


Image Analysis for Patient Management in Colorectal Cancer


Sarah Louise Bond
Keble College



Supervisor: Professor Sir J. M. Brady

Wolfson Medical Vision Laboratory
Department of Engineering Science
University of Oxford

A thesis submitted for the degree
Doctor of Philosophy
Trinity Term 2006



This thesis is submitted to the Department of Engineering Science, University of Oxford, for the degree of Doctor of Philosophy. The thesis is entirely my own work, and except where otherwise stated, describes my own research.

Sarah Louise Bond, Keble College, Oxford.

Image Analysis for Patient Management in Colorectal Cancer

Sarah Louise Bond

Abstract

submitted for the degree of
Doctor of Philosophy
Trinity Term 2006

Colorectal cancer is the second most common form of cancer in the western world and kills over 400,000 people each year worldwide. Early and accurate diagnosis is critical in order to assess the risks and determine the best course of treatment.

This thesis introduces a number of image analysis methods that can be used in such patient management decisions. The first stage is to remove the distortion caused by field inhomogeneities from the images. Having done this, a range of segmentation techniques are introduced, to automatically find the colorectum and mesorectum. They can then be visualised in three dimensions. A search for the lymph nodes is also implemented, and finds small regions within the mesorectum, then classifies them. A likelihood is assigned to each region that looks like a node, as to whether it is in fact likely to be cancerous.

Another major factor in patient management decisions is the changes that have occurred during chemo/radiotherapy, prior to surgery. We have developed a non-rigid registration method that can be used to compare these sets of images in a quantitative, accurate and useful way. We tested a variety of state-of-the-art, 'generic' non-rigid registration algorithms and found that they almost always fail to register the image sets, due to the large scale deformations that have occurred over the course of treatment. This thesis describes two key ways in which these algorithms can be improved by incorporating anatomical knowledge, and physiological knowledge.

Firstly, we set out a method of finding and representing the anatomical features in the images that are always present, such as the bones, the bladder and the colorectum. These representations are then used to constrain the general algorithms. This increases the robustness, such that the algorithms ran successfully on all datasets tested, that is, they provided a clinically useful registration result.

Secondly, we incorporate knowledge of the physiology, or how we expect the anatomy to change due to treatment. We can represent these changes using the Jacobian of the deformation, which describes the local size and type of change. This is used to regularise the registration, and can be incorporated simultaneously with the iterations of the registration. The final result is an accurate and robust registration result that is clinically useful for finding corresponding features on pre- and post-treatment datasets

Acknowledgements

Firstly, I would like to thank my supervisor Prof. Sir Michael Brady, for his direction and support throughout my DPhil. He has been constantly enthusiastic about this work, and inspired me with many ideas!

I would also like to thank the clinicians at the Oxford Radcliffe Hospitals, in particular Fergus Gleeson, Neil Mortensen, Helen Bungay and Anthony McIntyre, for their patience in explaining the medical aspects of the project, and their continued support throughout. I am especially thankful for the time they took helping me to validate the work in this thesis.

Thanks to the members of the Wolfson Medical Vision Laboratory in Oxford. Many have provided helpful hints and criticisms on this thesis and throughout the project, especially Alison Noble, Niranjana Joshi, Olivier Noterdaeme, Matt Kelly, Jeong-Gyoo Kim, and Styliani Petroudi. Thanks must also go to all my many friends and colleagues in the lab over the last 3 years: lab lunches, 'cake club' and trips to G&D's for ice-cream were always a welcome distraction while I was writing up.

Finally, thanks to my Mum and Dad, and the rest of my family, for their love and encouragement throughout.

Contents

Abstract	ii
Acknowledgements	iii
Table of Contents	iv
List of Figures	vii
List of Tables	xiii
1 Introduction	1
1.1 The Medical Problem	2
1.1.1 How a Colorectal tumour is staged	3
1.1.2 Current MRI practice in the Oxford Hospitals	7
1.1.3 The Patient Journey	10
1.2 The Computer Vision Problem	12
1.2.1 Image Analysis Methods for Lower Gastro-Intestinal region . .	14
1.3 Contribution	15
2 MRI Technology	17
2.1 Background	17
2.2 MRI Basics	18
2.3 Image Formation	21
2.4 Image Artefacts	25
2.4.1 Gradient-Related Artefacts	25
2.4.2 Patient-Related Artefacts	26
2.4.3 Image Processing Related Artefacts	26
2.4.4 RF Related Artefacts	28
2.5 Bias Field Correction	28
2.5.1 Guillemaud and Brady method	30
2.5.2 Styner method	33
2.5.3 Implementation	35
2.5.4 Results and Discussion	37
2.5.5 Summary	42

3	Anatomy Segmentation	44
3.1	Active Contours	45
3.1.1	Theory	45
3.1.2	Implementation	49
3.1.3	Results and Discussion	52
3.1.4	Summary	56
3.2	Shape Models	57
3.2.1	Theory	58
3.2.2	Implementation	61
3.2.3	Results and Discussion	67
3.2.4	Summary	75
4	Node Detection and Analysis	77
4.1	Anti-Geometric Diffusion for Image Segmentation	79
4.1.1	Anti-Geometric Diffusion and the Heat Flow Equation	79
4.1.2	Classification	81
4.1.3	Region Merging	81
4.1.4	Shape Criteria	83
4.1.5	Two Dimensional versus Three Dimensional Implementation	84
4.2	Node detection	85
4.2.1	Classification Criteria	85
4.2.2	Training and Likelihood graphs	86
4.2.3	Node Analysis	92
4.2.4	Visualisation	92
4.2.5	Results and Discussion	93
4.2.6	Summary	97
5	Registration	98
5.1	Validation	100
5.1.1	Statistical Measures	100
5.1.2	Simulations and Phantoms	101
5.1.3	Fiducials	101
5.1.4	Validation for pre- and post-treatment image alignment	102
5.2	Affine and Rigid Registration Methods	103
5.2.1	Theory	103
5.2.2	Applying Affine Registration to Colorectal Cancer Images	107
5.2.3	Summary	109
5.3	Non-Rigid Registration Methods	111
5.3.1	Theory	111
5.3.2	Applying Non-rigid Registration to Colorectal Cancer Images	122
5.3.3	Summary	129
5.4	Scale Saliency	130

