3.2.3.3, chapter 3). Foot switch activation is still required in the CASTMS system during phantom resection so that the cutting strokes are signalled to the PC. However, the two events (current activation and cutting) are no longer inherently coupled, and CASTMS relies on timely activation of the foot switch by the user.

7.4.1.2 Time delays

Monitoring in the OT requires timely information of the position and shape of the prostate so that accurate feedback can be given to the surgeon. It is not possible for the surgery to be suspended while the CASMS system catches up with reality. However, this restriction does not apply for the laboratory CASTMS trial.

The main bottleneck which restricts timely updating of the computer's model of the prostate phantom is due to the slow movement of the ultrasound probe platform, which has a maximum speed of 1.44mms^{-1} . The maximum speed was deliberately made to be slow, so that movement of the probe within the patient in the OT would be constrained to be slow by design and not cause damage. A faster probe speed could be used in the laboratory trial, however, the goal of real time update would still not be achievable so no attempt was made to increase the speed. The treatable volume within the phantom is 6cm in length (from verumontanum to bladder neck), therefore the time required for a complete traverse of this region by the probe is approximately 40 seconds. This gives a

²The commercial phantom from Limbs and Things Ltd, Bristol, UK.

minimum time of at least 40 seconds for each update of the computer model of the resection region.

Time delays due to processing of the ultrasound images are minimized by processing the images during the periods of motion of the ultrasound probe platform. The ultrasound images of the prostate phantom are significantly clearer than those of real prostates, therefore a faster image processing method may be used to detect the outline. Detection of the phantom prostate outlines is discussed in section 7.4.2.

During testing, each resection stroke was preceded by a period of no resectoscope movement to allow the CASMS computer model of the phantom to catch up with the position of the real phantom. This was achieved by moving the resectoscope to the desired position for a cut and then clamping it using a chemical retort clamp stand. Once the ultrasound probe platform had made a complete traverse, the resection stroke was made by squeezing the trigger, but without changing the gross position of the resectoscope.

For the CASTS trainer, delays in the resection do not make the trainer unfeasible, however they are clearly not ideal. Therefore, further testing of the system has been performed to gauge the accuracy of monitoring a resection performed in the usual manner, *i.e.* without the waiting period prior to each resection stroke. In this trial, the ultrasound probe platform was only used in the initial phase to obtain referenced ultrasound images of the phantom. During the cutting phase the probe was clamped at the verumontanum position. The experimental procedure for both types of testing is explained in section 7.5.2.

7.4.2 Phantom outline detection

Detection of the phantom outline is performed using an optimised PDM and GA combined algorithm as described in chapter 6. The PDM was constructed from a set of outlines drawn by the author by hand on a series of ultrasound images taken at intervals along the length of the prostate phantom. Acquisition of ultrasound images was also repeated with the phantom deformed by varying degrees using a clamped resectoscope. The complete set of outlines were automatically aligned using the simulated annealing algorithm described in section 6.3, chapter 6, and the optimised PDM obtained by using the algorithm of Davies et al [59] as described in section 6.5, chapter 6.

The phantom's PDM is much more compact than the model obtained from out-




Figure 7.11: Phantom ultrasound image after preprocessing and outline detection.

lines of real prostates, since variation in outline shape is much less. The phantom outlines are approximately circular or elliptical depending on the degree of deformation/movement. Following alignment, the outlines match each other very closely so that the final PDM is a highly compact model.

The compactness of the PDM means that the searching performed by the GA is predominantly to find the best pose parameters to fit the outline to the given ultrasound image. An additional advantage stems from the fact that the ultrasound image of the phantom is relatively clearer in comparison to ultrasound images of real prostates. The contribution of both these factors allows the GA to find an optimal solution using a smaller population of outlines (100 instead of 400), and convergence after fewer iterations (500 instead of 2500) when compared with the GA for finding outlines of prostates in real ultrasound images.

A further gain in terms of algorithm speed is achieved by using a simple, fast averaging preprocessing filter for the phantom images, rather than the computationally expensive speckle reduction algorithm which gives the best result for the real prostate ultrasound images. The final algorithm takes approximately 5 seconds to perform the preprocessing and genetic algorithm optimization on the PIII dual processor machine³ which is the PC for CASTMS. The gains in execution speed are such that computing the position of the outline in each image takes less than the time required to move the ultrasound probe to the next acquisition position (section 7.4.1.2). Figure 7.11 shows an ultrasound image of the phantom following smoothing with a spatial averaging filter, which is overlaid with the PDM/GA algorithm outline result.

³The outline detection algorithm runs in a single thread so can only use one of the PC's processors.




Figure 7.12: IRED tracker for the TRUS probe.

7.4.3 Modelling the treatment volume

Once a set of outlines is found from a series of ultrasound images taken at intervals through the phantom, the next step is to construct a 3D model of the phantom from these outlines. The method of achieving this goal is described in the proceeding sections. The modelling algorithm has been constructed with twin goals of fast construction of the 3D model, and the ability for rapid interrogation of the model to assess collision detection. Although speed is not a critical issue for the lab-based trial, the objectives were set to allow the same algorithm to be used in an in-vivo CASMS system.

Other algorithms are documented in the literature for performing the same task – generating a 3D mesh from a series of points [137, 139, 231]. However, no attempt has been made to compare the timeliness of other algorithms with the mesh generation algorithm presented in the proceeding sections.

7.4.3.1 Coordinate transformations

The first step in generating the model is to convert pixel points on the ultrasound image to points in a 3D coordinate system relative to the resectoscope. These transformations have been broken down into a series of steps. The first step relates the pixel position within the image to a co-ordinate system defined around the ultrasound probe, which introduces physical dimensions. The position and orientation of the TRUS probe is initially captured using an IRED tracker attached to the probe as shown in figure 7.12. Once the resection is underway, the resectoscope obscures the TRUS tracker, so the position of the ultrasound probe is calculated by counting the number of steps completed by the stepper motor, and computing the current position using the initial position and orientation. The second step defines the relationship between the ultrasound co-ordinate system




Figure 7.13: a) Image and b) ultrasound coordinate systems.

and the tracking system (world) co-ordinates. The tracking system also monitors the position of the resectoscope, which is used to initially pinpoint certain anatomical features of the phantom (or patient's) prostate, allowing definition of a co-ordinate system for the boundary. Thus the transformation from the world co-ordinate system to boundary co-ordinate system may be defined. Combining all the transformations achieves the goal of relating an outline on the ultrasound image to the position of the phantom.

7.4.3.1.1 Image to ultrasound coordinates The first requirement is to transform positions from the image co-ordinate system (shown in figure 7.13a) given in terms of pixels, to the ultrasound probe co-ordinate system (figure 7.13b) in real-world units.

The transformation requires knowledge of the probe centre within the image, and the scales in r and c directions respectively (in pixels per millimetre). One method for determining these values would be to measure them once on the ultrasound machine at a particular zoom setting, and then use the measured values in CASMS. However, this constrains the user to always select the same zoom factor when using the system; incorrect setting would result in a boundary model which did not match the true size and position of the treatment volume within the prostate.




Figure 7.14: World and ultrasound coordinate systems

For the B&K ultrasound system used in the lab trial, the position of the probe centre does not change with variations in the zoom setting. Therefore, the probe centre parameter has been manually set in the CASMS software. For ultrasound machines in which the probe position in the image varies with the zoom, the centre can be found rapidly by using a Hough transform [103] to identify the most likely position.⁴ The pixel per millimetre scale factors are found automatically by detecting the scale markings on the left and bottom of the image.

The position of a pixel $P(r, c)$ transforms to the ultrasound probe coordinate position $P(x_{US}, y_{US}, z_{US})$ according to:

$$x_{US} = (c - c_o)c_{scale}, \quad y_{US} = (r_0 - r)r_{scale}, \quad z_{US} = 0 \quad (7.1)$$

where (r_0, c_o) is the probe centre in the image and r_{scale}, c_{scale} are the respective pixel per millimetre scales.

7.4.3.1.2 Ultrasound coordinate system Figure 7.14 shows the relationship between the ultrasound coordinate system and the world coordinate system used by the Flashpoint tracking system. The Flashpoint system initially captures the ultrasound probe position from the IRED markers attached to the TRUS probe, and calculates the unit vectors corresponding to the long axis of the probe

⁴The Kretz 420 machine, which was used to capture the prostate image data set described in chapter 5, is an example of an ultrasound system in which the probe position in the image is governed by the zoom setting of the machine.

($\hat{\mathbf{u}}_{probe}$) and a perpendicular axis which corresponds to the vertical axis from the transducer, and the position of the probe tip (S_{World_0})⁵. As the probe moves on the probe platform, these unit vectors remain the same, but the tip position varies according to:

$$S_{World} = S_{World_0} + N_{steps} M_{scale} \hat{\mathbf{u}}_{probe} \quad (7.2)$$

where N_{steps} is the number of steps moved by the probe platform's stepper motor, and M_{scale} is the corresponding millimeters per step parameter of the probe platform.

CASMS uses these vectors to determine the position and orientation (the cosines of the angles) of the probe with respect to the world co-ordinates. The ultrasound co-ordinate system is defined such that the zero-point (i.e. the monitored position where $[x_{World}, y_{World}, z_{World}, \theta_{x_{World}}, \theta_{y_{World}}, \theta_{z_{World}}] = [0, 0, 0, 0, 0, 0]$) corresponds to the long axis of the probe lying along the positive z_{World} axis, with the transducer in the plane $z_{World} = 0$, and the x_{US} and y_{US} axes aligned with the x_{World} and y_{World} axes respectively. $\theta_{x_{World}}, \theta_{y_{World}}, \theta_{z_{World}}$ are the angles of rotation around each of the world axes, defined as positive clockwise when looking from the origin along the respective positive axes. For a point $P(x_{US}, y_{US}, z_{US})$ in ultrasound co-ordinates, the corresponding position in the world co-ordinates system is found by a series of rotations (R) and translations (T), given the probe tip origin $S_{World} = [S_{x_{World}}, S_{y_{World}}, S_{z_{World}}]$:

Rotate around z_{World} axis:

$$\begin{bmatrix} \cos(\theta_{z_{World}}) & -\sin(\theta_{z_{World}}) & 0 & 0 \\ \sin(\theta_{z_{World}}) & \cos(\theta_{z_{World}}) & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} x_{US} \\ y_{US} \\ z_{US} \\ 1 \end{bmatrix} = R(\theta_{z_{World}}) P \quad (7.3)$$

Rotate around x_{World} axis:

$$\begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & \cos(\theta_{x_{World}}) & -\sin(\theta_{x_{World}}) & 0 \\ 0 & \sin(\theta_{x_{World}}) & \cos(\theta_{x_{World}}) & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix} R(\theta_{z_{World}}) P = R(\theta_{x_{World}}) R(\theta_{z_{World}}) P \quad (7.4)$$

⁵The ultrasound image is formed from mechanically rotating transducer with the TRUS probe. The "tip" position refers to the centre of rotation of the transducer.

Rotate around y_{World} axis:

$$\begin{bmatrix} \cos(\theta_{y_{World}}) & 0 & \sin(\theta_{y_{World}}) & 0 \\ 0 & 1 & 0 & 0 \\ -\sin(\theta_{y_{World}}) & 0 & \cos(\theta_{y_{World}}) & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix} R(\theta_{x_{World}}) R(\theta_{z_{World}}) P \dots \quad (7.5)$$

$$\dots = R(\theta_{y_{World}}) R(\theta_{x_{World}}) R(\theta_{z_{World}}) P$$

Translate to the scan position:

$$\begin{bmatrix} 1 & 0 & 0 & S_{x_{World}} \\ 0 & 1 & 0 & S_{y_{World}} \\ 0 & 0 & 1 & S_{z_{World}} \\ 0 & 0 & 0 & 1 \end{bmatrix} R(\theta_{y_{World}}) R(\theta_{x_{World}}) R(\theta_{z_{World}}) P \dots \quad (7.6)$$

$$\dots = T(S_{World}) R(\theta_{y_{World}}) R(\theta_{x_{World}}) R(\theta_{z_{World}}) P$$

The world and ultrasound co-ordinate systems are aligned, therefore the position of P in the world co-ordinate system is:

$$P_{World} = T(S_{World}) R(\theta_{y_{World}}) R(\theta_{x_{World}}) R(\theta_{z_{World}}) P \quad (7.7)$$

7.4.3.1.3 Boundary coordinate system Figure 7.15 shows the co-ordinate system used to define the boundary and its relationship to the world co-ordinate system. The boundary co-ordinate system is defined such that the origin lies at the verumontanum (the lower bound for resection) with the z_{Gland} axis lying along the urethra the straight line joining verumontanum to bladder neck. The planes orthogonal to the z_{Gland} axis are defined using a polar co-ordinate system, with r_{Gland} defining the distance from the z_{Gland} axis in each plane, and ϑ_{Gland} the angle around the axis. The line $\vartheta_{Gland} = 0$ is arbitrarily defined such that transforming the boundary co-ordinate system by translating the origin (verumontanum) to the world origin, followed by rotating around the x_{World} and y_{World} axes to align the z_{Gland} axis with the z_{World} axis, causes the line $\vartheta_{Gland} = 0$ to lie along the x_{World} axis (assuming no rotation around the z_{World} axis).

The positions of the verumontanum and bladder neck must be defined in terms of the world co-ordinate system. At present, this is achieved initially by asking the user to mark the position of the verumontanum and four points on the bladder




Figure 7.15: Relationship between the prostate boundary, ultrasound and world coordinate systems. The prostate (blue dashed outlines) is shown elongated to aid clarity.

neck (topmost, leftmost, bottommost and rightmost) using the resectoscope as a pointing device. The world co-ordinates of these points are recorded using the Flashpoint optical tracker. The bladder position is calculated as the average position of the four defined points.

These initial positions for the verumontanum and bladder neck also provide the limits for the motion of the ultrasound probe platform, *i.e.* the traverse of the ultrasound probe starts at the position where the verumontanum point lies within the ultrasound image, and finishes once the bladder neck point lies within the image plane. The end stops for the ultrasound probe remain constant throughout the course of the resection. However, the positions of the verumontanum and bladder neck are updated with each traverse of the ultrasound probe. They are calculated from the mean of a set of points on the detected outline in the respective ultrasound scans so that the gland coordinate system moves as the phantom (or real prostate) moves.

Considering a point $P(x_{World}, y_{World}, z_{World})$ defined in the world co-ordinate system, the corresponding point in the boundary co-ordinate system is found as follows, given the verumontanum position $V_{World}(V_{x_{World}}, V_{y_{World}}, V_{z_{World}})$ and bladder neck position $B_{World}(B_{x_{World}}, B_{y_{World}}, B_{z_{World}})$:

Translate verumontanum back to origin:

$$\begin{bmatrix} 1 & 0 & 0 & -V_{x_{World}} \\ 0 & 1 & 0 & -V_{y_{World}} \\ 0 & 0 & 1 & -V_{z_{World}} \\ 0 & 0 & 0 & 1 \end{bmatrix} P = T(V_{World}) P \quad (7.8)$$

The angle needed to rotate around the x_{World} axis to bring the z_{Gland} axis to lie on the x_{World}, z_{World} plane is calculated as:

$$\phi_{x_{World}} = \text{atan2}(B_{y_{World}} - V_{y_{World}}, B_{z_{World}} - V_{z_{World}}) \quad (7.9)$$

where $\text{atan2}(y, x)$ returns the arctangent of $\frac{y}{x}$ in the range $-\pi$ to $+\pi$ radians using the signs of y and x to determine the quadrant.