6	Knowledge Representation	133
6.1	Incorporation of Knowledge	134
6.2	Relational Representation of Colorectal Cancer MR Images	142
6.2.1	Hips	142
6.2.2	Coccyx	145
6.2.3	Bladder	146
6.2.4	Bone structures	146
6.2.5	Colorectum	147
6.2.6	Mesorectum	147
6.3	Results	147
6.4	Summary	149
7	Mobilising Shape Knowledge	151
7.1	Shape Representation	151
7.1.1	Finding Landmark Points	152
7.1.2	Thin Plate Splines	152
7.2	B-splines and Mutual Information	154
7.2.1	Method	154
7.2.2	Results	155
7.3	Adaptive Registration	158
7.3.1	Method	158
7.3.2	Results	158
7.4	Fluid Registration	160
7.4.1	Method	160
7.4.2	Results	160
7.5	RevealMVS TM	162
7.5.1	Method	162
7.5.2	Results	162
7.6	Conclusions	166
8	Mobilising Physiological Knowledge	169
8.1	Physiological Knowledge	170
8.1.1	Principal Warps of the Deformation Field	170
8.1.2	Vector Field Analysis	173
8.1.3	Jacobian of the Colorectal Cancer Deformation Field	177
8.2	Simultaneous Methods	182
8.2.1	Simultaneous Segmentation and Registration	182
8.2.2	Simultaneous Landmark Detection and Registration	183
8.2.3	Implementation	186
8.2.4	Results and Discussion	188
8.2.5	Conclusions	190

9	Conclusions	193
9.1	Summary	193
9.2	Future Work	195
9.2.1	Dynamic Contrast-Enhanced MRI	196
9.2.2	Validation	199
9.2.3	Local Excision Surgery	202
A	Imaging Orientations	205
	Bibliography	207

List of Figures

1.1	Axial MR image showing anatomy of colorectal region.	4
1.2	Sagittal MR image showing anatomy of colorectal region.	4
1.3	T-staging of Colorectal Cancer.	6
1.4	A typical radiologist's staging - 'Image 4.14' (Axial T2-weighted). . .	9
1.5	A typical patient journey.	10
2.1	Precession of a vector in a static magnetic field.	19
2.2	Flipping and relaxation to equilibrium.	20
2.3	Decay and Recovery of B_1 and B_0 fields giving T2 and T1 relaxation times.	21
2.4	TR and TE times for T1 and T2 relaxations, indicating how a good contrast is achieved. The top graph shows how a short TR enhances the T1 contrast. The bottom graph shows that a longer TE can enhance the T2 contrast.	22
2.5	An example of a signal timing diagram for a spin echo sequence. . . .	23
2.6	An example of k-space and its corresponding MR image.	24
2.7	Effect of field inhomogeneities on a colorectal MR image.	25
2.8	Effect of aliasing on a colorectal MR image.	26
2.9	Effect of field chemical shift on a colorectal MR image.	27
2.10	Effect of partial volume effect on a colorectal MR image.	28
2.11	Diagram of Legendre polynomials for two dimensions (0th to 4th degree). .	33
2.12	Bias Removal.	36
2.13	Flow of anisotropic diffusion.	37
2.14	Before and after anisotropic diffusion.	38
2.15	Estimation of the Bias Field using the Styner method visualised in two dimensions and three dimensions.	39
2.16	Axial image before and after bias removal.	39
2.17	Bias removed images.	41
2.18	Left hand images are the original breast MR images, and right hand images are the corresponding bias corrected images.	42
3.1	NURBS algorithm.	48
3.2	Kalman Filter Algorithm.	51

3.3	Active contour outlining the rectum on the first image in a series of axial images.	52
3.4	Contours propagated through the set of axial images using a Kalman filter.	53
3.5	A slice from the data-set shown in Fig. 3.4 showing the partial volume effect.	54
3.6	A slice from the data-set shown in Fig. 3.4 showing the tracking error.	55
3.7	Three dimensional visualisation of the lower colon/rectum using the contours found automatically.	55
3.8	A typical axial slice showing the tumour outlined on the right hand image.	57
3.9	Three dimensional model showing a lymph node.	57
3.10	Outlined mesorectal fascia.	62
3.11	Frame of reference of MR image.	63
3.12	Tracking mesorectal fascia using a snake.	64
3.13	Fourier surface (blue) fit to training set (green).	68
3.14	Mean shape of the mesorectum calculated from 10 images, showing the Fourier surface in blue, and the mean of the training sets in green.	68
3.15	First estimate of the mesorectum (mean shape).	69
3.16	Estimated mesorectum (yellow), and the manual segmentation (red) for dataset 1.	70
3.17	Estimated mesorectum (yellow), and the manual segmentation (red) for dataset 5.	70
3.18	Estimated mesorectum (yellow), and the manual segmentation (red) for dataset 8.	71
3.19	Estimated mesorectum (yellow), and the manual segmentation (red) for dataset 10.	71
3.20	Volume plot of the estimated mesorectum (yellow) for dataset 10.	72
3.21	Volume plot of the estimated mesorectum (yellow) for dataset 8.	72
3.22	Dataset 8 showing the raw data (top row), and the estimated mesorectal fascia (bottom row).	73
3.23	Error due to unclear boundary.	74
3.24	Error in fitting to colorectum.	74
3.25	Instability due to brightness of bladder.	75
4.1	The area of the mesorectum surrounding the colorectum, where the lymph nodes are found.	78
4.2	Normal and tangential directions of iso-intensity contour.	80
4.3	Initial image and oversegmentation using 3 classes.	82
4.4	Region merging to give 70 and 500 regions.	83
4.5	Initial image showing nodes, and the segmentation with the nodes preserved, indicated by the arrows.	84
4.6	Problem with three dimensional implementation.	85

4.7	Probability that a node is involved given its size.	87
4.8	Probability that a node is involved given its mean intensity.	88
4.9	Nodes with heterogeneous signals and irregular edges.	89
4.10	Probability that a node is involved given its overall heterogeneity. . .	89
4.11	Visibility of nodes and blood vessels on a series of slices.	90
4.12	Probability that a region is a node given the number of slices it is visible on.	91
4.13	Probability that a region is a node given its eccentricity.	91
4.14	Colourmap of likelihoods of involved lymph nodes.	93
4.15	Top left - Original image; Top right - Bias removed image; Bottom left - all regions plotted according to their probability of being involved nodes; Bottom right - Regions with a 90% probability of being involved.	94
4.16	Left - all regions plotted according to their probability of being involved nodes; Right - Regions with a 90% probability of being involved. . . .	95
4.17	A typical three dimensional representation of the blood vessels in the mesorectum.	95
4.18	A three dimensional representation of the small regions in the mesorec- tum, coloured according to their probabilities of being an involved lymph node.	96
4.19	A three dimensional representation of the small regions in the mesorec- tum that have a 90% probability of being an involved lymph node. . .	97
5.1	Equivalent slices pre- and post-therapy.	99
5.2	Correspondance of lymph nodes in pre- and post-treatment images. .	103
5.3	Rigidly Moving Anatomical Features in Colorectal Images	107
5.4	Rigidly Moving Landmarks in Colorectal Images	108
5.5	Affine Registration using RevealMVS, showing the target image at the top, and the target image with the overlaid registered source image at the bottom.	110
5.6	Result of Rueckert's algorithm.	123
5.7	Local mismatch measure.	124
5.8	Areas of maximum misregistration.	124
5.9	Affinely Registered Intensity Difference Image.	126
5.10	Fluid Registered Intensity Difference Image.	126
5.11	Fluid Registered Intensity Difference Image Using Normalised Mutual Information.	127
5.12	Image registered using RevealMVS.	127
5.13	Intensity Difference Image Using RevealMVS.	128
5.14	Intensity difference image after local affine registration using ITK. . .	128
5.15	Illustration of salience across scale and space.	132
6.1	Simulated images of before and after treatment, that need to be regis- tered.	135

6.2	Registered image using B-splines and Mutual Information, showing registered source image (left) and difference image (right).	135
6.3	Simulated images of before and after treatment without the ‘noisy’ surrounds.	136
6.4	Registered central regions of the images using B-splines and Mutual Information, showing registered source image (left) and difference image (right).	136
6.5	Registration of whole images using just central information, showing registered source image (left) and difference image (right).	137
6.6	Real images before and after treatment, and their segmentations. . .	138
6.7	Registration of segmented images, showing registered source image (left) and difference image (right).	139
6.8	Registration of real data using same transform as that calculated for segmented images, showing registered source image (left) and difference image (right).	139
6.9	Mesorectal area showing disappearance of lymph node and change in intensity of tumour in the colorectum.	140
6.10	Directed Graph Model for Colorectal Anatomy.	143
6.11	Search regions for key features of interest within colorectal cancer MR images.	144
6.12	Illustration of multiscale approach.	145
6.13	The left hand column shows the original two dimensional slices and the right hand column show the corresponding anatomies.	150
7.1	Key landmark points on the relational representation of the colorectal anatomy.	152
7.2	Results using B-splines, Mutual Information and shape information for Patient 1: Top left – Target image, Top right – Initially aligned source image, Bottom left – Final warped source image	156
7.3	Results using B-splines, Mutual Information and shape information for Patient 2: Top left – Target image, Top right – Initially aligned source image, Bottom left – Final warped source image	157
7.4	Results using adaptive registration and shape information for Patient 7: Top left – Target image with points of maximal misalignment, Top right – Source image showing changes in points of maximal misalignment, Bottom left – Source image showing aligned points, Bottom right – Warped source image.	159
7.5	Results using fluid registration and shape information for Patient 4: Left – Target image, Right – Warped source image.	161
7.6	Results using fluid registration and shape information for Patient 7: Left – Target image, Right – Warped source image.	161