Rotate around the x_{World} axis:

$$\begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & \cos(\phi_{x_{World}}) & -\sin(\phi_{x_{World}}) & 0 \\ 0 & \sin(\phi_{x_{World}}) & \cos(\phi_{x_{World}}) & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix} T(V_{World}) P \dots \quad (7.10)$$

$$\dots = R(\phi_{x_{World}}) T(V_{World}) P$$

The angle needed to rotate around the y_{World} axis to bring the z_{Gland} axis to lie along the z_{World} axis is calculated as:

$$\begin{aligned} \phi_{y_{World}} &= -\text{atan2}(B_{x_{World}} - V_{x_{World}}, \dots) \\ \dots &= (B_{y_{World}} - V_{y_{World}}) \sin(\phi_{x_{World}}) + (B_{z_{World}} - V_{z_{World}}) \cos(\phi_{x_{World}}) \end{aligned} \quad (7.11)$$

Finally, rotate around y_{World} axis :

$$\begin{bmatrix} \cos(\phi_{y_{World}}) & 0 & \sin(\phi_{y_{World}}) & 0 \\ 0 & 1 & 0 & 0 \\ -\sin(\phi_{y_{World}}) & 0 & \cos(\phi_{y_{World}}) & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix} R(\phi_{x_{World}}) T(V_{World}) P \dots \quad (7.12)$$

$$\dots = R(\phi_{y_{World}}) R(\phi_{x_{World}}) T(V_{World}) P$$

Therefore, P_{Gland} (the point P in the boundary co-ordinate system) is defined as:

$$P_{Gland} = \begin{bmatrix} x_{Gland} \\ y_{Gland} \\ z_{Gland} \end{bmatrix} = R(\phi_{y_{World}}) R(\phi_{x_{World}}) T(V_{World}) P \quad (7.13)$$

and the r_{Gland} and ϑ_{Gland} coordinates are found as :

$$r_{Gland} = \sqrt{x_{Gland}^2 + y_{Gland}^2} \quad (7.14)$$

$$\vartheta = \text{atan2}(y_{Gland}, x_{Gland}) \text{ radians} \quad (7.15)$$

7.4.3.1.4 Outline contour transformation The transformation equations 7.1, 7.7 and 7.13 can be combined to transform the position of a point in the image

to the boundary co-ordinate system. The full transformation is:

$$P_{Gland} = \begin{bmatrix} x_{Gland} \\ y_{Gland} \\ z_{Gland} \\ 1 \end{bmatrix} = R(\phi_{y_{World}}) R(\phi_{x_{World}}) T(V_{World}) T(S_{World}) \cdots \quad (7.16)$$

$$\cdots R(\theta_{y_{World}}) R(\theta_{x_{World}}) R(\theta_{z_{World}}) \begin{bmatrix} (c - c_0) x_{scale} \\ (r_0 - r) r_{scale} \\ 0 \\ 1 \end{bmatrix}$$

with the boundary polar co-ordinates defined according to equations 7.14 and 7.15. However, the conversion of each outline to a polar coordinate system (as discussed in section 7.4.3.1.5) requires knowledge of the position of the urethra in each image, which is the solution for x_{US} , y_{US} and z_{Gland} given that $x_{Gland} = y_{Gland} = z_{US} = 0$.

Re-writing equation 7.16 as:

$$\begin{bmatrix} x_{Gland} \\ y_{Gland} \\ z_{Gland} \\ 1 \end{bmatrix} = \begin{bmatrix} A & B & C & D \\ E & F & G & H \\ I & J & K & L \\ M & N & O & P \end{bmatrix} \begin{bmatrix} x_{US} \\ y_{US} \\ z_{US} \\ 1 \end{bmatrix}$$

multiplying out, re-arranging and inserting known values gives:

$$x_{US} = U_{x_{US}} = \frac{DF - BH}{BE - AF}$$

$$y_{US} = U_{y_{US}} = \frac{AH - DE}{BE - AF} \quad (7.17)$$

$$z_{Gland} = U_{z_{Gland}} = I \frac{DF - BH}{BE - AF} + J \frac{AH - DE}{BE - AF}$$

where U is the position of the urethra in the image, and the distance along the z_{Gland} axis.

The transformations are adjusted so that they relate the position in the image with respect to the verumontanum position, to the boundary co-ordinate system. This is synonymous with moving the ultrasound probe position to the position




Figure 7.16: a) Outline which cannot be represented in polar form. b) Polar approximation of the outline.

of the urethra, i.e.

$$S_{x_{World}} = V_{x_{World}} + U_{z_{Gland}} \frac{(B_{x_{World}} - V_{x_{World}})}{L_{Gland}} \quad (7.18)$$

$$S_{y_{World}} = V_{y_{World}} + U_{z_{Gland}} \frac{(B_{y_{World}} - V_{y_{World}})}{L_{Gland}} \quad (7.19)$$

$$S_{z_{World}} = V_{z_{World}} + U_{z_{Gland}} \frac{(B_{z_{World}} - V_{z_{World}})}{L_{Gland}} \quad (7.20)$$

where L_{Gland} is the distance between the verumontanum and bladder neck:

$$L_{Gland} = \sqrt{(B_{x_{World}} - V_{x_{World}})^2 + (B_{y_{World}} - V_{y_{World}})^2 + (B_{z_{World}} - V_{z_{World}})^2} \quad (7.21)$$

7.4.3.1.5 Polar outline transformation The mesh generation algorithm described in section 7.4.3.2.2 requires a polar representation of the prostate outlines from section 7.4.2. The origin of the polar coordinate system is defined as the position of the urethra within each scan. The use of a polar representation $r = f(\vartheta)$ presents certain limitations in that each value of ϑ has a single radius associated with it. Figure 7.16a) shows an outline which cannot be represented in polar form, hence cannot be used in CASMS. However, inspection of the prostate ultrasound images and expert drawn outlines shows that very few are not capable of exact representation in polar form. When approximation is necessary as in figure 7.16b), the polar representation is defined such that the modified outline




Figure 7.17: a) Enlarged median lobe overhanging bladder neck. b) Ultrasound image through the median lobe.

does not lie outside the original outline.

A further limitation of the polar co-ordinate representation is that outlines which do not encircle the urethra, may not be defined. Such an occurrence can occur in prostates where the median lobe is so enlarged that it extends beyond the bladder neck as in figure 7.17. Whilst not a problem for the laboratory test of CASMS on the prostate phantom, in the real CASMS system it may present difficulties. However, it is likely that the volume of treatable tissue within such an overhang is small, and that with proper knowledge of CASMS, the surgeon will know when to disregard a proximity warning given in this region.

7.4.3.2 Boundary generation

Once outlining is completed for all the acquired TRUS scans, CASMS develops a 3D model of the treatment volume within the prostate. The points on each outline are gathered and transformed to the boundary cylindrical co-ordinate system. Thus, each point (interpolated from the complete outlines) has co-ordinates (z, r, ϑ) . Additional points, with zero radius and at incremental ϑ , are added at both ends of the boundary, to ensure that interpolation at the boundary ends is performed between the most appropriate points.

7.4.3.2.1 Data space transformation The objective of the boundary generation procedure is to develop a mathematical model of the boundary, such that

the radius can be calculated for any (z, ϑ) . This corresponds to a two dimensional interpolation, between three defined points on at least two TRUS scans. Various schemes exist for performing such an interpolation, but the best way is to form triangles across the data space, and interpolate a point according to the vertices of the triangle that it falls inside. The problem with the interpolation is that the data points are defined with one linear (z_{Gland}) and one angular (ϑ_{Gland}) dimension, implying a complicated calculation involving the radius of each point in order to compare relative distances. CASMS overcomes this problem by transforming the angular dimension to a linear one (ω), by multiplying the angle of each point by a constant. The transformation corresponds to unrolling the cylindrical co-ordinate system onto a 2D plane. The assumption is that the radius does not vary significantly between points. The transformation constant is chosen so that the range of the dimension ω is the same as the dimension z_{Gland} , i.e. the data-space is square.

7.4.3.2.2 Triangulation The data points now lie on a 2D plane. A Voronoi tessellation [218] of the data-space is obtained by drawing polygons around each point such that the area within the polygon is closest to that point. Each side of each polygon bisects the line joining the two data points whose polygons are bounded by that side. Therefore, each vertex of the polygons is the centre of a circle which passes through the three nearest data points. The triangles formed by the data points are the basis for interpolation for any point on the boundary. This triangular mesh is known as a Delaunay triangulation [62]. The criterion for the mesh is that no data-point falls within the inscribing circle of any triangle. Figure 7.18 shows a sample data-space with Voronoi tessellation and Delaunay triangulation.

Constructing the mesh is a stepwise process of introducing a new point, checking all triangles to find those where the new point falls within the inscribing circle, deleting the found triangles, and reforming the triangles with the new point. The usual initial starting point is to create one triangle from three extra points that encompasses the whole data-space; two triangles (and four points) are used in CASTMS since the data-space is square. The initial triangles are shown in figure 7.19.

The software implementation of the triangulation routine contains an extra criterion beyond a simple Delaunay triangulation of the data points. As each point is added to the mesh, an additional check is made to ensure that new triangles




Figure 7.18: a) Data points. b) Areas closest to each data point (Voronoi tessellation). c) Inscribed circles inside which no point may lie. d) Delaunay triangles.




Figure 7.19: Initial triangles formed by four extra points (blue) completely surround the true data points (red).

are constructed from points which derive from at least two different ultrasound images. The check is necessary because points from the same ultrasound image will only lie on a straight line (and therefore not produce a triangle) in the data space transformation if the z_{Gland} and z_{US} axes happen to be parallel. The additional check extends the calculation time required to construct the mesh, however, the execution speed is still sufficiently rapid for a mesh of 600 data points⁶ to be performed in under a second on the CASTMS PC (using a single 1GHz processor).

7.4.3.2.3 Radius interpolation The purpose of constructing the triangle mesh is to interpolate the radius at any point from the vertices of the triangle around the point. Thus, the first task is to find the triangle that contains the point in question. A simple method for performing this function is to sum the angles between lines joining the point to the triangle's vertices. However, such a method requires trigonometric calculations, and is therefore too slow, (since it must be applied many times). Instead, a line from the point to infinity in any arbitrary direction is drawn. This line will cross the containing triangle once, and twice or not at all for any other triangle. There are two possibilities for error (since the computer has a limit of precision). The first is caught by an explicit check for the point lying on a vertex - the radius of the vertex is returned. If the point lies on or near a triangle edge, then the lack of precision could possibly cause the wrong triangle or no triangles to be found. Wrong triangles are rejected by performing a second edge-crossing test in the orthogonal direction. Extensive testing of the algorithm has shown that this mechanism is adequate to prevent any false finds or search failures.

Once the surrounding triangle is found, calculating the radius is a simple interpolation. The triangle lies in the plane defined by its vertices (x_1, y_1, r_1) , (x_2, y_2, r_2) , (x_3, y_3, r_3) such that:

$$r = x \frac{dr}{dx} + y \frac{dr}{dy} c$$

which may be solved by substituting for the known values at the vertices.

The software implementation of the radius interpolation routine is important since the position of the resectoscope's cutting loop must be constantly checked against the mesh to detect proximity to the resection limits. The triangle search method is accelerated by starting each new search from the previously found triangle.

⁶The 600 data points are formed from 100 interpolated outline points on each of six ultrasound images at 10mm spacings through the resectable region in the phantom prostate.

Thus extensive searching of the mesh is only required when the resectoscope moves from one side of the z_{Gland} axis to the other (during which the cutting loop cannot be close to the resection limits, so extended searching time is not a problem). Testing of the software implementation shows that radius interpolation can be performed at rates of 50Hz using a mesh containing 600 data points and running on the CASTMS PC.

The radius interpolation routine does not guarantee that the returned point is necessarily the closest point on the mesh to the cutting loop, *i.e.* exact collision detection between the resection limits and the cutting loop is not performed. However, for smoothly varying shapes such as the prostate (both the real and phantom varieties) this only presents a problem when the cutting loop is at either end of the resection region (near to the verumontanum or bladder neck). Additional checks are made to find the true closest point when the cutting loop gets close to these ends.

Fast explicit collision detection routines which always find the closest point on a surface do exist, such as the method proposed by Starkie [200], where a NURBS surface is used as the collision boundary. However, direct generation of a NURBS surface from a cloud of points is not a simple task. For the purposes of CASTMS, boundary proximity rather than collision detection is the required knowledge, and the current interpolation routine is considered to be sufficient.

7.4.4 Software architecture

The software for CASMS has been created according to the principles of good object-oriented design [133] using the Unified Modelling Language [73] approach to analyze the problem and assess possible usage scenarios prior to coding [81]. The software classes have been written with strong encapsulation of the required tasks, and provide a level of abstraction to the underlying hardware. The result is a highly modular structure within which many individual components can be re-used in subsequent minimally invasive training or monitoring systems. An overview of the software is given in figure 7.20, which shows the class hierarchy, inter-class relationships and the number of instances corresponding to each relationship (one to one or one to many⁷).

The PC used for the laboratory CASMS system is a Dell⁸ Precision 220 workstation with dual 1GHz Pentium III processors running the MicrosoftTM Windows

⁷One to many relationships are denoted by an asterisk on the multiple instance class.

⁸Dell Computer Corporation, Bracknell, Berkshire, UK.




Figure 7.20: Software class hierarchy.

2000 operating system. The additional components for the PC are the DT302 data acquisition card for detecting foot switch activation (described in chapter 3, section 3.2.3.3) and a Matrox Meteor-II frame grabber card for capturing video signals from the ultrasound system. Unfortunately, both these cards are only supplied with drivers for the Microsoft Windows family of operating systems, so that the choice of operating system for the CASTMS PC is restricted to these. This does not pose a significant problem for the laboratory CASTMS setup, however, the preemptive and closed-source nature of the Windows operating system means that it is unsuitable for use in a real time application such as a complete CASMS system for use in the OT. The author advocates the use of alternative operating systems, such as one of the Linux distributions, as a more suitable alternative to Windows since they are open-source, capable of true real-time performance and highly configurable to the required task.

To enable progress towards adoption of an alternative operating system, the software has been designed for cross-platform compatibility where possible. In particular, the design provides a clear boundary between the graphical user interface (GUI) and the remainder of the software. The GUI itself has been designed using the freely available FLTK [71] cross platform library, which provides a fast toolkit without the unnecessary overheads of the native Windows GUI framework. Instances where specific Windows instructions are necessary, such as access to the boards and ports, have been completely encapsulated to enable a change in operating system with minimal rewriting of code.

The following sections describe the major classes within the CASTMS software design.

7.4.4.1 Frame Factory class

The frame factory class controls the automated aspects of the CASTMS system. Figure 7.21 shows the flow diagram for the frame factory's active thread. The frame factory generates requests for movement of the ultrasound probe and passes them to the motion controller class (section 7.4.4.4). Once completion of each movement is signalled by the motion controller, the frame factory collects an ultrasound image from the grabber class (section 7.4.4.2), and the current ultrasound probe location from the motion controller, which are combined in a frame class. Once processing of the previous frame⁹ into an ultrasound slice ob-

⁹The frame contains the ultrasound image and the location of the probe when the image was captured.




Figure 7.21: Frame Factory thread flow diagram.

ject is completed by the frame processor class (section 7.4.4.5), this ultrasound slice is retrieved and the new frame passed to the frame processor. The processed ultrasound slice is added to a set of slices for the current traverse of the ultrasound probe. Once the ultrasound slice set is complete, it is passed to the gland processor (section 7.4.4.7) and the probe motion changes direction.

7.4.4.2 Grabber class

The grabber class provides an abstracted interface to the low-level driver for the Matrox¹⁰ Meteor-II frame grabber card. The grabber class captures image data from the card and performs some preliminary calculations including establishing the scale setting and probe location in the ultrasound image.

7.4.4.3 3D Position class

The 3D position class provides information on the positions and status of the resectoscope, cutting loop and foot pedals. The class provides an abstraction layer to the flashpoint class which provide access the Flashpoint hardware through a serial port class. The 3D position class has an in built timer which triggers requests to update the positional information. On each tick, the current position of the resectoscope and cutting loop are acquired from the flashpoint class, and the status of each foot pedal is checked to detect cutting and coagulation actions.

7.4.4.4 Motion Controller class

The motion controller class is an abstraction layer for interfacing to the underlying hardware which controls the motion of the ultrasound probe platform. The class contains a list of requested moves which are carried out in sequence; completion of each move is signalled to the requesting class through a *functor*¹¹ so that the class is encapsulated according to the set design criteria.

The motion controller class contains a timer which triggers at intervals of 25 milliseconds¹². On each tick, if a move is requested or already underway, a pulse

¹⁰Matrox Electronic Systems Ltd, Dorval, Quebec, Canada.

¹¹A functor is an object which may be called as though it is a function. In the motion controller class, calling the functor releases a semaphore in the Frame Factory class, although the Motion Controller does not (or need to) know this.

¹²The timer interval is adjustable and alters the velocity of the stepper motor. The current value of 25ms is the minimum interval which still ensures correct activation of the stepper motor, without loss of steps.

is sent through a pin on the parallel port to the stepper motor controller hardware, and the ultrasound probe platform moves through one step. The desired direction is controlled by changing the state of a second parallel port pin. Control of the parallel port is abstracted in a separate class which uses a freeware utility to access the parallel port under Microsoft WindowsTM.¹³

The flow of the motion controller class is presented in figure 7.22.

7.4.4.5 Frame Processor class

The frame processor class receives grabbed frames from the frame factory and uses the genetic algorithm/PDM method described in section 7.4.2 to find the boundary of the prostate/prostate phantom. The prostate boundary information is stored in an outline class, which together with the frame comprise an ultrasound slice class. Completion of processing is signalled by calling a functor which releases a semaphore in the frame factory class. Figure 7.23 shows the flow of the frame processor thread

7.4.4.6 Ultrasound Slice class

The ultrasound slice class is an abstract parent class which allows other classes access to the ultrasound image and outline information. The class is split into two child classes, a real ultrasound slice class which is generated by the frame processor and contains the functionality for processing the image and creating the outline information, and a proxy ultrasound class which allows other classes access to the data but without the ability to change it, such as the gland processor and GUI classes (sections 7.4.4.7 and 7.4.4.10). The real/proxy divide ensures that only one copy of the ultrasound slice information is stored, and that it remains valid while other classes are accessing it. Once all the proxies are released, the real ultrasound slice data is deleted.

7.4.4.7 Gland Processor class

The gland processor class receives a set of processed ultrasound slices from the frame factory, and uses the detected outlines to generate a 3D mesh of the allowable operating zone within the prostate. The outlines in each frame are converted to polar form, then interpolated to a set of points which are transformed to the

¹³The parallel port driver is freely available from www.mtnmicro.com/resources.