7.7	Results using RevealMVS TM and shape information for patient 2: Top left – Target image, Top right – Source image, Bottom left – Warped Source image with overlaid warp field.	163
7.8	Results using RevealMVS TM and shape information for patient 6: Top left – Target image with overlaid warp field, Top right – Source image, Bottom left – Warped Source image with overlaid warp field.	165
7.9	Spatial distribution of the errors measured from the centre of the region of interest, for 8 datasets, using RevealMVS TM	166
7.10	Qualitative comparison of registration of data-set 7 using each of the four methods. Top left: Target image, Top right: Source image, Middle left: Registration result using B-Splines, Middle right: Registration result using Adaptive registration, Bottom left: Registration result using Fluid registration, Bottom right: Registration result using RevealMVS TM	167
8.1	The 76 manually placed landmarks on the pre- and post-treatment images, and the corresponding registered image.	171
8.2	An example of registration with RevealMVS TM showing the extent of the deformations.	172
8.3	The deformations in the x and y directions.	173
8.4	The principal warps of the deformation field.	174
8.5	The central section of $\log(\text{Jacobian})$ of a typical deformation of colorectal cancer images.	178
8.6	A cross section of $\log(\text{Jacobian})$ of a typical deformation of colorectal cancer images.	178
8.7	The Jacobian of the colorectal cancer images with the 76 manually placed landmarks.	179
8.8	The Jacobian of the colorectal cancer images with 67 of these manually placed landmarks.	179
8.9	The original colorectal cancer images with a selected group of 17 manually placed landmarks, and the corresponding Jacobian of the deformation just using these landmarks.	180
8.10	The Jacobian of the colorectal cancer images using a random selection of 26 of the manually placed landmarks	181
8.11	A visual comparison of Jacobians using manually selected landmark points and automatically generated points.	181
8.12	Pseudo-code of the Simultaneous Landmark Detection and Registration algorithm	187
8.13	Effects of changing the sub-block and step size on the results of the algorithm.	188
8.14	Effects of changing the values of λ on the results of the algorithm, with the dashed line indicating the optimum value of λ	189
8.15	Typical distribution of landmarks on an image.	189

8.16	Reliability and Precision of non-rigid registration algorithms with and without the use of shape knowledge.	192
9.1	Idealised signal time curves for Gd-DTPA uptake.	197
9.2	Pharmacokinetic model setup.	198
9.3	Three consecutive macro slices.	200
9.4	Aligned macro slices viewed as a 3-D stack, courtesy of Chen [70]. . .	201
9.5	A typical trans-anal ultrasound image.	204
A.1	MR Imaging Directions.	205
A.2	Sagittal MRI showing the direction of the off-plane axial images, perpendicular to the rectum.	206

List of Tables

1.1	TNM staging of rectal cancer.	5
1.2	Staging of cancer.	6
2.1	Comparison of standard deviation/mean, averaged over all tissue classes.	40
3.1	Results of finding the mesorectum with the shape model for each dataset.	73
5.1	Node Misalignment (mm) using Affine Registration.	109
5.2	Node Misalignment Comparison.	129
7.1	Node Misalignment using B-Splines, Mutual Information and Shape Information.	155
7.2	Node Misalignment using Adaptive registration and Shape Information.	158
7.3	Node Misalignment using Fluid Flow registration and Shape Information.	160
7.4	Node Misalignment using RevealMVS TM and Shape Information. . . .	164
7.5	Node Misalignment using Non-rigid registration techniques with and without shape information.	168
8.1	Accuracy of Simultaneous Landmark Detection and Registration Algorithm.	191
8.2	Comparison of the node misalignments (mm) for each of the 10 datasets, running the Park-Meyer algorithm on its own, the Park-Meyer algorithm with shape initialisation, and finally using the simultaneous algorithm.	191

Chapter 1

Introduction

Medical Image Processing and Analysis is an essential stage in aiding the diagnosis and treatment planning (patient management) of cancer. This thesis applies and further develops some of the leading computer vision techniques in order to detect and diagnose colorectal cancer, and to help clinicians assess the success of treatment, both chemo/radiotherapy and surgery.

This includes automatically finding regions of anatomical and clinical interest, such as the colorectum and mesorectum, where the primary colorectal tumour is located. We also describe how we can automatically detect the lymph nodes, and discriminate them from blood vessels, which often appear almost identical in two-dimensional images. The assessment of lymph nodes is of primary importance in the management of colorectal cancer because they often contain cancer cells. If such lymph nodes are close to the surgical resection margin, there is a much greater risk that the tumour will remain, recur and/or metastasise. We outline how we can use computer vision and imaging techniques to provide an estimate of the diagnosis of these nodes.

We also develop a system for the quantitative comparison of pre- and post-neo-adjuvant chemo/radiotherapy MRI scans. This too is important, as it enables clinicians to assess the changes due to treatment, before they proceed to surgery. From a computer vision standpoint, the key challenge is that the changes are large scale and non-rigid, both due to treatment, and also due to the fact that the images are typically

taken six months apart. The patient posture is unlikely to be exactly the same, nor is the extent, for example, of bladder distension. Hence this is a difficult comparison and alignment problem. We incorporate anatomical and physiological knowledge in order to solve this alignment and provide a robust and accurate registration.

It is obviously important that the methods we employ throughout this thesis are accurate, but robustness is also a key factor to consider for medical imaging. A clinician will not use a piece of software that gives a gross error in its calculation of alignment for example. We therefore require and obtain 100% robustness in our final registration algorithms.

1.1 The Medical Problem

Cancer is already the second most common cause of death in the Western world, indeed it is expected that it will soon overtake cardiovascular diseases as the most common cause [61]. It is a major problem facing the UK, with approximately 16,000 people dying from the disease each year. More specifically, colorectal cancer is the second most common form of cancer death in the UK, accounting for over 10% of all cancer deaths [77]. The problem extends beyond the UK, with two thirds of a million people newly diagnosed with colorectal cancer each year worldwide. Each year there are approximately 400,000 colorectal cancer related deaths worldwide. It affects men (11%) and women (12%) approximately equally [25]. Even more grim is the fact that the patient survival rate at 5 years is only 50%, largely because of metastatic spread, primarily to the liver and pelvis. This reinforces the point made earlier about assessing lymph node status. Currently, surgery is the only curative therapy: current surgical practice is traumatic, and local excision surgery has been slow to take off. Chemotherapy is currently used in 65% of cases, reinforcing the need for quantitative comparison between pre- and post-chemotherapy images.