Figure 7.22: Motion Controller flow diagram.




Figure 7.23: Frame Processor and Gland Processor threads flow diagram.

gland coordinate system, as described in section 7.4.3. The resulting points are triangulated to form a gland object (also described in section 7.4.3).

7.4.4.8 Gland class

The gland class is an abstract parent class allowing other classes access to the current 3D model of the resectable region within the prostate. The class follows a similar structure to the ultrasound slice class with a real gland class which is generated by the gland processor containing the functionality for generating the model, and a proxy gland class which allows for interrogation of the model by other classes, such as the gui class and the operation class, but without the ability to change the underlying gland model. The real/proxy divide ensures that only one copy of the model information is stored, and that it remains valid while other classes are accessing it. Once all the proxies are released, the real gland data is deleted. Since the gland data is continuously updated, the real/proxy structure provides an excellent method for allowing many classes access to the same information without unnecessary duplication of the data.

7.4.4.9 Operation class

The operation class provides the monitoring aspect of the CASTMS system. The class contains an embedded timer which provides timed requests for resectoscope status information to the 3D position class. During resection, the operation class interrogates the current gland class model to assess the proximity of the resectoscope to the model boundary. Resection strokes are captured and stored and warnings are generated when the cutting loop approaches the resection limits. The flow of the operation class thread is shown in figure 7.24.

7.4.4.10 GUI class

The gui (graphical user interface) class provides the visual display to the user and provides the various buttons and dialogs to allow set-up of the CASTMS system. The gui class contains a timer which causes refresh of the display so that the position and status of the resectoscope, the current gland 3D model, and the progress of the resection are shown at interactive rates. The user interface is similar to that used in the CASTS trainer (chapter 3, section 3.2.4.2) except that the planes through the gland model are superimposed onto the ultrasound image




Figure 7.24: Operation thread flow diagram.




Figure 7.25: The GUI of the CASTMS software.

captured nearest to the current position of the cutting loop (although these were not displayed when the screen shot of figure 7.25 was captured).

The GUI is shown in figure 7.25. The left side of the GUI shows the resectoscope (in green since cutting is not being performed) overlaid on the modelled treatment volume (the red, rendered portion) of the prostate phantom. The blue/black wire frame regions of the phantom are static images which convey the overall phantom shape, but are not updated as the treatment volume model changes. The horizontal lines show the locations of the ultrasound scans, with the blue line corresponding to the closest ultrasound scan to the position of resectoscope cutting loop.

The right hand side of figure 7.25 shows the ultrasound image taken at the location of the blue line. The green curve shows the detected phantom outline, and the red curve shows the periphery of the resection cavity. The yellow cross-hairs show the position of the cutting loop in the plane of the ultrasound image.

7.4.5 Modelling the resection cavity

As the resection progresses, CASTMS uses the positions of the resectoscope and cutting loop during each cutting stroke and the current model of the gland to calculate and store the resection cuts in terms of the gland coordinate system (described in section 7.4.3.1.3). These resection cuts appear as circles in the cutting planes through the gland which are shown as part of the GUI.

In a real prostate, there is no (or at least very little) cavity within the prostate prior to resection. However, the phantom does have a significant cavity already




Figure 7.26: CASTMS physical setup.

before any cuts have been made which must be included in the resection cavity data.

An additional image processing phase is performed to detect the initial cavity during generation of the first gland model. As each ultrasound image is processed, the region of the image which lies within the detected outline is thresholded to produce a binary image of the cavity. Points on the periphery of the cavity are then added to the list of resection cuts.

The cut position data is used to create the resection cavity presented in figure 7.36, on page 285.

7.5 CASTMS system testing

Figure 7.26 shows the CASTMS physical setup. Apart from the addition of the ultrasound probe platform, the main change from the CASTS system is the modification of the mock abdomen (perspex box) so that it may be completely filled with water during the resection. The water ensures a good conduction medium for the ultrasound imaging. Flexible latex ports on the front of the abdomen allow entry and movement of the ultrasound probe (lower port) and



Figure 7.27: Gelatin phantom set into perspex base.

the resectoscope (upper port) without leakage.

7.5.1 CASTMS setup

Figure 7.27 shows a gelatin phantom prior to resection in the CASTMS system. The phantom is set onto a perspex plate which screws into the base of the mock abdomen and anchors the phantom. The metal rod shown in figure 7.27 is a former for the hole to allow ultrasound access, and is removed prior to use.

Once the phantom is placed inside the abdomen, the ultrasound probe and resectoscope are inserted through the latex ports. The resectoscope port contains an o-ring which clamps onto the resectoscope shaft and achieves a watertight seal whilst still allowing the resectoscope to slide in and out. Sealing of the ultrasound probe port is achieved with a cable tie. The final step in the physical setup is to fill the mock abdomen with water.

On initialization, the CASTMS software proceeds through a series of internal checks to ensure that correct handshaking with the Flashpoint system, the frame grabber, the data acquisition card and the stepper motor controller are achieved. The stepper motor is activated so that the ultrasound probe platform moves to its home position. CASTMS acquires the position of the ultrasound platform home position and the vector of the axis of travel using the Flashpoint IRED probe tracker (see figure 7.26). The resectoscope must be held to one side during this process, to allow direct line-of-sight between the IRED markers and the Flashpoint camera array (the resectoscope obscures the probe tracker during use as described in section 7.4.3).

The final task in the setup is that the user is prompted to pinpoint the positions of




Figure 7.28: CASTMS in use in the laboratory.

the verumontanum and bladder neck using the resectoscope as a pointing device, and select the number of scans to be acquired on each pass of the ultrasound probe through the phantom. The ultrasound probe then moves to the verumontanum position and begins the first traverse. CASTMS is ready for resection to proceed once the first traverse is completed.

Figure 7.28 shows the CASTMS system during resection of a phantom.

7.5.2 Experimental procedure

Two experiments were conducted using the CASTMS laboratory setup, with the first being a test of using ultrasound tracking as a method of improving the registration of the CASTS system, and the second an assessment of the potential accuracy of ultrasound tracking. The main objective of the first experiment was to assess the performance of the system without inflicting changes to the operation protocol on the user. For the second experiment, this criterion was lifted – any change to the normal procedure of resecting the prostate (phantom) was deemed permissible.

7.5.2.1 Static probe experiment

Following setup of the CASTMS equipment and software, the first experiment was conducted with the ultrasound probe platform clamped at the verumontanum,

i.e. the position at which the specified position of the verumontanum coincide with the image plane. The verumontanum position was chosen since this portion of the resectable region undergoes the most movement during the course of the phantom's resection.

The ultrasound images were processed using the same outline detection algorithm described in section 7.4.2. However, only the centroid of the detected outline was used to update the computer model of the prostate phantom. Following outline detection, the verumontanum position in the gland coordinate system was updated to be the calculated outline centroid, and the z_{Gland} axis adjusted accordingly. Thus only the position of the phantom model changes without any deformation.

The resection was conducted in the same manner as the CASTS system, with the exception that resection of the upper area of the phantom near to the verumontanum was hindered by the presence of the ultrasound probe.

7.5.2.2 Clamped resectoscope experiment

Following setup of the CASTMS equipment and software, a partial resection of the prostate phantom was performed. Prior to each cut being made, the resectoscope was held in a fixed position by clamping it in a chemical retort stand. The clamping allowed the software to obtain an up-to-date model of the phantom which correctly reflected the deformation and position changes of the phantom. Once a complete pass of the ultrasound probe had completed (*i.e.* the model was updated), a resection stroke was made using the trigger action alone and with the cutting foot pedal depressed. During the course of the resection, the software saved all the acquired information to disk, including the ultrasound images, the detected outlines in these ultrasound images, the generated 3D models and the positions of the cutting strokes in both world and gland coordinates.

7.5.2.3 Post resection analysis

At the end of the resection in both experiments, the gelatin phantom was removed from the mock abdomen and placed in a Perspex box, as shown in figure 7.29a). The box contains a guide hole so that the metal rod (shown in figure 7.27) can be inserted ensuring correct position of the phantom within the box. The box was filled with gelatin, and left to set. Once solid, the Perspex sides of the box were replaced by aluminium plates containing a series of slots (figure 7.29b). The




Figure 7.29: Prostate phantom in the Perspex cutting box, a) with Perspex sides for filling, b) with cutting guide sides for slicing.




Figure 7.30: A slice of the resected phantom.

slots allow the phantom to be cut accurately into 10 mm thick slices. These slices were removed from the box and digitized using a flat bed scanner¹⁴. Figure 7.30 shows a slice from the resected phantom.

The phantom slice images were loaded into a Matlab based software application and manually outlined by selecting points on the boundary through which a smooth spline was fitted. The periphery of the resection volume was also high-

¹⁴The slices from the static probe experiment were photographed using a digital camera on a tripod placed directly over the prostate portion of each slice. Scanning was used in the second experiment since it provides a more accurate image dimensionally. The poorer accuracy of the camera is apparent in figure 7.31 on page 281 where the inner surface of the ultrasound probe hole is visible (*cf* figure 7.30).



Figure 7.31: Phantom slice with overlaid hand drawn outlines (green), tracked outlines (blue), untracked outline (red) and initial resection cavity (magenta).

lighted by hand using the same spline fitting process. These two sets of splines were then interpolated to provide a regular grid of points in 3D space (6 outlines each with 50 points). The 3D points were then loaded into a 3D modelling application¹⁵ to produce rendered images of the phantom internal boundaries.

7.5.3 Results

Figure 7.31 shows a slice of the phantom from the static ultrasound probe experiment overlaid with various outlines. The magenta outline shows the initial resection cavity detected from the first set of ultrasound images through the whole phantom. The green outlines show the outer boundary of the prostate and the periphery of the resection cavity, found manually. The outer blue outline shows the boundary of the prostate detected in the ultrasound image taken through the same plane in the phantom as the subsequent slice, and with no deformation of the phantom. The inner blue outline shows the boundary of the resection volume detected by the CASTMS system. This outline was found by first finding the points of intersection between the plane of the phantom slice and the resection strokes. Circles of 3mm radius (the same as the resectoscope's cutting loop) were then placed at each of the intersection points, and a smooth outline fitted around the outside of the circles and initial resection boundary. Figure 7.32 shows the hand drawn outline (blue) surrounding the initial resection boundary (magenta)

¹⁵Rhinoceros 3D, Robert McNeel & Associates, Seattle, Washington, USA.




Figure 7.32: Finding the resection boundary.

and resection circles (black) drawn at the resection intersections (red crosses)¹⁶. The red outline in figure 7.31 shows the outline of the resection detected by CASTMS without using the ultrasound tracking. This outline was found in the same manner as the tracked blue outline, except that the resection intersections were calculated using the world coordinates of the resection strokes.

The outlines shown in figure 7.31 show a good correspondence between the actual boundary of the prostate, and the outline detected by image processing of the corresponding ultrasound image. The tracked resection outline shows an over-estimation of the resection cavity in comparison to the actual resection periphery. However, the tracked resection outline shows an improvement over the resection cavity outline formed with no tracking.

Figure 7.33 shows a plot of the average distance (see section 5.4.2.2, chapter 5) between the actual resection outline and the tracked/untracked detected resection outlines, and the areas of these outlines for each slice throughout the phantom. The graph shows that the over-estimation of the resection cavity is apparent through the whole of the phantom, however, the tracked resection cavity is consistently closer to the actual resection than the respective region formed without tracking.

Figure 7.34 shows a slice of the phantom from the resection experiment using a clamped resectoscope and moving ultrasound probe. The figure shows the outlines of the prostate and resection cavity, found by hand-outlining (green), tracking (blue), no tracking (red) and initial cavity (magenta) in the same manner as described previously. Whilst still over-estimating the resection cavity, the blue tracked outline gives a closer match to the true outline when compared to the result from the static probe experiment (figure 7.31).

¹⁶Figure 7.32 shows only those resection intersections and circles which contribute to the final position of the detected outline to preserve clarity.




Figure 7.33: Plots of outline area and average distance between the tracked/untracked outlines and the true resection cavity, for the static ultrasound probe experiment.




Figure 7.34: Phantom slice with overlaid hand drawn outlines (green), tracked outlines (blue), untracked outline (red) and initial resection cavity (magenta).





Figure 7.35: Plots of outline area and average distance between the tracked/untracked outlines and the true resection cavity, for the clamped resectoscope experiment.

Figure 7.35 shows the plot of average distance and outline area for the tracked and untracked outlines in the clamped resectoscope experiment. The graph shows an improved match between the tracked and true resection cavities compared to the static ultrasound probe experiment, however, there is still an overestimation in the resection area for each slice through the phantom.


The 3D models of the true resection cavity (green), tracked resection cavity (blue) and untracked resection cavity (red) are shown in figure 7.36.

7.6 Discussion


The results presented in section 7.5.3 show that the outline cavity detected whilst using ultrasound tracking is closer to the true resection cavity than the untracked resection cavity, for both the experiments. It is also clear that the detected resection cavity during the clamped resectoscope experiment is more accurate than the cavity in the static ultrasound probe experiment, with average and Hausdorff distance measures of approximately 3mm and 5mm respectively between the tracked boundaries and hand outlined boundaries in the worst cases. The corresponding values for the static probe experiment are 5mm and 13mm, the latter value comparable to the estimated Hausdorff distance error in the CASTS system using predetermined deformation models and phantom clamping (see section 3.3 on



(a)



(b)



(c)

Figure 7.36: a) Resection boundary calculated using US tracking. b) Hand outlined resected volume. c) Resection volume without tracking.

page 75)¹⁷.

The static ultrasound probe experiment demonstrates that in its current form, the CASTS system does not achieve sufficient benefit from ultrasound tracking to justify its use. There are several reasons for the discrepancies between the hand drawn resection cavity outlines and those from the computer system.

When cutting the phantom (or real prostate), each resection stroke is a smooth motion whereby the phantom is moved and deformed by the resectoscope as it is moved to the initial position of the cut, and the cutting motion is completed without hesitation. Between each cut the user often moves the resectoscope away from the current site of resection in order to observe the resected area. At present, the ultrasound image processing takes too long for the computer model of the phantom to be moved in time to follow the phantom's real movements. A further source of error is due to the assumption of no deformation of the phantom, however, this is overshadowed by the mismatch in the position of the computer model and the real phantom.

There are some errors introduced in the hand outlining of the phantom slices. In the static probe experiment, the phantom slice images were acquired using a digital camera. Although a tripod was used to try and ensure accurate imaging, the resulting pictures are not ideal. This can be observed in figure 7.31 where the sides of the ultrasound access hole are visible, and the sides of the slice are not straight (not shown in figure 7.31). For the clamped resectoscope experiment, these initial errors were overcome by using a scanner to image the slices rather than a camera. However, several further aspects may have introduced errors into the hand outlining. The thickness of each slice and the transparency of the surrounding gelatin can make it difficult to discern the correct outline – when the prostate section of the phantom on the visible side of the slice is smaller than the section on the underside, then the coloured lower regions of the prostate are visible and obscure the outline. This is compounded by leaching of the colour from the prostate into the surrounding gelatin. Melting of the phantom during the filling of the phantom slicing box was minimized by ensuring that the gelatin was sufficiently cool prior to pouring.

The error of the Flashpoint system has been reported as 1 ± 0.54 mm [130]. Calibration data for the Flashpoint system used in the CASTMS trial shows a spatial

¹⁷Although accurate assessments of the errors in the CASTS system were not performed, the Hausdorff distance can be estimated as 12 mm by examination of the user interface screen shots.



Figure 7.37: “Pistol” type transrectal ultrasound probe.

accuracy of 0.25 ± 0.13 mm. Spatial accuracy in the ultrasound system was experimentally determined as 1 ± 0.3 mm. This was evaluated by using a phantom of thin wires at 10 mm spacings in a water bath and measuring the positions of the echos on the resulting ultrasound image.

The problem of the inability to perform a complete resection of the phantom prostate could be overcome by redesigning the phantom to increase the vertical distance between the ultrasound access hole and the resectable region. This would allow full resection of the phantom without interference from the ultrasound probe. In the OT, a modified probe design would prevent the obstruction to the resectoscope motion. Figure 7.37 shows an older “pistol” type transrectal probe which would cause less interference to the resectoscope’s movement.

The use of a modern 3D ultrasound system could bring multiple benefits to the CASTMS system. Firstly, modern machines produce digital images of far greater resolution and clarity than the outdated ultrasound systems used for the laboratory trials and to acquire the data set for the outline image processing assessment. This would improve the accuracy of the outlining algorithms, and probably reduce the inter-observer error of the ultrasound outlining experts, since a clearer outline is likely to be present in the prostate images.

3D ultrasound would also remove the necessity for moving the ultrasound probe during the course of resection. The 3D ultrasound probe could be clamped to the operating table, and its position registered to the CASTMS system initially, without the necessity to keep track of the position through the operation.

Some problems would still need to be addressed before a 3D ultrasound system could be incorporated in to the CASTMS set-up. At present, only one 3D machine exists in the market which provides a dedicated 3D transrectal probe – the




Figure 7.38: a) The Kretz Voluson 730 3D ultrasound system. b) Example prostate volume image and rendered prostate segmentation.