It is critical that cancer is diagnosed early and accurately in order to assess the risks and determine the best course of treatment. The staging of cancers (that is the process of determining the size and spread) informs clinical decisions about how

to proceed with treatment. It gives an indication of the patient's chance of survival as well as the appropriate forms of management of the cancer. Magnetic Resonance Imaging (MRI) is widely used to determine the local extent of colorectal cancer, after a cancer has been diagnosed. Prior to performing the MRI scan, the presence of a cancer in the colorectal region is verified using one or more of the following methods:

1. Endoscopy,
2. Barium Enema, and
3. Computed Tomography (CT) scan.

The MR scan is then performed in order to stage the cancer, to ascertain how far the tumour has spread locally and to identify any metastases (spreading to distal sites). However, current techniques need to be improved because nodal metastases in particular are often under or over reported. This means that the best course of treatment is not always clear. In order to appreciate fully the issues faced by clinicians when staging colorectal cancer, we have, throughout the course of this thesis, participated in the weekly multi-disciplinary colorectal cancer meetings, to observe the current practice in the Oxford hospitals (John Radcliffe, Churchill). We also consulted the key clinicians, in particular the radiologists, on a regular basis. Details of the current procedures are given below and further details can be found in [77].

1.1.1 How a Colorectal tumour is staged

A brief introduction to the anatomy of the region of interest can be seen in the Figs 1.1 and 1.2. They show the colorectum, which is where the primary tumour is located. It is surrounded by the mesorectum, which consists primarily of a layer of fat, the boundary of which is the mesorectal fascia. Also visible in the images are the hips, the spine, the coccyx, and the bladder. An explanation of the image orientations used to form these images can be found in the Appendix.




Figure 1.1: Axial MR image showing anatomy of colorectal region.




Figure 1.2: Sagittal MR image showing anatomy of colorectal region.

Table 1.1: TNM staging of rectal cancer.

Stage	Definition
T1	Tumour invades submucosa
T2	Tumour invades circular or longitudinal muscle layers
T3	Tumour invades through muscularis propria subserosa or into non-peritonealised perirectal tissues
T4	Tumour directly invades other organs or structures or perforates visceral peritoneum
N0	No regional lymph node metastasis
N1	Metastases in one to three perirectal lymph nodes
N2	Metastases in four or more perirectal lymph nodes
M0	No distant metastases
M1	Distant metastases present

Lymph nodes and blood vessels are usually found in the mesorectum, which is the area of fat surrounding the rectum. For patient management, the key questions are (1) how big is the tumour, and how far has it extended into the surrounding fat?; (2) are there any involved (that is, containing cancer cells) lymph nodes?; and (3) has the primary tumour in the colorectum metastasised? Conventionally, the tumour stage is determined using three major criteria:

1. T - Primary tumour.
2. N - Regional Lymph Nodes.
3. M - Metastases.

These stages are summarised in Table 1.1.

The primary tumour in rectal cancer can be staged using the MR image, as is shown in Fig. 1.3. The TNM values can then be used to specify the overall stage of the cancer as shown in Table 1.2.

Lymph nodes have to be analysed closely in order to determine whether or not they are cancerous. However, to date, there are no standards for guiding assessment: different criteria for the assessment of nodal involvement are utilised by different radiologists. It is known, for example, that size is a very poor discriminator. It is also




Figure 1.3: T-staging of Colorectal Cancer.

Table 1.2: Staging of cancer.

Stage	TNM (E.g.)	Prognosis
1	$T_1N_0M_0$	Tumour is confined to primary organ therefore is usually operable.
2	$T_2N_1M_0$	Local spread. Operable but with risk of further spread of disease.
3	$T_3N_2M_0$	May not be operable and hence may require a combination of different treatments.
4	$T_4N_2M_1$	Distant metastases and primary site may not be operable.

suggested that if a node demonstrates a chemical shift artefact (potentially visible on MRI) then it is likely to be benign. Conversely, if there is evidence of contour irregularity or signal heterogeneity then this is suggestive of malignancy. Because of the limited spatial resolution of MRI, typically 0.4mm x 0.4mm x 3mm voxels, there are inevitably problems judging whether or not very small nodes, particularly those that are only 1-2mm in diameter, are cancerous; at this size it is very difficult to tell if the signal is heterogeneous. If several cancerous lymph nodes are found and the cancer is locally extensive (say T4), then there is a high risk that it has spread

from the colorectal area into other organs of the body such as the liver. (The liver is routinely imaged using CT). Evidently, the staging of colorectal cancer needs to be rigorous and precise; but there do not currently exist systems that are able to do this. When post-therapy scans are taken, it is critical to be able to determine whether or not chemo/radiotherapy has been successful, and whether or not to continue treatment with the aim of down-staging the tumour further. If treatment has not been successful, then the tumour may well have grown, instead of decreasing in size, and the tumour may need to be re-staged. Post-therapy scans are a relatively new idea. They are done so that the surgeon can be informed of the current stage of the cancer before starting the operation. It cannot be assumed that down-staging (by chemo/radiotherapy) will always work. In fact in 50% of cases there is little effect. Much work on staging colorectal cancers has been done by the collaborating radiologists at the Royal Marsden and Oxford hospitals, Dr. Gina Brown *et al.* [14], [33], [16], and Dr. Fergus Gleeson [34].

1.1.2 Current MRI practice in the Oxford Hospitals

After a tumour has been detected, the patient is normally referred to have an MR scan. This enables the extent and size of the tumour (described previously) to be determined from the images. The usual method of performing this staging is as follows. There are four main sets of images that are taken, and these are all T2-weighted.

1. Initial thick section ‘localiser’ images are taken, with five images in each of the axial, coronal and sagittal planes ($TE = 1.6\text{ms}$, $TR = 58\text{ms}$, $\alpha = 90^\circ$, slice thickness = 5mm). Each slice is 256 x 256 pixels, and the pixel size is 1.875mm x 1.875mm. These images ensure that the correct area of the body is being imaged and to enable the radiographer to select from where the more detailed images should be taken.
2. Thin section coronal images are taken in the area of the tumour and this is used in staging the tumour ($TE = 95.5\text{ms}$, $TR = 3000\text{ms}$, $\alpha = 90^\circ$, slice

thickness = 3.5mm). Each slice is 512 x 512 pixels, and the pixel size is 0.43mm x 0.43mm.

3. Thick section sagittal images are also taken ($TE = 86\text{ms}$, $TR = 7500\text{ms}$, $\alpha = 90$ deg, slice thickness = 5mm). Each slice is 512 x 512 pixels, and the pixel size is 0.47mm x 0.47mm. This is an optional, though frequently used, precautionary measure and is not vital in staging the cancer. Sometimes a smoothing function is used on the sagittal images. This evens out the density and applies a noise reduction algorithm to make the images appear clearer.
4. Thin section off-plane axial small field of view images are always taken ($TE = 120\text{ms}$, $TR = 6500\text{ms}$, $\alpha = 90$ deg, slice thickness = 3mm). Each slice is 512 x 512 pixels, and the pixel size is 0.39mm x 0.39mm. They are off-plane since the rectum is at an angle and hence to best stage the cancer the images are taken perpendicular to it, as shown in Appendix A. These images are used predominantly to determine the size and stage of the cancer.

Any other images that are taken are extra and usually just the standard 4 sets are taken. However some centres choose to take thick axial section T1-weighted images that go higher up in the body. These help when looking for nodal disease that may have spread further up the body.

An example of the images obtained using the current practice is shown in Fig. 1.4, along with the report by the radiologist, indicating the extent of the cancer. The radiologists' reports are purely based on the MR images and they view these images either using a printed slide or directly on the computer screen. The report is shown below:




Figure 1.4: A typical radiologist's staging - 'Image 4.14' (Axial T2-weighted).

MRI PELVIS

Indication: Rectal cancer; for local staging.

Coronal and axial T2 images obtained.

There is an eccentric mass present within the distal rectum extending from approximately 10 o'clock - 5 o'clock. There is effacement of the muscularis propria from 12 o'clock - 2 o'clock and there is extension of tissue, isointense with tumour into the mesorectal fat anteriorly at 2 o'clock superiorly (image 4.14) and also between 11-2 o'clock more inferiorly (images 4.16-18). The appearances suggest early T3 disease. The minimum separation between tumour and the mesorectal fascia is 8 mm superiorly (image 4.14) and 4 mm inferiorly (images 4.17-18). There is a 5 mm mesorectal lymph node laterally on the left (image 4.14) and other smaller posteriorly placed mesorectal nodes. At this size they are indeterminate.

There are many ways in which medical imaging techniques can be used to improve the staging of cancers. One obvious example can be seen by noting the brightness variation from the bottom to the top of the image: this is due to the 'bias field', an MRI artefact that is discussed in Chapter 2. This intensity deficiency in MR images can be removed, which then enables measurements of the data to be more