Voluson 730 from Kretz¹⁸ (shown in figure 7.38). This machine can produce 3D volume images of the prostate at a rate of up to 4 volumes per second, which is fast enough for monitoring prostate position and deformation during the course of a resection. However, discussions with the Kretz technical department reveal that the only way to obtain the volume information at present is to pause volume acquisition, save the acquired data to the machine's internal hard drive, and copy it onto a networked computer or removable media drive. Therefore it is not possible to access the data during the acquisition process at present. Collaboration with the manufacturers would be needed to allow on-line access to the data during imaging.

The problem of the detrimental effect of the electrosurgical current on the ultrasound image is unlikely to be affected by the change to a modern ultrasound system. However, there are possible solutions which deserve investigation. The purpose of the current is to induce heating in the tissue during cutting, so that fluid within the cells vaporize and the cells explode. The current for cutting is an undamped, continuous sinusoidal waveform of constant frequency, therefore filtering of ultrasound echo signals at the same frequency should restore the ultrasound image. Pulsing the current on and off so that ultrasound images are acquired during the off periods could provide a solution, however, it is likely that the pulsing would result in a cutting action which is not smooth. In fact, pulsing the current on and off is the method used for coagulating tissue - the heating pe-

¹⁸KretzTechnik AG, Zipf, Austria.

riods are not long enough for vaporization to occur, but the tissue is denatured, although the pulsing period is much shorter than would be used in an interleaved cutting/ultrasound acquisition set-up. The solution to the interference problem would again require collaboration with the ultrasound manufacturers to insert additional filtering into the image acquisition process.

One further possibility for overcoming the corruption to the ultrasound would be to move away from electrosurgery as the method for resecting tissue, and employ an alternative cutting method such as laser [146] or cryosurgery [33, 134]. Further investigation is necessary to assess the potential of these therapies in comparison to the current gold standard of electrosurgery, and their influence (if any) on the ultrasound image.

Section 7.6 of chapter 6 discussed various methods which could be used to increase the speed of the image processing to find the outline of the prostate in the ultrasound scans. Adoption of a 3D ultrasound system does not mean that the image processing method is made redundant. The statistical shape models presented in chapter 6 are fully extensible to 3D; the principles behind the optimization of the corresponding points and the genetic algorithm method of finding the optimal solution would remain the same. The accuracy of the image processing is also likely to improve because of the additional information from adjacent scans which are not available in the current set-up where the scans are offset from one another.

The final accuracy of the clamped resectoscope experiment is promising when the inherent errors of the Flashpoint tracking and ultrasound imaging systems are taken into account. A maximum Hausdorff distance of 5mm for a final CASMS system, would suggest that resection could be performed safely to within 7.5mm of the capsule and verumontanum without the possibility of damage to these sensitive structures. This would provide the ability to perform an adequate resection to overcome the symptoms of benign prostatic hyperplasia. However, repeated experiments (both lab and OT based) are necessary to confirm the accuracy of ultrasound tracking. This further work is included in section 8.5.1, chapter 8, but falls outside the scope of this thesis.

7.7 Conclusion

This chapter has described initial attempts to modify the original CASTS TURP training system to produce an OT based CASMS system for monitoring the

progress of TURP operations. This goal was deemed unattainable at present due to the problems experienced during initial OT trials, with corruption of the ultrasound image during cutting and coagulation being the most prominent difficulty. Therefore, the emphasis of the research was changed to a laboratory based simulation of the operation using a prostate phantom allowing the principle of ultrasound monitoring to be assessed without the inherent OT difficulties. The lab system also allowed the investigation of ultrasound imaging as a potential method for maintaining registration between real and computer phantom models in the CASTS system – the dual role lab system has been dubbed CASTMS.

The chapter described the additional and the adaptations to hardware components required for the CASTMS system. The method for obtaining the outline of the phantom from the ultrasound system was described and an example result given. The process of generating a registered computer model of the phantom from these outlines has been explained. The chapter also introduced the design and implementation of the software controlling the CASTMS system.

The results of two experiments have been given. The first experiment was an attempt to investigate the possibility of using ultrasound to improve the registration between the cutting of the real phantom, and the modelled cutting displayed on the CASTS user interface. The result of the experiment showed that the time delay required to process the ultrasound images was too long for timely updating of the computer model, and the modelled resection volume did not correspond to reality.

The second experiment showed that ultrasound-based tracking is a viable method for monitoring the prostate (phantom or real) during both training and the real TURP operation, provided the problems of time delay and image interference are overcome. Various solutions to these problems have been discussed, as has the outline for further work which is required before the final goal of a full monitoring system during in-vivo operations is achievable.

Chapter 8

Conclusions and future work

8.1 Introduction

This thesis has described research in the fields of training and monitoring systems for minimally invasive surgery. Although much of the research is relevant to these generic fields, the prototype systems which have been developed are specific to one minimally invasive procedure – transurethral resection of the prostate.

The chapter begins with an overall summary of the research in terms of the original objectives, given in section 8.2. The remainder of the chapter provides conclusions and directions for further work in the three major areas of the research – the TURP CASTS trainer (section 8.3), outline detection in prostate ultrasound images (section 8.4) and the TURP CASMS monitoring system (section 8.5).

8.2 Summary of conclusions

This research has made significant advances in the areas of MIS training, prostate ultrasound outline detection and a novel application MIS monitoring. The outcomes of the research, in terms of the original research objectives (presented on page 27), are as follows:

- A hybrid training simulator for TURP has been implemented and has received favourable reviews by users.
- Previously presented methods for detecting the prostate outline in ultrasound images have been investigated, but none provided adequate results.

Significant gains in accuracy have been achieved by the optimized outlining algorithm presented in this thesis.

- An optimal indicator for distinguishing the prostate outline in ultrasound images has been found by correlating the expert drawn outlines to the image pixel information.
- 2D ultrasound is not a viable method for monitoring prostate motion and deformation during most surgical applications. 3D ultrasound is a potential solution, but will require collaboration with the manufacturer, especially for electrosurgical applications such as TURP.
- A fast algorithm for generating and interrogating a 3D model of the treatment model from prostate outlines has been designed.
- Significant progress has been made towards a fully working, OT based monitoring system for TURP. A laboratory based system has shown promising results as a proof of concept.

In summary, a novel hybrid training system has been presented which incorporates the benefits of existing virtual reality and physical phantom based trainers, using a new phantom which is both permeable by ultrasound and can be resected in a similar manner to a real prostate. An original application – a system for monitoring the progress of minimally invasive surgery – has been described and significant progress has been made towards this goal. An optimized algorithm for detecting the outline of the prostate in ultrasound images has been presented and shown to perform better than previously reported methods.

8.3 MIS training system

Chapter 2 described many of the training systems currently available which allow trainee surgeons to practice the skills required for minimally invasive surgery. These systems were categorized into two types depending on the method of modelling the surgical target, i.e. using virtual-reality or physical phantoms. The advantages and disadvantages associated with each type have been presented.

The work presented in chapter 3 describes the development of a training system for TURP, the computer assisted surgical training system (CASTS). CASTS uses a novel hybrid approach, incorporating aspects of both virtual reality and physical

phantom based systems; the prostate phantom provides realistic force feedback particularly during resection, and requires the use of real surgical instruments in the same manner as the actual operation. The virtual reality component provides a computer generated model of the physical phantom which allows additional navigational information to be presented to the surgeon, although the usual VR goal of realistic representation is not implemented since it is not applicable.

Initial trials of the system received favourable reviews from the participating urology surgeons suggesting that the additional information provided by CASTS is beneficial to the process of training in the TURP operation. However, the trainer was marred by a lack in registration between the computer model of the prostate and the physical prostate phantom, resulting in the incorrect display of resection strokes on the CASTS computer. The registration problems were found to be caused by a combination of factors: variation in the manufacturing tolerances of the phantom, swelling of the phantom due to absorption of irrigant, and deformation under the action of the resectoscope. Initial attempts to restrain or pre-calculate the deformation of the prostate phantom gave some improvement, but were not sufficient to eliminate the registration issues, and did not tackle the differences caused by variable manufacturing and fluid absorption.

Several methods of tracking the prostate phantom were discussed as potential solutions to the problems of mis-registration, with ultrasound imaging found to be the most promising option. To this end, a novel prostate phantom has been developed which may be resected but is also pervious by ultrasound. The prostate phantom can give relatively clear ultrasound images (when compared to clinical prostate images), but consistent quality is difficult due to the difficulties of manufacture, particularly the problems of eliminating trapped air bubbles in the gelatin mixture.

Use of ultrasound for tracking the prostate phantom closes the gap between the training system and a laboratory mock-up of the monitoring system. The laboratory system was discussed in chapter 7, which also presented the results of an experiment using ultrasound to track the prostate phantom in the training system. However, despite the success of an automated image processing algorithm in detecting the phantom prostate outline (described in section 7.4.2 of chapter 7), the experiment showed that the trainer does not currently achieve increased registration accuracy through the use of ultrasound tracking. The lack of success in solving the registration issues is due to a combination of time delays resulting from the processing of the ultrasound images and latency in movement

of the ultrasound probe. These problems could be solved by the adoption of a 3D ultrasound imaging system to eliminate the requirement for probe movement and optimization of the prostate outline detection algorithm run on a high end workstation to reduce the processing time.

The original objective for the research on the TURP trainer was to develop a system which “provides additional navigation information and enhances the learning process, with the ability to evaluate training performance”. The system which has been developed fulfills the first two parts of the objective, but the problems of registration reduce the system’s benefits. Evaluating training performance is also possible using the system, since data such as the resectoscope movements and associated timings, and the extent of resection, can be saved for review. However, the mechanisms for interpreting this data have not been investigated, and remain a direction for further research.

8.3.1 Future work

3D ultrasound imaging has been proposed as a method for obtaining images of the prostate phantom, without the need (and associated time delay) for moving the ultrasound probe. However, such ultrasound hardware is very expensive, and it is unlikely that a CASTS TURP system which uses 3D ultrasound as a method of tracking the prostate phantom, would find much interest outside the realms of a research project. Therefore, given the benefit of hindsight, this author advocates the adoption of a purely VR reality based system, incorporating haptic feedback to provide the tactile sensation, if the objective of the research is solely toward training systems for MIS. There is a large scope for research in the haptic field, both in terms of the mechanical design of manipulators with multiple axes of force-feedback, and in the mechanical modelling of soft tissues to provide realistic force information.

8.4 Prostate ultrasound image processing

The work of chapters 4, 5 and 6 investigated methods of obtaining the outline of the prostate from ultrasound images. The basis for the work was to use the outline information to monitor changes in shape and position of the prostate as it is influenced by the resectoscope during the course of TURP surgery. Chapter 4 described many techniques from the literature which have been used previously

to detect objects of interest in images. These algorithms covered the range from standard, commonplace image processing techniques through to specialized routines for detecting objects in ultrasound images, and transrectal prostate ultrasound in particular. Several algorithms were implemented and assessed in terms of their qualitative result when applied to a prostate ultrasound image of high quality.

Five algorithms (all previously used for prostate detection in ultrasound images) were deemed worthy of further investigation. Chapter 5 described the research undertaken to determine the best algorithm across a set of ultrasound images covering the range in quality encountered during routine prostate ultrasound examinations. Acquisition of the data set was described, and several possible methods for deciding the position of the correct outline were presented, with averaging of multiple experts' outlines chosen as the most suitable approach. An improved method for assigning correspondency between outlines was described, which forms a more robust algorithm in comparison to previously described corresponding point routines. Statistical methods for assessing the algorithms' outlines were presented, and the results of the analysis presented.

A criticism of most of the previous prostate outline detection literature is that the only presented results show the algorithm applied to a single image, which is always of high quality in ultrasound terms. In the few cases where multiple results are given, the number of trials is still relatively low. The statistical analysis described in this thesis represents the first rigorous comparison of prostate outline detection algorithms on a set of ultrasound images.

One outline detection method showed a clear margin of success in its best results when compared to the other algorithms. However, this method has a degree of randomness which can also result in very poor results – far worse than any other method. Therefore the direction of the research was directed toward optimizing this best method to minimise the occurrence of the poor results; the optimization proceeded along two separate avenues which were presented in chapter 6. The two candidates for optimization were the prostate outline model which dictates the permissible shapes of the outline results, and the test used to assess the fit of a potential outline to the ultrasound image data. The model optimization work followed a similar path to previous work in the literature, but is the first use of an optimised model for prostate outline detection in ultrasound images.

The optimization of the method for testing the fit of a candidate outline to the ultrasound image was achieved by performing a statistical analysis of many

potential indicator functions. The basis for the assessment was a novel direction for outline detection algorithms and a reversal of current methodologies; rather than assuming some characteristic of an image gives a good indication of the presence of an outline, start from the outline and find the indicator which best highlights the outline region in comparison to non-outline regions. The techniques for finding the best indicator have a potential role in many other applications where low signal-to-noise ratios and poor edge definitions are apparent, such as other ultrasound image processing tasks.

The results of the optimization give an algorithm whose outlining performance is significantly better than those reported previously, even in its worst set of results. This algorithm has many applications in urological surgery, where prostate outlining is currently performed or where its use provides better estimates of clinically relevant factors over other methods, such as for prostate volume determination. However, the algorithm also has potential uses for processing ultrasound images of other areas of the anatomy, e.g. the ventricles of the heart, the bladder and the kidney. Indeed, the algorithm should give better results for these applications over the prostate since their ultrasound images display more contrast at the desired outline locations.

The second research objective listed in the introductory chapter of this thesis, to assess current image processing methods to find one suitable for prostate outline detection in ultrasound images, has been fulfilled. However, none of the currently available methods was sufficiently accurate so the additional work to optimize the outline detection represents a significant advance beyond the bounds of the original objectives. The optimization has also fulfilled the third research objective – identification of the best indicators for the presence of the prostate outline.

The final part of chapter 6 described the testing of the algorithm's execution speed following implementation in C++. The result demonstrates that in its present form, the algorithm is too slow to process ultrasound images intra-operatively such that the outlines obtained can provide the desired prostate deformation and movement information. Therefore, the outcome of the fourth research objective, whether ultrasound is a viable method for tracking the prostate, is a negative response at present.

8.4.1 Future work

The optimization of the method for assessing candidate outlines resulted in an indicator function which gives significantly different results for the outline regions when compared to areas where an outline is not present. However, the statistical analysis can only discern the best indicator from those which are included in the assessment. Many other image processing techniques could be included with the potential to provide better results than those reported in this thesis, *e.g.* textural energy measures. Other possibilities include using multiple indicators to obtain a combined test which performs better than the separate components.

One area where further research should be directed is in the study of the role of edge direction in comparison to ultrasound beam direction. The underlying physics of ultrasound image formation indicates that a boundary between two regions will appear differently in the image as the angle between the boundary and the ultrasound path is altered. Therefore, it is probable that an indicator which clearly discerns edges at one angle will perform less well at different angles, and may be better replaced by a different indicator at those angles.

The current implementation of the algorithm executes faster than any previously reported method for detecting the prostate outline in ultrasound images. However, further speed gains are required before the algorithm can be used in its intended role for intra-operative ultrasound processing. Several potential methods for decreasing the processing time were discussed in chapter 6, which form a basis for future work.

The final area for further work questions the assumption made in this thesis (in section 4.1.1, chapter 4) that the best method of detecting deformation and motion of the prostate with ultrasound imaging is to detect the prostate outline in the images. During the latter stages of the research the author discovered work in the literature which attempts to estimate the motion between ultrasound image frames by correlating speckle patterns, *e.g.* the work reported by Strintzis [203] and Kokkinidis [110]. However, the presented applications are for diagnostic purposes (*e.g.* of tumours in echocardiograms) and neither paper provides results which demonstrate that the detected motion correlates to the true motion. Therefore, this author remains skeptical that motion estimation is a viable method for tracking the prostate, particularly since deformation and out-of-(image)-plane motions occur, but further investigation is warranted.

8.5 MIS monitoring system

Chapter 7 presented the work undertaken with the goal of developing a monitoring system for TURP. The early sections of the chapter describe the components required for the monitoring system, many of which are similar to those used in the training system. Initial trials of various components of the monitoring system were conducted in the OT to assess their suitability for use in the final monitoring system. The major findings of these OT trials were that intra-operative ultrasound was not a feasible method for tracking the prostate due partly to the latency involved in moving the ultrasound probe so that images across the whole treatment volume can be obtained, but primarily because the electrosurgical current required for resecting tissue causes severe disruption to the ultrasound image, at the time when outline information is most required. Therefore, the hypothesis presented as the fourth original research objective is resolved – at present, ultrasound imaging is not a viable method for tracking the prostate during TURP.

Despite this setback, there are potential solutions to the problems, which were described in section 7.6 of chapter 7. In addition, prostate based operations which do not involve electrosurgery, such as brachytherapy, would also benefit from tracking of prostate motion and deformation during the procedure. Therefore, the remainder of the work of chapter 7 was directed towards implementing a laboratory based monitoring system with the goal of investigating the potential of ultrasound tracking on a prostate phantom, in the absence of the image degradation caused by the electrosurgery.