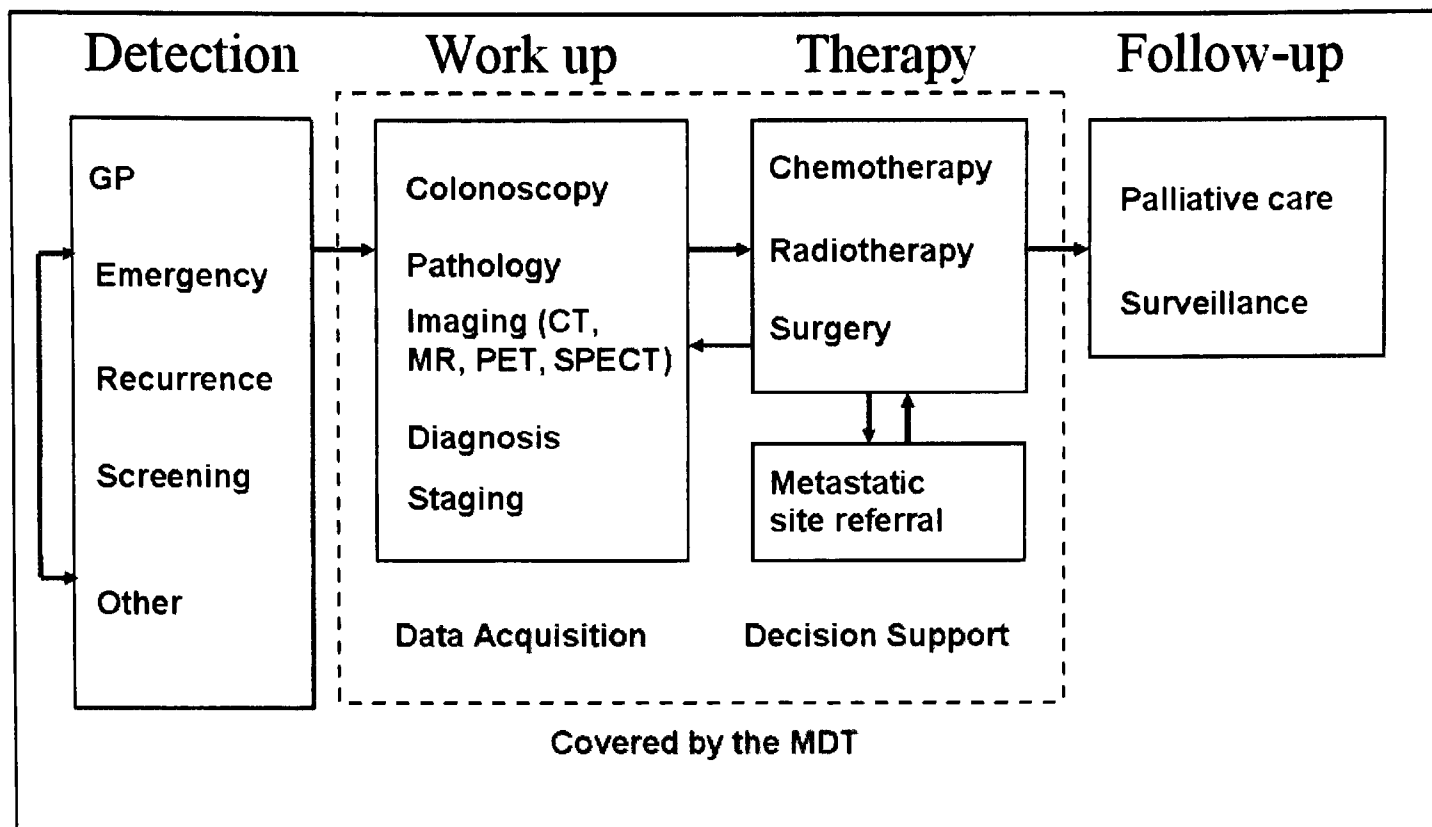


Figure 1.5: A typical patient journey.

quantitative. Another example is by automatically detecting and diagnosing the lymph nodes which radiologists find difficult to discriminate from blood vessels. This data can also be viewed in three dimensions.

1.1.3 The Patient Journey

Once the presence of a colorectal tumour has been confirmed and the scans analysed, the patient is discussed at a multidisciplinary team (MDT) meeting. In the John Radcliffe Hospital this is done as a weekly meeting, where the MR, CT and occasionally PET images are compared with the pathological data and other relevant information about the patients. This meeting is attended by radiologists, surgeons, oncologists, histopathologists and senior nursing staff. A representation of the set of typical patient journeys can be seen in Fig. 1.5

When a patient is first presented at this meeting a decision will be taken as what course of treatment the patient will proceed with. The most common options are:

1. Surgery - If the tumour is very small and localised, then the patient will often proceed straight to surgery. In such cases the patient is usually reasonably fit and the tumour is deemed unlikely to have metastasised. This surgery may be followed by a short course of chemo or radiotherapy, to remove any remaining tumour.
2. Pre-surgery treatment - If the tumour is more advanced, the patient will undergo a course of (hopefully) downstaging chemotherapy, radiotherapy, or both, prior to having surgery. The aim is to ensure that the tumour shrinks such that it is well within the margin that will be resected during surgery, and that any mesorectal lymph nodes are cleared.
3. Palliative Care - If a patient has a tumour that cannot be fully treated they will often have palliative care, which typically involves courses of chemotherapy or radiotherapy to slow the progression of the cancer and to alleviate pain during the remaining lifetime of the patient, which is expected to be short.

The majority of patients fall into the second category and have some level of downstaging treatment prior to surgery. If this is the case they will then have another series of MRI scans after treatment in order to assess the success of this treatment before the patient proceeds to the theatre.

Before surgery the clinicians need to plan exactly which section they are removing, the edge of which is referred to as the circumferential resection margin. In order to plan this they need to know exactly how far away the tumour and any involved lymph nodes are from this margin. The MERCURY project [72] was a recent study aimed at assessing this resection margin based on the pre-surgery MRI, by comparing the scans with the gold-standard histopathology post-surgery.

Neo-adjuvant chemotherapy and radiotherapy is successful, at least to a certain extent, in 50% of cases; indeed, in 10% of cases there is no residual tumour. Successful (albeit partially successful) response causes the T2-weighted MR signal to reduce in intensity. Any involved lymph nodes are also reduced in size. There is a predictable

density change that occurs in the mesorectal fat: it becomes much ‘streakier’ in appearance in the post-therapy image, largely due to fibrous tissue. The volume of mesorectal fat also decreases, though to varying degrees. The wall of the colorectum often thickens and becomes adematous despite the overall reduction in size.

Finally, we note that even in the absence of cancer this area of the body is particularly prone to localised displacement, not least because of bowel and bladder movement. For all of these reasons, the differences between the images pre- and post-therapy are most often substantial, to the point where even skilled clinicians have difficulty evaluating patient response to therapy. This motivates the need for image registration to align the anatomy and tumour in the pre- and post-therapy datasets. The expected changes are substantial, and the clinical descriptions are qualitative, however they do refer to anatomical structures which can be identified.

Most of the diagnosis and comparison stages predominantly use the off-plane axial, small field of view, T2 weighted ($TE = 90\text{-}120\text{ms}$, $TR = 3500\text{-}6500\text{ms}$, $\alpha = 90^\circ$, slice thickness = 3mm) MR images, and these are the images used throughout this thesis. The images were taken on a 1.5T GE machine at the Radiology Department of the Churchill Hospital, Oxford. The datasets are 3 dimensional, consisting of 20-30 slices, of 512 x 512 pixels each. Each voxel is 0.4mm x 0.4mm x 3mm.

1.2 The Computer Vision Problem

Some of the computer vision problems faced to aid clinical application of these images are clear. Bias field correction for example is a well researched area and many algorithms exist to correct for this effect. There are also many methods available for segmentation and registration.