The processing time required for finding the phantom prostate outline was significantly reduced in the lab system. This was achieved by reducing the amount of searching performed by the algorithm, however, this did not reduce accuracy since the ultrasound images of the phantom show clearer boundaries at the periphery of the prostate when compared to real prostate ultrasound images. The degradation to the ultrasound image was removed by abandoning electrosurgery as a resection method in favour of mechanical slicing using the cutting loop as a blade. However, the delay in obtaining ultrasound images could not be removed since a period of time is required to translate the ultrasound between image capture locations.

The fifth objective for the research from chapter 1 was to design an algorithm which takes the outlines detected from a set of ultrasound images covering the treatment volume, and generates a 3D model of this volume, which can then

be interrogated to find the closest boundary to a given point (the resectoscope's cutting loop). The modelling algorithm presented in chapter 7 is an adaptation of an established technique, exploiting the smoothness of the prostate shape so that the dimensions of the problem are reduced with an associated reduction in processing time. However, the algorithm does not give an exact measure of the closest boundary point.

The results of the experiment conducted to investigate assess ultrasound as a potential method for tracking the prostate show encouraging results in terms of the overall errors between the modelled and actual resection volumes, when factors such as the positional accuracy of the tracking system and ultrasound hardware are taken into account. However, further experimental work must be performed to corroborate these initial results.

The final original objective of the research was the development of a monitoring system for use in the OT during TURP procedures. Clearly this was an ambitious target which has not been achieved by the current research. However, significant advances and findings have been made during the course of the research which move a step closer to this ultimate goal.

8.5.1 Further work

The problems of interference with the ultrasound image and the time delay in obtaining images across the treatment volume preclude the possibility of tracking the prostate during TURP using the ultrasound hardware available for this research project. However, potential solutions to both problems were discussed in section 7.6 of chapter 7, principally the adoption of a modern 3D ultrasound system. Significant involvement of the ultrasound manufacturer would be required, since changes to the ultrasound hardware would be necessary, but should be possible since the benefit (and therefore lucrativeness) of an accurate MIS monitoring system for TURP is large. The application also gives further justification for the existence of 3D ultrasound for prostate imaging – until now the benefit of 3D ultrasound for prostate imaging has not yet been commercially demonstrated.

Modern scanners also provide clearer ultrasound images due to advances in the technology and adoption of digital images. The optimised outline detection algorithm presented in this thesis is extendable to 3D, although a large amount of work would be required in obtaining a set of volumetric images, the expert segmentations of the images into prostate and non-prostate volumes, building

and optimizing the volumetric shape model, and finding the optimal indicator for distinguishing the prostate boundary. Alternatively, the volume image data could be split into a set of parallel 2D slices and processed in the identical manner to those presented in this thesis.

The algorithm for modelling the treatment volume was found to be adequate for the laboratory based experiments. However, further work should be done to obtain comparisons of the algorithm's performance with other model generation/collision detection algorithms.

Finally, further validation work will be required to judge the performance of the monitoring system at each stage of its evolution, including both evaluation in the laboratory, and assessment during TURP surgery in the operating theatre.

Bibliography

- [1] Aarnink R, Geisen R, Huynen A, Rosette J, Debruyne F, Wijkstra H 1994. A practical clinical method for contour determination in ultrasonographic prostate images. *Ultrasound Med Biol*, vol 20, no 8, pp 705-717.
- [2] Aarnink R, Huynen A, Giesen R, Rosette J 1995. Automated prostate volume determination with ultrasonographic imaging. *J Urol*, vol 153, pp 1549-1554.
- [3] Aarnink R, Pathak S, Rosette J, Debruyne F, Kim Y, Wijkstra H 1998. Edge detection in prostatic ultrasound images using integrated edge maps. *Ultrasonics*, vol 36, pp 635-642.
- [4] Arambula-Cosio F 1997. Towards a safe and effective robotic prostatectomy. PhD thesis, Imperial College, London, UK.
- [5] Arambula-Cosio F, Davies B 1999. Automated prostate recognition: a key process for clinically effective robotic prostatectomy. *Med Biol Eng Comput*, vol 37, pp 236-243.
- [6] Baker J 1987. Reducing bias and inefficiency in the selection algorithm. *Proceedings 2nd International Conference on Genetic Algorithms*, Hillsdale, NJ, USA, pp 14-21.
- [7] Bargar W, Bauer A, Borner M 1998. Primary and revision total hip replacement using the Robodoc (R) system. *Clinical Orthopaedics and Related Research*, pp 82-91.
- [8] Basset O, Sun Z, Mestas J, Gimenez G 1993. Texture analysis of ultrasonic images of the prostate by means of co-occurrence matrices. *Ultrasonic Imaging*, vol 15, pp 218-237.
- [9] Ballard D, Brown C 1982. *Computer Vision*, Prentice-Hall, New York.

- [10] Bamber J and Daft C 1986. Adaptive filtering for reduction of speckle in ultrasonic pulse-echo images. *Ultrasonics*, vol 24, pp 41-44.
- [11] Bauer J, Stein R, McPherson D 1991. Segmentation of high-frequency ultrasound images of atherosclerotic coronary arteries. *Comput Biol Med*, vol 21, no 5, pp 275-287.
- [12] Baumberg A and Hogg D 1994. Learning flexible models from image sequences. *Proceedings European Conference on Computer Vision*, Sweden, pp 299-308.
- [13] Baur C, Guzzoni D, Georg O 1998. VIRGY: a virtual reality and force feedback based endoscopic surgery simulator. *Studies in Health Technology & Informatics*, vol 50, pp 110-116.
- [14] Belohlavek M and Greenleaf J 1997. Detection of cardiac boundaries in echocardiographic images using a customized order statistics filter. *Ultrasonic Imaging*, vol 19, pp 127-137.
- [15] Benayoun A, Ayache N, Cohen I 1994. Adaptive meshes and nonrigid motion computation. *Proceedings International Conference on Pattern Recognition*, Jerusalem, Israel, pp 730-732.
- [16] Courtesy of Biocomputation Centre, Stanford University, California, USA. <http://biocomp.stanford.edu>
- [17] Birkfellner W, Watzinger F, Wanschitz F, Ewers R, Bergmann H 1998. Calibration of tracking systems in a surgical environment. *IEEE Transactions on Medical Imaging*, vol 17, no 5, pp 737-742.
- [18] Blake A, Zisserman 1987. Visual reconstruction. MIT Press, Cambridge, Mass, USA.
- [19] Bland J, Altman D 1986. Statistical methods for assessing the agreement between two methods of clinical measurement. *Lancet*, vol 1, pp 307-310.
- [20] Blandy J 1983. Benign prostatic hypertrophy, chapter 85. Springer Verlag, New York.
- [21] Bovik A, Huang T, Muson D 1983. A generalization of median filtering using linear combinations of order statistics. *IEEE Trans Acoust Speech Signal Processing*, vol 31, pp 1342-1349.

-
- [22] Bovik A 1988. On detecting edges in speckle imagery. *IEEE Trans Acoust Speech Signal Processing*, vol 36, no 10, pp 1618-1627.
- [23] Courtesy of Boston Dynamics Incorporated, Massachusetts, USA. <http://www.bdi.com>
- [24] Boyle R and Thomas R 1988. Computer vision: a first course. Blackwell Scientific Publications, Oxford, UK, pp 48-50.
- [25] Breen S 1999. The evaluation of automated boundary detection algorithms for ultrasound images of the prostate: development of a gold standard and statistical methods for comparisons. Internal report, Nanyang Technological University, Singapore.
- [26] Brown J, Montgomery K, Latombe J, Stephanides M 2001. A microsurgery simulation system. *MICCAI 2001*, Utrecht, The Netherlands.
- [27] Burckhardt C 1978. Speckle in ultrasound B-mode scans. *IEEE Transactions on Sonics and Ultrasonics*, vol SU25, no 1, pp 1-6.
- [28] Burckhardt C, Flury P, Glauser D 1995. Stereotaxic brain surgery – integrated Minerva system meets demanding robotic requirements. *IEEE Engineering in Medicine and Biology Magazine*, vol 14, pp 314-317.
- [29] Bude R, Adler R 1995. An easily made, low-cost, tissue-like ultrasound phantom material. *J Clin Ultrasound*, vol 23, pp 271-273.
- [30] Bush N, Hill C 1995. Gelatine-alginate complex gel: a new acoustically tissue-equivalent material. *Ultrasound Med Biol*, vol 9, no 5, pp 479-484.
- [31] Çakmak H, Kühnapfel U, Bretthauer G 2000. Virtual Reality Techniques for Education and Training in Minimally Invasive Surgery. *Proceedings of VDE World Micro Technologies Conference MICROtec*, pp. 395-400.
- [32] Canny J 1986. A computational approach to edge detection. *IEEE Trans Pattern Analysis Machine Intelligence*, vol 8, pp 679-697.
- [33] Carroll P, Presti Jr J, Small E, Roach M 1997. Focal therapy of prostate cancer 1996: maximizing outcome. *Urology*, vol 49, pp 84-94.
- [34] Catalona W, Beiser J, Smith D 1997. Serum free prostate specific antigen and prostate specific antigen density measurements for predicting cancer in men with prior negative prostatic biopsies. *J Urol*, vol 158, pp 2162-2167.

- [35] Cattolica E, Sidney S, Sadler M 1997. The safety of transurethral prostatectomy: a cohort study of mortality in 9416 men. *J Urol*, vol 158, pp 102-104.
- [36] Chakravarti, Laha, Roy 1967. Handbook of methods of applied statistics, volume 1. John Wiley & Sons, New York, USA, pp 392-394.
- [37] Chalana V, Winter T, Cyr D, Haynor D, Kim Y 1996. Automatic fetal size measurements from ultrasound images. *Academic Radiol*, vol 3, pp 628-635.
- [38] Chalana V, Kim Y 1997. A methodology for evaluation of boundary detection algorithms on medical images. *IEEE Trans Med Imag*, vol 16, no 5, pp 642-652.
- [39] Chassat F and Lavalle S 1998. Experimental protocol for accuracy evaluation of 6-D localizers for Computer-Integrated Surgery: Application to four optical localizers. *MICCAI 1998*, Boston, USA, pp 358-367.
- [40] Clark J, Nicholls J, Cooper M, Bates S, Frappell J, Byrne D 2001. An evaluation of a laparoscopic ectopic simulation by trainees. *Gynaecological Endoscopy*, vol 10, no 5-6, pp. 309-314(6).
- [41] Cohen L 1991. On active contour models and balloons. *CVGIP: Image Understanding*, vol 52, no 2, pp 211-218.
- [42] Cinquin P, Lavallée S, Troccaz J 1992. IGOR: Image guided operating robot. Methodology, applications. *Proc Ann Int Conf IEEE Eng Med Biol Soc*, vol 14, pp 1048-1049.
- [43] Cipolla R, Blake A 1990. The dynamic analysis of apparent contours. *Proceedings 3rd international computer vision conference*, Osaka, Japan, pp 616-623.
- [44] Cooperberg P 1991. Artifacts in US. Ultrasound syllabus, *Radiological Society of North America Inc*, Illinois USA.
- [45] Cootes T, Taylor C, Cooper D, Graham J 1995. Active shape models - their training and application. *Comput Vision Image Understanding*, vol 61, pp 38-59.
- [46] Cos L 1990. Simulator for transrectal ultrasound of prostate. *Urology*, vol 35, no 5, pp 450-451.

-
- [47] Crivianu-Gaita D, Miclea F, Gaspar A, Margineatu D, Holban S 1997. 3D reconstruction of prostate from ultrasound images. *Int J Medical Informatics*, vol 45, pp 43-51.
- [48] Czerwinski R, Jones, O'Brien W 1995. Ultrasound speckle reduction by directional median filtering. *Proceedings IEEE international conference on image processing*, Washington, DC, USA, pp 358-360.
- [49] Czerwinski R, Jones, O'Brien W 1998. Line and boundary detection in speckle images. *IEEE Trans Imag Processing*, vol 7, no 12, pp1700-1714.
- [50] Czerwinski R, Jones, O'Brien W 1999. Detection of lines and boundaries in speckle images - application to medical ultrasound. *IEEE Trans Med Imag*, vol 18, no 2, pp 126-136.
- [51] D'Ancona F, Francisca E, Witjes W, Welling L, Debruyne F, de la Rosette J 1997. High energy thermotherapy versus transurethral resection in the treatment of benign prostatic hyperplasia: results of a prospective randomized study with 1 year followup. *J Urol*, vol 158, pp 120-125.
- [52] Darzi A, Smith S, Taffinder N 1999. Assessing operative skill: Needs to become more objective. *British Medical Journal*, vol 318, pp 887-888.
- [53] Davies E 1990. Machine vision: theory, algorithms and practicalities. Academic Press, San Diego, California, USA, ch 5.
- [54] Davies B, Hibberd R, Timoney A, Wickham J 1991. The development of a surgeon robot for prostatectomies. *J Eng Med (Proc Inst Mech Engrs)*, vol 205, part H, pp 35-38.
- [55] Davies B, Hibberd R, Ng W, Timoney A, Wickham J 1991. A surgeon robot for prostatectomies. 5th Int Conf Adv Robotics, pp 871-875.
- [56] Davies B, Ho S, Hibberd R 1994. The use of force control in robot assisted knee surgery. *Proc 1st Int Symp Med Robotics Comp Assisted Surgery*, vol 2, pp 258-262.
- [57] Davies B, Harris S, Lin W, Hibberd R, Middleton R, Cobb J 1997. Active compliance in robotic surgery - the use of force control as a dynamic constraint. *Proc Inst Mech Engrs*, vol 211, pp285-292.

- [58] Davies R, Cootes T, Twining C, Taylor C 2001. An information theoretic approach to statistical shape modelling. *British Machine Vision Conference*, Manchester.
- [59] Davies R, Cootes, T, Waterton J, Taylor C 2001. An efficient method for constructing optimal statistical shape models. *MICCAI 2001*, Utrecht, Netherlands, pp 57-65.
- [60] Davies R, Twining C, Cootes T, Waterton J, Taylor C 2002. A minimum description length approach to statistical shape modelling. *IEEE Trans Med Imag*, vol 21, no 5, pp 525-537.
- [61] De Boor C 1978. A practical guide to splines. Springer-Verlag, New York, USA.
- [62] Delaunay B 1934. Sur la sphere. Vide *Izvestia Akademia Nauk SSSR*, VII Seria, Otdelenie Matematicheskii i Estestvennyka Nauk, vol 7, pp 793-800.
- [63] Delingette H, Hébert M, Ikeuchi K 1992. Shape representation and image segmentation using deformable surfaces. *Image Vision Comput*, vol 10, no 3, pp 132-144.
- [64] Detmer P, Bashein G, Martin R 1990. Matched filter identification of left-ventricular endocardial borders in transesophageal echocardiograms. *IEEE Trans Med Imag*, vol 9, pp 396-404.
- [65] Drake J, Joy M, Goldenburg A, Kreindler D 1991. Computer and robot assisted resection of thalamic astrocytomas in children. *Neurosurgery*, vol 29, pp 27-33.
- [66] Drake J, Joy M, Goldenburg A, Kreindler D 1991. Computer and robotic assisted resection of brain tumours. *Proc Ann int Conf IEEE Eng Med Biol Soc*, vol 13, pp 888-891.
- [67] Duda R, Hart P 1972. Use of the Hough Transformation to detect lines and curves in pictures. *Comm ACM*, vol 15, no 1, pp 11-15.
- [68] Dyer C 2001. GMC extends restrictions on Bristol heart surgeon. *BMJ*, vol 322, no 7300, pp 1441.
- [69] Eliasziw M, Young S, Woodbury M, Fryday-Field K 1994. Statistical methodology for the concurrent assessment of interrater and intrarater reliability: using goniometric measurements as an example. *Phys Ther*, vol 74, pp 777-788.

- [70] European Association of Endoscopic Surgeons 1994. Training and assessment of competence. *Surg Endoscopy* vol 8, pp 721-722.
- [71] Fast and Light Toolkit. Available from www.ftk.org.
- [72] Fowler F et al 1988. Symptom status and quality of life following prostatectomy. *J Am Med Assoc*, vol 259, pp 3018-3022.
- [73] Fowler M, Scott K 1999. UML distilled 2nd ed. Addison-Wesley. Massachusetts, USA.
- [74] Friedland N, Adam D 1989. Automatic ventricular cavity boundary detection from sequential ultrasound images using simulated annealing. *IEEE Trans Med Imag*, vol 8, pp 344-353.
- [75] Fredfeldt K 1986. An easily made ultrasound biopsy phantom. *J Ultrasound Med*, vol 5, pp 295-297.
- [76] Geiser E, Conetta D, Limacher M, Stockton V, Olivier H, Jones B 1990. A second-generation computer-based edge detection algorithm for short-axis two-dimensional echocardiographic images: accuracy and improvement in interobserver variability. *J Amer Soc Echocardiol*, vol 3, pp 79-90.
- [77] Gibson S, Fyock C, Grimson E, Kanade T, Kikinis R et al 1998. Simulating surgery using volumetric object representations, real-time volume rendering and haptic feedback. *Medical Image Analysis*, vol 2, pp 121-132.
- [78] Courtesy of Global Technologies, Florida, USA. <http://www.global-technologies.net>
- [79] Goodman J 1977. Statistical properties of laser speckle patterns. *Laser Speckle and Related Phenomena*, edited by Dainty J, Springer-Verlag, Berlin, Germany, pp 9-75.
- [80] Gomes M et al 1993. Computer assisted ultrasound segmentation for robotic prostatectomy. *Proceedings 5th Portuguese conference on pattern recognition*. Maia, Portugal.
- [81] Gomes M, Barrett A, Timoney A, Davies B 1999. A computer-assisted training/monitoring system for TURP structure and design. *IEEE Trans Inf Tech Biomed*, vol 3, no 4, pp 242-251.

-
- [82] Gonzales R and Woods R 1992. Digital image processing. Addison-Wesley, Massachusetts, USA, pp 414-428.
- [83] Grefenstette J, Fitzpatrick J 1985. Genetic search with approximate function evaluations. *Proceedings international conference genetic algorithms and applications*, pp 112-120.
- [84] Grimm P, Blasko J, Ragde H 1994. Ultrasound guided transperineal implantation of iodine 125 and palladium 103 for the treatment of early stage prostatic cancer. *Atlas Urol Clin North Amer*, vol 2, pp 113-125.
- [85] Grimm P, Blasko J, Ragde H, Sylvester J, Clarke D 1996. Does brachytherapy have a role in the treatment of prostate cancer?. *Hemat Oncol Clin North Amer*, vol 10, pp 653-672.
- [86] Hammoude A 1999. Edge detection in ultrasound images based on differential tissue attenuation rates. *Ultrasonic Imaging*, vol 21, no 1, pp 31-42.
- [87] Haralick R 1984. Digital step edges from zero crossing of second directional derivatives. *IEEE Trans Pattern Analysis Machine Intelligence*, vol 6, no 1, pp 58-68.
- [88] Haralick R and Shapiro L 1993. Computer and robot vision 2nd ed. Addison-Wesley, Massachusetts, USA.
- [89] Harris S, Lin W, Fan K, Hibberd R, Cobb J, Middleton R, Davies B 1997. Experiences with robotic systems for knee surgery. *Proc 1st Int Conf Comp Vision, VR & Robotics in Medicine & Medical Robotics & Comp Assisted Surgery*, pp 757-766.
- [90] Harris S, Jakopc M, Hibberd R, Cobb J, Davies B 1998. Interactive pre-operative selection of cutting constraints, and interactive force controlled knee surgery by a surgical robot. *1st Int Conf Medical Image Comp & Comp Assisted Intervention*, pp 996-1006.
- [91] Harris S, Jakopc M, Cobb J, Davies B 1999. Intra-operative application of a robotic knee surgery system. *Proc 2nd Int Conf Med Im Comp & Comp Assisted Intervention*, pp 1117-1124.
- [92] van Herk M, Bruce A, Guus Kroes A, Shouman T, Touw A, Lebesque J 1995. Quantification of organ motion during conformal radiotherapy of the prostate

- by three dimensional image registration. *Int J Radiation Oncology Biol Phys*, vol 33, no 5, pp 1311-1320.
- [93] Hill A and Taylor C 1994. Automatic landmark generation for point distribution models. *Proceedings British Machine Vision Conference*, pp 429-438.
- [94] Hill C and ter Haar G 1995. High intensity focussed ultrasound: potential for cancer treatment. *Br J Radiol*, vol 68, pp 1296-1303.
- [95] Courtesy of Hillway Surgical Ltd, West Sussex, UK.
<http://www.surgimodels.com>
- [96] Courtesy of Human Interface Technology Laboratory, University of Washington, Washington, USA. <http://www.hitl.washington.edu>
- [97] Ho S, Hibberd R, Cobb J, Davies B 1995. Force control for robotic surgery. *Proc 7th Int Conf Adv Robotics*, vol 1, pp 21-31.
- [98] Ho S, Hibberd R, Davies B 1995. Robot assisted knee surgery. *IEEE Eng Med Biol*.
- [99] Hollander J, Diokner A 1996. Prostatism: benign prostatic hyperplasia. *Urol Clin North Am*, vol 23, pp 75-86.
- [100] Holtgrewe H, Mebust W, Dowd J, Peters P, Proctor C 1989. Transurethral prostatectomy: practice aspects of the dominant operation in American urology. *J Urol*, vol 162, pp 1633.
- [101] Horowitz S, Pavlidis T 1974. Picture segmentation by a directed split-and-merge procedure. *Proc 2nd Int Joint Conf Pattern Analysis*, pp 424-433.
- [102] Houck C, Joines J, Kay M 1995. A genetic algorithm for function optimization: a Matlab implementation. Technical report NCSU-IE TR 95-09, North Carolina State University, USA.
- [103] Hough P 1962. A method and means for recognizing complex patterns. US Patent 3069654.
- [104] Huttenlocher D, Klanderman G, Rucklidge W 1993. Comparing images using the Hausdorff distance. *IEEE Trans Pattern Analysis Machine Intelligence*, vol 15, pp 850-863.
- [105] Jansen J 1997. MEng thesis, Imperial College, London, UK.

- [106] Kambhamettu C and Goldgof D 1992. Point correspondence recovery in non-rigid motion. *Proceedings IEEE Conference on Computer Vision and Pattern Recognition*, pp 222-227.
- [107] Kass M, Witkin A, Terzopoulos D 1987. Snakes: active contour models. *Int J Computer Vision*, vol 1, no 4, pp 321-331.
- [108] Kirkpatrick, S. Gelatt C, Vecchi, M 1983. Simulated annealing. *Science*, vol 220, pp 671-680.
- [109] Knoll C, Alcaniz-Raya M, Monserrat C, Grau-Colomer V, Juan M 1999. Multiresolution segmentation of medical images using shape restricted snakes. *SPIE Med Imag*, vol 3661, pp 287-292.
- [110] Kokkinidis I, Strintzis M, Maglaveras N 1995. Motion estimation in ultrasound B-mode images using a maximum likelihood matching criterion. *17th Annual Conference, IEEE Engineering in Medicine and Biology Society*, vol 1, pp 593-594.
- [111] Koo J, Park S 1991. Speckle reduction with edge preservation in medical ultrasonic images using a homogeneous region growing median filter. *Ultrasonic Imaging*, vol 13, pp 211-237.
- [112] Kotcheff A and Taylor C 1998. Automatic construction of eigenshape models by direct optimisation. *Medical Image Analysis*, vol 2, pp 303-314.
- [113] Kremkau F and Taylor K 1986. Atrifacts in ultrasound imaging. *Journal of Ultrasound in Medicine*, vol 5, pp 227-237.
- [114] Krivanek A, Sonka M 1998. Ovarian ultrasound image analysis: follicle segmentation. *IEEE Trans Med Imag*, vol 17, no 6, pp 935-944.
- [115] Kwoh Y 1985. A new computerised tomographic aided robotic stereotaxis system. *Robotics Age*.
- [116] Kwoh Y, Hou J, Jonckheere A, Hayati S 1988. A robot with improved absolute positioning accuracy for CT guided stereotactic brain surgery. *IEEE Trans Biomed Eng*, vol 35, pp 153-159.
- [117] Ladak H, Downey D, Steinman D, Fenster A 1999. Semi-automatic technique for segmentation of the prostate from 2D ultrasound images. *Proceedings 1st joint BMES/EMBS conference*. Atlanta, GA, USA, p 1144.

- [118] Ladak H, Mao F, Wang Y, Downey D, Steinman D, Fenster A 2001. Prostate segmentation from 2D ultrasound images. *Medical Physics*, vol 27, pp 1777-1788.
- [119] Laing F 1983. Commonly encountered artifacts in clinical ultrasound. *Seminars in Ultrasound*, vol 4, pp27-43.
- [120] Larsson A 2001. An open and flexible framework for computer aided surgical training. *Studies in Health Technology & Informatics*, vol 81, pp 263-265.
- [121] Lavallée S, Cinquin P, Demongeot J, Benabid A, Marque I, Djaid M 1989. Computer assisted driving of a needle into the brain. *Proc CAR'89*, pp 416-420.
- [122] Lavallée S, Cinquin P 1989. Computer assisted medical interventions. *Proc CAR'89*, pp 301-312.
- [123] Lavallée S 1989. A new system for computer assisted neurosurgery. *Proc Ann Int Conf IEEE Eng Med Biol Soc*, pp 926-927.
- [124] Laws K 1980. Textured image segmentation. PhD thesis, University of Southern California, Los Angeles, CA, USA.
- [125] Lee Y, Kassam S 1985. Generalized median filtering and related nonlinear filtering techniques. *IEEE Trans Acoust Speech Signal Processing*, vol 33, pp 672-683.
- [126] Lee Y, Rhodes W 1990. Feature detection and enhancement by a rotating kernel min-max transformation. *Hybrid image and signal processing II, proceedings SPIE*, vol 1297, 154-159.
- [127] Leroy B, Herlin I, Cohen L 1996. Multiresolution algorithms for active contour models. *12th Int Conf Analysis and Optimization of Systems*, pp 58-65.
- [128] Levienaise-Obadia B, Gee A 1999. Adaptive segmentation of ultrasound images. *Image Vision Comput*, vol 17, pp 583-588.
- [129] Levine M 1981. A modular computer vision system for picture segmentation and interpretation. *IEEE Trans Pattern Anal Machine Intelligence*, vol PAMI-3, no 5, pp540-558.

- [130] Li Q, Zamorano L, Jiang Z, Gong J, Pandya A, Prez R, Diaz F 1999. Effect of optical digitizer selection on the application accuracy of a surgical localization system - a quantitative comparison between the Optotrak and Flashpoint tracking systems. *Computer Aided Surgery*, vol 4, no 6, pp 314-321.
- [131] Courtesy of Limbs and Things Ltd, Bristol, UK. <http://limbsandthings.com>
- [132] Liao G, Nodes T, Gallagher N 1985. Output distributions of two dimensional median filters. *IEEE Trans Acoust Speech Signal Processing*, vol 33, pp 1280-1295.
- [133] Liberty J 1998. Beginning object-oriented analysis and design. Wrox Press Ltd, Birmingham, UK.
- [134] Littrup P, Sparschu A, Grignon D 1995. Prostatic cryotherapy: current experiences. *Cancer*, vol 75, pp 1957-1962.
- [135] Litynski G 1996. Highlights in the history of laparoscopy : the development of laparoscopic techniques- a cumulative effort of internists, gynecologists, and surgeons. Bernert Verlag, Frankfurt, Germany.
- [136] Liu Y, Ng W, Teo M, Lim H 1997. Computerised prostate boundary estimation of ultrasound images using radial bas-relief method. *Med Biol Eng Comput*, vol 35, pp 445-454.
- [137] Lo S 1991. Volume discretization into tetrahedra - II. 3D triangulation by advancing front approach, *Computers and Structures*, vol 39, no 5, pp 501-511.
- [138] Lobregt S, Viergever M 1995. A discrete dynamic contour model. *IEEE Trans Med Imag*, vol 14, no 1, pp 12-24.
- [139] Lohner R 1996. Progress in grid generation via the advancing front technique. *Engineering with Computers*, vol 12, pp 186-210.
- [140] Loupas T, McDicken N, Allan P 1989. An adaptive weighted median filter for speckle suppression in medical ultrasonic images. *IEEE Trans Circuits Systems*, vol 36, no 1, pp 129-135.
- [141] Madersbacher S, Pedevilaa M, Vingers L, Susani M, Marberger M 1995. Effect of high-intensity focussed ultrasound on human prostate cancer in vivo. *Cancer Res*, vol 55, pp 3346-3351.

- [142] Magic software source code for computer graphics, image analysis and numerical methods. Available from <http://www.magic-software.com>.
- [143] McCarthy A and Hollands R 1997. Human factors related to virtual reality surgical simulator: the Sheffield knee arthroscopy training system. *Proceedings of the UK Virtual Reality Special Interest Group Conference 1997*, Uxbridge, UK, pp 150-158.
- [144] Martin J, Regehr G, Reznick R, MacRae H, Murnaghan J, Hutchison C, et al 1997. Objective structured assessment of technical skill (OSATS) for surgical residents. *Br J Surg*, vol 84, pp 273-278.
- [145] Martini F 1998. Fundamentals of anatomy and physiology, 4th edition, chapter 28. Prentice Hall, New Jersey, USA.
- [146] Matsuoka K, Iida S, Tomiyasu K, Shimada A, Suekane S, Noda S, 1998. Holmium laser resection of the prostate. *Journal of Endourology*, vol 12, no 3, pp 279-82.
- [147] McCloy R, Wilson M, Sutton C, Middlebrook A, Chater P, Stone R 1997. MIST VR: a part-task virtual reality trainer for laparoscopic surgery. *Journal of Telemedicine and Telecare*, 2 June 1997, vol 3, no 1, supp 1, pp. 97-97.
- [148] Mebust W, Holtgrewe H, Cockett A, Peters P 2002. Transurethral prostatectomy: immediate and postoperative complications. a cooperative study of 13 participating institutions evaluating 3,885 patients. *Journal of Urology*, vol 167, pp 999-1004.
- [149] Courtesy of Mentice Medical Simulation, Gothenburg, Sweden. <http://www.mentice.com>
- [150] Mikic I, Krucinski S, Thomas J 1998. Segmentation and tracking in echocardiographic sequences: active contours guided by optical flow estimates. *IEEE Trans Med Imag*, vol 17, no2, pp 274-284.
- [151] Michel M, Knoll T, Kohrmann K, Alken P 2002. The URO Mentor: development and evaluation of a new computer-based interactive training system for virtual life-like simulation of diagnostic and therapeutic endourological procedures. *BJU International*, vol 89, no 3, pp 174-177.
- [152] McNamara M, McNamara M 1989. Preparation of a homemade ultrasound biopsy phantom. *J Clin Ultrasound*, vol 17, pp 456-458.

- [153] Mojsilovic A, Popovic M, Amodaj N, Babic R, Ostojic M 1997. Automatic segmentation of intravascular ultrasound images: a texture-based approach. *Ann Biomed Eng*, vol 25, no 6, pp 1059-1057.
- [154] Montgomery K, Bruyns C, Wildermuth S, Heinrichs L, Hasser C, Ozenne S, Bailey D 2001. Surgical simulator for operative hysteroscopy. *IEEE Visualization 2001*, San Diego, California.
- [155] Montilla G, Barrios V, Subacius V, Rangel N, Roux C 1993. Model-based epicardial boundary detection using genetic algorithms. *Proceedings EMBS*, pp 226-227.
- [156] Muller W, Ziegler R, Bauer A, Soldner E 1995. Virtual reality in surgical arthroscopic training. *Journal of Image Guided Surgery*, vol 1, no 5, pp 288-294.
- [157] Ng W, Davies B, Timoney A, Hibberd R 1993. The use of ultrasound in automated prostatectomy. *Med Biol Eng Comput*, vol 31, pp 349-354.
- [158] Nicholson R, Crofton M 1997. Training phantom for ultrasound guided biopsy. *Br J Radiol*, vol 70, pp 192-194.
- [159] Page B 1980. The pathological anatomy of digital enucleation for benign hyperplasia and its application to endoscopic resection. *Br J Urol*, vol 52, pp 111-126.
- [160] Patel A, Harrington T, Saunt K, Jones W 1996. Construction of an ultrasound biopsy phantom. *Australasian Radiology*, vol 40, pp 185-186.
- [161] Pathak S, Chalana V, Haynor D, Kim Y 1999. Edge guided delineation of the prostate in transrectal ultrasound images. *Proceedings 1st joint BMES/EMBS conference*. Atlanta, GA, USA, p 1056.
- [162] Pathak S, Chalana V, Haynor D, Kim Y 1999. Edge-guided boundary delineation in prostate ultrasound images. *IEEE Trans Med Imag*, vol 19, no 12, pp 1211-1219.
- [163] Perona P, Malik J 1990. Scale-space and edge detection using anisotropic diffusion. *IEEE Trans Pattern Analysis Machine Intelligence*, vol 12, no 7, pp 629-639.
- [164] Pitas I, Venetsanopoulos A 1992. Order statistics in digital image processing. *Proceedings of the IEEE*, vol 80, no 12, pp 1893-1921.

- [165] Prager, R, Gee H, Berman L 1999. Stradx: real time acquisition and visualization of freehand three-dimensional ultrasound. *Medical Image Analysis*, vol 3, no 2, pp 129-140.
- [166] Prater J, Richard W 1992. Segmenting ultrasound images of the prostate using neural networks. *Ultrasonic Imaging*, vol 14, pp 159-185.
- [167] Pratt W, 1978. Digital Image Processing. Wiley, New York, USA, pp 330-333.
- [168] Pratt W, 1991. Digital Image Processing 2nd ed. Wiley, New York, USA.
- [169] Press W, Teukolsky S, Vetterling W, Flannery B 1993a. Numerical recipes in C: the art of scientific computing. Cambridge University Press, Cambridge, pp 50-51.
- [170] Press W, Teukolsky S, Vetterling W, Flannery B 1993b. Numerical recipes in C: the art of scientific computing. Cambridge University Press, Cambridge, pp 73-75.
- [171] Press W, Teukolsky S, Vetterling W, Flannery B 1993c. Numerical recipes in C: the art of scientific computing. Cambridge University Press, Cambridge, pp 504-510.
- [172] Press W, Teukolsky S, Vetterling W, Flannery B 1993c. Numerical recipes in C: the art of scientific computing. Cambridge University Press, Cambridge, pp 444-455.
- [173] Presti Jr J, Hovey R, Bhargav V, Carroll P, Shinohara K, Moul J 1997. Prospective evaluation of prostatic specific antigen and prostate specific antigen density in detection of carcinoma of prostate: ethnic variations. *J Urol*, vol 157, pp 907-912.
- [174] Rabben S, Torp A, Stoylen A, Slordahl S, Bjornstad K, Haugen B, Angelsen B 2000. Semiautomatic contour detection in ultrasound M-mode images. *Ultrasound Med Biol*, vol 26, no 2, pp 287-296.
- [175] Ramsey E, Miller P, Parsons K 1997. A novel transurethral microwave thermal ablation system to treat benign prostatic hyperplasia: results of a prospective multicentre clinical trial. *J Urol*, vol 158, pp 112-118.
- [176] Rangarajan A, Chui H, Bookstein F 1997. The Softassign procrustes matching algorithm. *Proceedings of IPMI'97*, pp 29-42.

- [177] Richard W, Grimmel C, Bedigian K, Frank K 1993. A method for three-dimensional prostate imaging using transrectal ultrasound. *Comput Med Imag Graph*, vol 17, no 2, pp 73-79.
- [178] Richard W, Keen C 1996. Automated texture-based segmentation of ultrasound images of the prostate. *Comput Med Imag Graph*, vol 20, no 3, pp 131-140.
- [179] Rifkin M 1988. Ultrasound of the prostate, chapter 1. Raven Press, New York.
- [180] Van-Rij A, McDonald J, Pettigrew R, Putterill M, Reddy C, Wright J 1995. Cusum as an aid to early assessment of the surgical trainee. *Br J Surg*, vol 82. pp 1500-1503.
- [181] Rissanen J 1983. A universal prior for integers and estimation by minimum description length. *Annals of Statistics*, vol 11, no 2, pp 416-431.
- [182] Rissanen J 1989. Stochastic complexity in statistical inquiry. *World Scientific Series in Computer Science*, vol 15, World Scientific Publishing Co, Singapore.
- [183] Roberts L 1965. Machine perception of 3D solids. Optical and electro-optical information processing, MIT Press.
- [184] Rosenfeld A, Kak A 1982. Digital picture processing, 2nd edition. Academic Press, New York.
- [185] Rubin J, Adler R, Bude R, et al 1991. Clean and dirty shadowing at US: a reappraisal. *Radiology*, vol 181, pp 231-236.
- [186] Russell C and Cheslyn-Curtis S 1994. Minimal Access General Surgery - General Principles.
- [187] Sagiv C, Akselrod S, Tepper R 1999. Application of a semiautomatic boundary detection algorithm for the assessment of amniotic fluid quantity from ultrasound images. *Ultrasound Med Biol*, vol 25, no 4, pp 515-526.
- [188] Sarty G, Liang W, Sonka M, Pierson R 1998. Semiautomated segmentation of ovarian follicular ultrasound images using a knowledge-based algorithm. *Ultrasound Med Biol*, vol 24, no 1, pp 27-42.

- [189] Satava R 1998. Current and future applications of virtual reality for medicine. *Proceedings of the IEEE*, vol 86, no3, pp 484-489.
- [190] Scallan SE, Fairholm DJ, Cochrane DD, Taylor DC 1992. Evaluation of the operating room as a teaching venue. *Can J Surg*, vol 35, pp 173-176.
- [191] Schijven M and Jakimowicz J 2002. Face-, expert, and referent validity of the Xitact LS 500 Laparoscopy Simulator. *Surgical Endoscopy*, vol 16, no 8.
- [192] Schmidt RA 1991. Processing information and making decisions. *Motor Learning and Performance*. Champaign, Ill: Human Kinetics Publishing; pp 26-28.
- [193] Courtesy of Select-IT VEST Systems AG, Bremen, Germany.
<http://www.select-it.de>
- [194] Courtesy of Simulab Corporation, Washington, USA.
<http://www.simulab.com>
- [195] Courtesy of Simbionix USA Corporation, Ohio, USA.
<http://www.simbionix.com>
- [196] Slater G, Jourdan I, Folscher D, Snook A, Cooper M, D'Allessandro P, Rangeley C, Bailey M 2001. The Guildford MATTU TEP hernia model. *Surgical Endoscopy*, vol 15, no 5, pp 493-496.
- [197] Sonka M, Zhang X, Siebes M, Bissing M, DeJong S, Collins S, McKay C 1995. Segmentation of intravascular ultrasound images: a knowledge-based approach. *IEEE Trans Med Imag*, vol 14, pp 719-732.
- [198] Spencer F 1978. Teaching and measuring surgical techniques: the technical evaluation of competence. *Bull Amer Coll Surg* 1978, vol 63, pp 9-12.
- [199] Staib L, Duncan J 1992. Boundary finding with parametrically deformable models. *IEEE Trans Pattern Analysis Machine Intelligence*, vol 14, no 11, pp 1061-1075.
- [200] Starkie, S 2002. The application of virtual environments as constraints in robotic surgery. PhD thesis, Imperial College, UK.
- [201] Stevens J 1996. Robodoc: a kinder, gentler surgeon's assistant. *Technology Review*, vol 99, pp 14-17.

- [202] Steen E, Olstad B 1994. Scale-space and boundary detection in ultrasonic imaging using nonlinear signal-adaptive anisotropic diffusion. *Proceedings SPIE medical imaging*, Newport Beach, USA, vol 2167, pp 116-127.
- [203] Strintzis, G, Kokkinidis I 1997. Maximum likelihood motion estimation in ultrasound image sequences. *IEEE Signal Processing Letters*, vol 4, no 6, pp 156-157.
- [204] Szekely G, Brechcuhler C, Hutter R, Rhomberg A, Schmid P 1998. Modelling of soft tissue deformation for laparoscopic surgery simulation. *MICCAI1998*, Boston, USA.
- [205] Taffinder N, Sutton C, Fishwick R, McManus I, Darzi A 1998. Validation of virtual reality to teach and assess psychomotor skills in laparoscopic surgery: results from randomised controlled studies using the MIST VR laparoscopic simulator. *Studies in Health Technology & Informatics*, vol 50, pp 124-30.
- [206] Taffinder N, McManus I, Jansen J, Russell R, Darzi A 1998. An objective assessment of surgeons' psychomotor skills: validation of the MIST-VR laparoscopic simulator. *British Journal of Surgery*, vol 85(suppl 1) pp 75.
- [207] Tagare H 1999. Shape-based non-rigid correspondence with application to heart motion analysis. *IEEE Transactions on Medical Imaging*, vol 18, no 7, pp 434-439.
- [208] Tammela T 1997. Benign prostatic hyperplasia. Practical treatment guidelines. *Drugs Aging*, vol 10, pp 349-366.
- [209] Tang S, Kwoh C, Teo M, Ng W, Ling K 1998. Augmented reality systems for medical applications. *IEEE Engineering in Medicine and Biology*, vol 17, no 3, pp 49-58.
- [210] Tendick F, Downes M, Goktekin T, Cavusoglu M, Feygin D, Wu X, Eyal R, Hegarty M, Way L 2000. A virtual environment testbed for training laparoscopic surgical skills. *Presence*, vol 9, no 3, pp 236-255.
- [211] Terzopoulos D, Platt J, Barr A, Fleischer K 1987. Elastically deformable models. *Comput Graphics*, vol 21, no 4, pp 205-214.
- [212] Tong S, Cardinal H, McLoughlin R, Downey D, Fenster A 1998. Intra- and inter-observer variability and reliability of prostate volume measurement via

- two-dimensional and three-dimensional ultrasonic imaging. *Ultrasound Med Biol*, vol 24, no 5, pp 673-681.
- [213] Topper A, Jernigan M 1991. Estimation of bladder wall location in ultrasound images. *Med Biol Eng Comput*, vol 29, pp 297-303.
- [214] Tsai L, Chen T, Chen T, Chen J 1997. Application of automatic boundary detection for computerized quantitative analysis of left ventricular regional wall motion by two-dimensional echocardiography. *J Ultrasound Med*, vol 16, pp 177-182.
- [215] Wang Y, Peterson B, Staib L 2000. Shape-based 3D surface correspondence using geodiscs and local geometry. *Proceedings of CVPR*, vol 2, pp 644-651.
- [216] Wall M 1994. GAlib – genetic algorithm library. Available from <http://lancet.mit.edu/ga>.
- [217] Watanabe H and Kojima M 1997. The role of prostatic planimetry using transrectal sonography in prostatic diseases. *Semin Sur Oncol*, vol 13, pp 425-430.
- [218] Watson D 1981. Computing the Delaunay tessellation with application to Voronoi polytopes. *The Computer Journal*, vol 24, no 2, pp 167-172.
- [219] Weghorst S, Airola C, Oppenheimer P, Edmond C, Patience T, Heskamp D, Miller J 1998. Validation of the Madigan ESS simulator. *Studies in Health Technology & Informatics*, vol 50, pp 399-405.
- [220] Whittington R, Malkowicz S, Machtay M, van Arsdalen K, Barnes M, Broderick G, Wein A 1997. The use of combines radiation therapy and hormonal treatment in the management of lymph node-positive prostate cancer. *Int J Radiat Oncol Biol Phys*, vol 39, pp 673-680.
- [221] Wills D, Abblard E, Sherman K 2002. A Generic Arthroscopy Simulator Architecture. *Proceedings of Medicine Meets Virtual Reality 2002*, IOS Press, pp 573-579.
- [222] Williams G 1976. Comparing the joint agreement of several raters with another rater. *Biometrics*, vol 32, pp 619-627.
- [223] Woo M, Neider J, Davis T 1998. OpenGL Programming Guide: the official guide to learning OpenGL, ver 1.1, 2nd ed, Addison Wesley Longman Inc.

- [224] Wu R 1999. Automated prostate boundary recognition. MEng thesis, Nanyang Technological University, Singapore.
- [225] Wu R, Ling K, Ng W 2000. Automated prostate recognition in ultrasound images using feature modelling and genetic algorithm. *J Ultrasound Med*, vol 19, pp771-782.
- [226] Xu C, Prince J 1998. Snakes, shapes and gradient vector flow. *IEEE Trans on Image Processing*, pp 359-369.
- [227] Xu C, Yezzi A, Prince J 1999. On the relationship between parametric and geometric active contours. Technical Report JHU/ECE 99-14, The John Hopkins University, Baltimore.
- [228] Xuan J, Wang Y, Sesterhenn I, Moul J, Mun S. 3-D model supported prostate biopsy simulation and evaluation. *MICCAI1998*, Boston, USA.
- [229] Yamauchi Y, Suzuki M, Mochimaru M et al, 2001. A Training System for Endoscopic Sinus Surgery with Skill Evaluation. *Proceedings of CARS 2001*, Berlin.
- [230] Yeung F, Levinson S, Fu D, Parker K 1998. Feature-adaptive motion tracking of ultrasound image sequences using a deformable mesh. *IEEE Trans Med Imag*, vol 17, no 6, pp 945-956.
- [231] Yerry M and Shephard M, 1984. 3D mesh generation by modified octree technique. *International Journal for Numerical Methods in Engineering*, vol 20, pp 1965-1990.
- [232] Zaman M, Moloney C 1993. A comparison of adaptive filters for edge-preserving smoothing of speckle noise. *IEEE international conference on acoustics, speech and signal processing*, vol 5, pp77-79.
- [233] Zeng J, Kaplan C, Bauer J, Sesterhenn I, Moul J, Mun S. Visualization and evaluation of prostate needle biopsy. *MICCAI1998*, Boston, USA.
- [234] Zhao F, deSilva C 1998. Use of the laplacian of gaussian operator in prostate ultrasound image processing. *Proceedings 20th IEEE/EMBS conference*. Hong Kong, pp 812-815.
- [235] Zimmer Y, Tepper R, Akselrod 1996. A two-dimensional extension of minimum cross entropy thresholding for the segmentation of ultrasound images. *Ultrasound Med Biol*, vol 22, no 9, pp 1183-1190.

- [236] Ziskin M, Thickman D, Goldenberg N et al 1982. The comet-tail artifact.
Journal of Ultrasound in Medicine, vol 1, pp 1-7.

Appendix A

Speckle reduction algorithms

This appendix reviews many of the speckle reduction algorithms found in the literature. The algorithms fall into two categories; those which attempt to reduce speckle by smoothing of areas with near constant grey level, and those which reduce speckle by emphasizing probable lines within the image.

The algorithms have been developed using the MatlabTM¹ programming platform and its Image Processing Toolbox. This package was chosen since it allows fast implementation of algorithms without the need to code much of the functionality, such as loading, saving and displaying images. The Image Processing Toolbox also provides many useful functions, such as morphological operators and filters.

Each speckle reduction algorithm has been applied to a transrectal ultrasound prostate image for testing (figure A.1). Most of the algorithms have adjustable parameters which affect the amount of speckle smoothing in the image. The best results are presented for each algorithm, based on a qualitative analysis of the set of results produced by varying the filter parameters.

A.1 Speckle smoothing algorithms

A.1.1 Anisotropic diffusion filter

Perona and Malik [163] describe a noise smoothing technique based on the heat conduction equation, called *anisotropic diffusion*. The basic criterion of the technique is that as smoothing is applied to the image, there is a preference for

¹The Mathworks, Inc., Massachusetts, USA.




Figure A.1: Transrectal ultrasound image of the prostate

smoothing within regions rather than between regions, i.e. that edges between regions remain sharp until their point of disappearance, as the image resolution diminishes. This is in sharp contrast to the Gaussian filter² which provides good speckle smoothing characteristics at the expense of blurring of edge details, with accompanied loss of spatial location.

The process is stated mathematically as:

$$\frac{\delta}{\delta t} I(x, y, t) = \nabla \cdot (c(x, y, t) \nabla I(x, y, t)) \quad (\text{A.1})$$

where $I(x, y)$ is the image, t is the iteration step, ∇ is the gradient operator, and c is a monotonically decreasing diffusion function. Perona and Malik [163] offer two suitable diffusion functions:

$$c(x, y, t) = e^{-\left(\frac{|\nabla I(x, y, t)|^2}{K}\right)} \quad (\text{A.2})$$

and

$$c(x, y, t) = \frac{1}{1 + \left(\frac{|\nabla I(x, y, t)|^{1+\alpha}}{K}\right)} \quad (\text{A.3})$$

where $\alpha > 0$.

Steen and Olstad [202] offer a slightly different function. They argue that the edge enhancement property of the Perona and Malik [163] diffusion functions is unsuitable for ultrasound images, since the true edges are not step functions. They remove the enhancement by setting the flux function to be constant if the local gradient is greater than the standard deviation of the assumed local noise

²The Gaussian filter is often used in ultrasound image processing [80, 177, 234]

level. The flux is defined as:

$$\nabla_N I(x, y) \equiv I(x-1, y) - I(x, y) \quad (\text{A.4})$$

and the diffusion functions as:

$$c(x, y, t) = e^{\left(-\frac{|\nabla I(x, y, t)|^2}{2\sigma^2}\right)} \quad (\text{A.5})$$

if $|\nabla I(x, y, t)| \leq \hat{\sigma}$ or:

$$c(x, y, t) = \frac{e^{-\frac{1}{2}\hat{\sigma}}}{|\nabla I(x, y, t)|} \quad (\text{A.6})$$

otherwise. $\hat{\sigma}$ is calculated by assuming that the image is corrupted with signal dependent noise [21]:

$$I = I_0 + I_0^{\frac{1}{2}} \varepsilon \quad (\text{A.7})$$

where I is the observed image, I_0 the original uncorrupted image, and ε is a zero-mean Gaussian with standard deviation σ^3 . $\hat{\sigma}$ is calculated by assuming that the previous anisotropic diffusion filter (ADF) output gives a reasonable approximation to the signal level, such that:

$$\hat{\sigma} = \sigma \hat{I}^{\frac{1}{2}} \quad (\text{A.8})$$

where $\hat{I}$ is the local mean of the signal, calculated over a small window.

In the discrete form equation A.1 becomes:

$$\begin{aligned} I_{t+1}(x, y) = & I_t(x, y) + \lambda [c_{tN}(x, y) \cdot \nabla_N I_t(x, y) + c_{tS}(x, y) \cdot \nabla_S I_t(x, y) \\ & c_{tE}(x, y) \cdot \nabla_E I_t(x, y) + c_{tW}(x, y) \cdot \nabla_W I_t(x, y)] \end{aligned} \quad (\text{A.9})$$

where:

$$\begin{aligned} \nabla_N I(x, y) &\equiv I(x-1, y) - I(x, y), & \nabla_S I(x, y) &\equiv I(x+1, y) - I(x, y) \\ \nabla_E I(x, y) &\equiv I(x, y+1) - I(x, y), & \nabla_W I(x, y) &\equiv I(x, y-1) - I(x, y) \end{aligned}$$

and $0 \leq \lambda \leq \frac{1}{4}$.

Steen and Olstad[202] use a maximum likelihood estimate of the original signal for the initial local signal mean in equation A.8 (3.7), calculated over a 3 by 3 window W :

$$\hat{I} = -\frac{\sigma^2}{2} + \sqrt{\frac{\sigma^4}{4} + \frac{1}{9} \sum I(x, y)^2}, \quad (x, y) \in W \quad (\text{A.10})$$

³ σ is a user-defined parameter




Figure A.2: ADFs applied to the test ultrasound image.

Perona and Malik [163] determine the value of K in equations A.2 and A.3 either as a user-defined parameter, or as 90% of the Canny [32] "noise estimator" – form the histogram of the absolute gradient and take the value at 90% of the cumulative area under the curve.

Figure A.2 shows the result of applying the ADF to the test ultrasound image of figure A.1. The filtered image in figure A.2a) results from using the diffusion function of equation A.2. with $K = 16$, $\lambda = 0.25$, and $n = 100$, where n is the number iterations. The result shown in figure A.2b) was obtained using equation A.3 as the diffusion function with, $K = 4$, $\lambda = 0.25$ and $n = 100$.

A.1.2 Homogeneous region growing median filter

Koo and Park [111] present the *homogeneous region growing mean filter* (HRGMF) for speckle reduction with edge preservation. The mean and variance of the grey level for an initial region centred at the pixel of interest are calculated, and their ratio provides a measure of the local homogeneity:

$$H = \frac{\sigma^2}{\mu} \quad (\text{A.11})$$

where H is the homogeneity, μ is the *current* region mean:

$$\mu = \frac{1}{A} \sum_A I(x, y), \quad (x, y) \in A \quad (\text{A.12})$$

and σ^2 the variance, calculated as:

$$\sigma^2 = \frac{1}{A} \sum_A [I(x, y) - \mu]^2, (x, y) \in A \quad (\text{A.13})$$

where I is the input image, A is the size of the current region, and the initial region size is set slightly larger than the average speckle size.

If the homogeneity is greater than a specified threshold then the region size is iteratively reduced until the threshold criterion is met. Region size reduction is achieved in the following manner. First the top row of pixels in the region is removed, and the homogeneous criterion tested. If the test succeeds (i.e. region is homogeneous) then this region is selected for region growing; if not, then the top row is replaced, and the right hand column removed, and the remainder tested. This process is repeated for each side (top row, right column, bottom row, left column) of the region; if removal of any of the sides fails to produce a homogeneous result, then all sides are removed, and the cycle continues. If the region shrinks to a single pixel (a high noise impulse), then the final grey level is set to the mean of the eight neighbours.

The homogeneity threshold should be calculated for each particular ultrasound machine by determining the value whilst insonating a homogeneous region of a tissue-mimicking phantom. At present, it is set as a user-defined parameter.

Once a homogeneous region (of more than one pixel) is obtained, the algorithm iteratively extends the region whilst the homogeneity criterion is still met. Each side of the region is taken in turn, and the adjacent pixels to that side are considered as a neighbouring region (of size side length by one pixel). Then the relative variance of the neighbouring region is calculated with respect to the mean of the current region:

$$\sigma_N^2 = \frac{1}{D} \sum_D [I(x, y) - \mu_C]^2, (x, y) \in D \quad (\text{A.14})$$

where D is the size of the neighbouring region and μ_C is the mean of the current region (not the combined region). A similarity measure is calculated from the ratio of the relative variance to current region mean, as:

$$H_N = \frac{\sigma_N^2}{\mu_C} \quad (\text{A.15})$$

If this ratio is less than the homogeneity threshold, then the two regions are statistically similar and are joined. Once the final homogeneous region is found, then the new pixel value is set to the mean of the homogeneous region.




Figure A.3: HRGMF result.



Figure A.4: Standard MF result.

Figure A.3 shows the result of applying the HRGMF to the test ultrasound image of figure A.1. The parameters used are $x = 7$, $y = 7$, so that the initial region size is 7 by 7, and $H_T = 0.3$ as the hysteresis threshold.

A.1.3 Median filter

The algorithms in the following sections (A.1.4 – A.1.8) are all adaptations of the standard *median filter* (MF) [167]. The MF is commonly used for impulse noise reduction with edge preservation in image processing applications. The filter replaces the value of a candidate pixel by the median value of the pixels in the n by n region around the pixel.

Figure A.4 shows the result of applying the standard median filter to the test ultrasound image of figure A.1, with $n = 15$.

A.1.4 Separable median filter

The *separable median filter* (SMF) is simply a two pass linear MF applied over the columns of the image, and then over the rows. Liao et al [132] report that




Figure A.5: SMF result.




Figure A.6: VLMF result.

the SMF gives improved signal preservation when compared to the MF.

Figure A.5 shows the result of applying the SMF to the test ultrasound image of figure A.1, with $n = 15$.

A.1.5 Variable length median filter

Loupas et al [140] describe the *variable length median filter* (VLMF). The VLMF applies a standard MF to the neighbourhood of each pixel, but adapts the size of the filter window according to the local grey level statistics. The size of the window n is given by:

$$n = \left\lceil K \left(1 - \frac{C\sigma^2}{\mu} \right) \right\rceil \quad (\text{A.16})$$

where $(2K + 1)$ by $(2K + 1)$ is the maximum window size, σ^2 is the variance of the maximum window, μ is the mean of the maximum window and C is the filter constant. K and C are defined by the user.

Figure A.6 shows the result of applying the variable length median filter to the test ultrasound image of figure A.1, with parameters $K = 7$ and $C = 0.2$.



Figure A.7: DWMTMF result.

A.1.6 Double-window modified trimmed median filter

The *double-window modified trimmed median filter* (DWMTMF) is presented in [125]. The filter assigns a value to a candidate pixel based on the mean value of pixels in an n by n window, which lie within a certain range. The range of pixel values is centred on the median value x of an m by m window around the candidate pixel (where $m < n$), and extends from $-q$ to $+q$, where:

$$q = Cx^{\frac{1}{2}} \quad (\text{A.17})$$

The advantage of the DWMTMF filter is that it does not require calculation of the local variance and therefore should perform faster.

Figure A.7 shows the result of applying the DWMTMF to the test ultrasound image of figure A.1. The user defined parameters are $m = 5$, $n = 9$ and $C = 2$.

A.1.7 Adaptive weighted median filter

Loupas et al [140] present a filter for reducing speckle, known as the *adaptive weighted median filter* (AWMF). The AWMF is an adaptation of the weighted median filter (WMF), which is itself an adaptation of the MF. In the WMF, each pixel in the filter window is added to a vector the number of times prescribed by its associated (integer) weight. If

$$\begin{bmatrix} a & b \\ c & d \end{bmatrix}$$

are the pixels in the filter window, and:

$$w_1 = 2, \quad w_2 = 3, \quad w_3 = 3, \quad w_4 = 2$$




Figure A.8: AWMF result.

are the filter weights, then the filter result is:

$$\text{med}(a, a, b, b, b, c, c, c, d, d)$$

where $\text{med}()$ is the median operator, and $a \leq b \leq c \leq d$. The standard median filter is achieved by setting each weight to unity.

The problem of the WMF is that it applies the same smoothing regardless of the local statistics of the image. This has been addressed by the AWMF, where the weights are adjusted as follows:

$$w(x, y) = f\left(w(K+1, K+1) - cd\frac{\sigma^2}{\mu}\right) \quad (\text{A.18})$$

where σ^2 and μ are the variance and mean respectively of the $(2K+1)$ by $(2K+1)$ window, d is the distance of the point (x, y) from the centre of the window $(K+1, K+1)$, $w(K+1, K+1)$ is the assigned central pixel weight, and c is a scaling constant. The function $f(x)$ gives the nearest integer value to x if x is positive, else zero.

Figure A.8 shows the result of applying the AWMF to the test ultrasound image of figure A.1, with $K = 7$, $w(K+1, K+1) = 99$ and $c = 20$.

A.1.8 Directional median filter

Czerwinski et al [48] present a different type of median filter known as the *directional median filter* (DMF). The DMF uses a set of sticks – binary line segments of constant thickness and varying orientation – to approximate the local features of the image. The line segments for one pixel thick, size five sticks are shown in figure A.9.




Figure A.9: Size five, one pixel thick stick masks

In general, the number of stick masks m for a particular stick size n is calculated as:

$$m = 2(n - 1) \quad (\text{A.19})$$

In the standard MF, the filter window is rectangular and the filter does not take any variation in the local image statistics into account. The AWMF addresses this problem, but still uses a fixed shaped window (circular) which can give poor results when used on images composed mainly of lines. This is the justification for implementing the DMF.

For a stick size n , each stick is applied over the n by n window centred at the candidate pixel. The median is calculated from the set of the pixels which lie along the length of the stick. The DMF output is set to the value of the largest median. In summary:

$$I_{out}(x, y) = \max_{\theta} (\text{med}_{\xi, \eta} [I(x - \xi, y - \eta) s_{\theta}(\xi, \eta)]) \quad (\text{A.20})$$

where $I(x, y)$ is the input image and s_{θ} is a stick of orientation θ .

Figure A.10 shows the result of applying the DMF to the test ultrasound image of figure A.1. The filter uses a set of 44 stick masks of size $n = 23$, and stick thickness three pixels.

A.1.9 Sticks filters

Czerwinski et al [49, 50] present a class of speckle filters using the stick masks developed as part of the DMF. They differ by calculating the mean value of the pixels which lie along the stick in each orientation, rather than the median value. This is equivalent to performing a 2D convolution of the image and stick mask,




Figure A.10: DMF result.

which may be implemented as a quicker algorithm, than the equivalent DMF algorithm.

Different filter results may be obtained depending on how the orientation means are combined. Czerwinski [49] proposes setting the filter output to the maximum value of the mean over all orientations (*max stick filter* MSF), i.e. :

$$I_{out}(x, y) = \max_{\theta} (\text{mean}_{\xi, \eta} [I(x - \xi, y - \eta) s_{\theta}(\xi, \eta)]) \quad (\text{A.21})$$

Czerwinski [50] also presents a different stick filtering operation called *decision-directed sticks* (DDS). The image is again mean filtered with each stick, but the output value is calculated based on the likelihood of a particular stick orientation at the candidate pixel. The most likely stick orientation is found for each point according to the maximum mean value. Next a prior probability is calculated for each angle at each point. Finally, the result is calculated based on the values of the means and the most likely orientation based on the pixel's neighbours.

The process is described mathematically as:

$$J_{\theta}(x, y) = \text{mean}_{\xi, \eta} [I(x - \xi, y - \eta) s_{\theta}(\xi, \eta)] \quad (\text{A.22})$$

$$\alpha(x, y) = \arg \left[\max_{\theta} [J_{\theta}(x, y)] \right] \quad (\text{A.23})$$

$$\pi_{\theta}(x, y) = \sum_{\xi, \eta} f(\alpha(x - \xi, y - \eta) = \theta) \quad (\text{A.24})$$

$$I_{out}(x, y) = \max_{\theta} \left[\log(\pi_{\theta}) + \frac{1}{\sigma^2} J_{\theta}(x, y) \right] \quad (\text{A.25})$$

where $\arg(z)$ is the θ value associated with z , $f(z)$ is the indicator function equal to one if z is true or zero if z is false, π_{θ} is the histogram of stick orientations around the candidate pixel (over the n by n window), and σ^2 is the variance of




Figure A.11: Sticks results using a) the MSF and b) the DDS filter.

the assumed noise (set as a user-defined parameter).

Figure A.11 shows the results of applying the sticks filters to the test ultrasound image of figure A.1. The image on the left (figure A.11a) uses the MSF of equation A.21, while the right hand result uses DDS filtering with the assumed noise variance set as $\sigma^2 = 0.03$. Both results were obtained using a set of stick masks of size $n = 23$ and three pixel thickness.

A.2 Edge enhancing algorithms

A.2.1 Order statistics filter

Pitas and Venetsanopoulos [164] describe a class of digital filters known as order statistics filters. Belohlavek and Greenleaf [14] have used such filters for edge detection in echocardiographic images. Aarnink et al [1] have used a particular implementation of the filters as a step in detecting the prostate outline (Section 4.3.2). The generalized form of an order statistics filter [164] is:

$$y_i = \sum_{j=1}^n a_j x_{(j)} \quad (\text{A.26})$$

where $x_{(j)}$ are ordered pixel intensities within an m by m window, i.e. $x_1 \leq x_2 \leq \dots \leq x_n$, a_j are appropriate weights which determine the filter characteristics and $n = m^2$. Setting $a_j = \frac{1}{n}$ gives the *ordered mean filter*, which produces the same result as a standard mean filter. However, replacing n by βn , $0 \leq \beta \leq 1$,




Figure A.12: Order statistics filter result.

results in the *ranked order mean filter*, with β defining the rank.

$$y_i = \sum_{j=1}^{\beta n} \frac{1}{\beta n} x_{(j)} \quad (\text{A.27})$$

In a discrete filter, the value of β is set according to the selected rank. For an m by m window, rank zero sets $\beta = \frac{1}{n}$, which produces the standard minimum filter (as used by Aarnink [1]), rank k , $1 \leq k \leq m$, sets $\beta = \frac{km}{n}$ with $k = m$ resulting in the standard mean filter.

Similarly,

$$y_i = \sum_{j=n}^{n(1-\beta)+1} \frac{1}{\beta n} x_{(j)} \quad (\text{A.28})$$

produces a filter whose output is biased toward the high value pixels. Rank zero of this filter performs the same operation as the standard maximum filter; rank $k = m$ is the standard mean filter. The high and low value pixel filters are usually combined as follows:

$$\begin{aligned} y_i &= \left| \sum_{j=n}^{n(1-\beta)+1} \frac{1}{\beta n} x_{(j)} - \sum_{j=1}^{\beta n} \frac{1}{\beta n} x_{(j)} \right| \\ &= \frac{1}{\beta n} \left| \sum_{j=n}^{n(1-\beta)+1} x_{(j)} - \sum_{j=1}^{\beta n} x_{(j)} \right| \end{aligned} \quad (\text{A.29})$$

Figure A.12 shows the results of applying the order statistics filter to the rest ultrasound image from figure A.1, with rank $\beta = 6$ and $m = 23$.

A.2.2 Rotating kernel transformation

This algorithm uses the sticks approach from section A.1.9 with the filter output set to the difference between the maximum mean and minimum mean over all



Figure A.13: RKT result.

stick orientations:

$$I_{out}(x, y) = \max_{\theta} (\text{mean}_{\xi, \eta} [I(x - \xi, y - \eta) s_{\theta}(\xi, \eta)]) - \min_{\theta} (\text{mean}_{\xi, \eta} [I(x - \xi, y - \eta) s_{\theta}(\xi, \eta)]) \quad (\text{A.30})$$

This is equivalent to the constant *rotating kernel transformation* (RKT) proposed by Lee and Rhodes [126]. Other RKT may be obtained by using grey-level sticks rather than binary, e.g. with triangular or Gaussian contours along the length of the stick.

Figure A.13 shows the result of applying the RKT filter to the test ultrasound image of figure A.1 using size 23 stick masks of width 3 pixels.

A.2.3 Ratio of averages filter

Bovik [22] presents a filter for speckle suppression and edge detection known as the *ratio of averages filter* (ROAF). Bovik discusses use of the Laplacian of Gaussian (LOG) filter (Section 4.3.1) for detecting edges in speckle images and notes the large number of zero-crossings produced by speckle, rather than edges. A common approach to selecting the true edges is to apply gradient measures to the LOG filtered image. However, the variation in signal-to-noise ratio (SNR) inherent in the image means that the gradient statistics will vary across the image, making the edge strength scheme unreliable [22]. Bovik's approach is based on a comparison of the mean grey-level in regions either side of the candidate pixel.

Zaman and Maloney [232] present a simplified form of the ROAF. For each candidate pixel, the means are calculated of pairs of neighbouring regions in the n by n neighbourhood centred at the pixel of interest. The neighbouring region pairs are based on four directions such that the boundaries between them run




Figure A.14: ROAF result.

W-E, N-S, NE-SW, NW-SE (where N,E,S,W are the usual compass directions), and the regions are contiguous.

For each pair of regions the ratio measure R is defined as:

$$R_i = \min \left(\frac{P_i}{Q_i}, \frac{Q_i}{P_i} \right) \quad (\text{A.31})$$

where P, Q are the means of the two regions and $i = 1, \dots, 4$ are the four pairs.

The filter output is set to the minimum of the four ratios:

$$I_{out} = \min (R_1, R_2, R_3, R_4) \quad (\text{A.32})$$

Figure A.14 shows the result of applying the simplified ROAF to the test ultrasound image of figure A.1, with $n = 23$.