Many computer vision techniques are applied to medical images in such a way as to find ‘generic’ algorithms. That is one algorithm is developed with the intention of solving many medical imaging problems. However, we show that this approach does not allow for the large variation in patients and disease. Every case is different and we need to account for those differences in our analysis. We need to know what it is

we are trying to find in the images, i.e. what the anatomy is likely to look like, and where different features are in relation to one another. We also need a clear idea of what we are trying to find, i.e. which parts of the images do we need to align in order to quantify the changes due to chemo/radiotherapy. As such, knowledge needs to be added to the generic algorithm.

Robustness is also a key aspect in this field. Failure to detect disease has obvious consequences, but failure to detect lack of disease can result in patients, who are often very frail, being put through unnecessary and potentially fatal treatment or surgery. Also, a clinician is unlikely to use an algorithm or piece of software that gives an obviously wrong result, so we need to develop methods that give sensible results 100% of the time. A clinician's confidence in such a tool needs to be built up slowly, but can be totally eroded by a bad result. Therefore, we need to develop both accurate and robust methods to analyse these images, and we do this by incorporating anatomical and physiological knowledge.

The previous section outlined some of the clinical challenges we face. Some of the constraints we may then need to apply to our algorithms, particularly in the alignment phase are given below:

1. Tumours and lymph nodes are restricted to the colorectal and mesorectal regions, which are the only regions of clinical interest.
2. Chemotherapy and radiotherapy may reduce not only the size of the tumour but also the sizes of lymph nodes, in certain cases to the point where they become undetectable.
3. In addition, oedema occurs in the colorectal and mesorectal regions, so that they reduce in size, even in cases of partial response, and almost never increase in size.
4. The appearance of mesorectal tissue changes predictably to become more streaky (that is, it has a textured appearance that is primarily locally uni-directional).

5. The wall of the colorectum often thickens and becomes adenomatous.
6. Large changes in shape and size of bladder and rectum are to be expected, even for the same patient, and across the population.

Point 2 means that whereas we hope, for the patients sake, that the pre- and post-chemotherapy images are *not* homologous, this is an assumption that is incorporated into most ‘generic’ algorithms.

Note also that only patients who have presented with symptoms (for example, painful bowel movements, swelling, or general discomfort in the lower back), whose primary care physician has determined that they need specialist treatment, and who have been CT scanned to confirm the presence of a colorectal cancer, are given an MRI scan. On the basis of this scan, the decision whether to proceed straight to surgery or to undergo a course of chemo/radiotherapy will be made. The imaged population is a hugely biased sample and their images pre- and post-therapy are far from normal.

As such the medical challenges are completely interleaved into the computer vision, image processing tasks and we consider both throughout this thesis.

1.2.1 Image Analysis Methods for Lower Gastro-Intestinal region

Little known work has been done to date in computer vision to aid the problems associated with Colorectal cancer MR imaging. Work has been done in the radiology field looking at the circumferential resection margins and the evaluation of lymph nodes using both conventional MR techniques as well as contrast agents such as Gd-DTPA or USPIO’s [52, 33, 16]. However no quantitative image analysis has been performed on these images.

Segmentation and Registration methods for medical images are rarely applied to this region of the body due to its mobility and huge variance in shape. There are a couple of exceptions, for example Tsai et al. [86] applied a coupled shape model for segmentation to the bladder and lower gastro-intestinal region, but this was only

applied to a small number of healthy cases. Colorectal cancer sufferers often have huge pathologies in the mesorectal region, and as such there are no 'normal' or 'typical' cases.

This problem was recognised by Christensen et al. [19] who attempted registration for the application of treatment planning using an anatomical template. This technique worked well, but required manual segmentation of the images prior to the registration. It required a skilled clinician to do this.

The majority of the computer vision work for colorectal cancers, to date, has focussed around CT colonography [81, 82]. This is a method most used to detect cancers in the first instance, as well as polyps that may potentially form malignancies. CT colonography cannot be used to determine the extent and spread of a tumour or identify the lymph nodes or circumferential resection margin used by surgeons.

1.3 Contribution

This thesis sets out a number of methods that can be used to detect anatomical features within the images, and highlight regions of interest such as possible involved lymph nodes. We also study how best to align images pre- and post-treatment (neoadjuvant chemo/radiotherapy), incorporating knowledge of the anatomical features and how we expect them to change. Finally, we interleave the anatomy detection and alignment to give an accurate and robust integrated registration algorithm.

Chapter 2 describes how MR images are formed, and the fundamental physics behind the NMR (Nuclear Magnetic Resonance) technology. We discuss some of the artefacts associated with MRI and go on to describe how we remove the bias field from the axial colorectal cancer images.

Chapter 3 outlines some of the initial image processing techniques we can apply to the colorectal MRI scans. We describe the segmentation of the colorectum using a combination of active contours and a Kalman filter, and finally the segmentation of the mesorectum using shape models.

Chapter 4 continues by describing the segmentation of the images using the

method of anti-geometric diffusion to segment out small regions that are likely to be lymph nodes. We describe how these regions are classified, in order to find an estimate as to whether each region is a node, and if so, whether it is likely to be involved, i.e. contain tumour.

We then look at the issue of registering the images pre- and post-treatment in Chapter 5. We describe a selection of state-of-the-art, ‘generic’ affine and non-rigid registration algorithms and how they behave when run on the colorectal MRI scans. Since the results are not good, we discuss why this is the case and what can be done to assist the performance of these algorithms. We conclude that application-specific knowledge pertaining to the images needs to be incorporated.

Chapter 6 outlines a method to find a relational representation of the anatomy within the images, using the knowledge we have about the shapes of the anatomical features and their positions relative to one another. This results in an accurate set of points on the images that correspond to features of anatomical interest within the images.

We then describe how this knowledge is incorporated into the ‘generic’ algorithms that were described in Chapter 5. Chapter 7 gives the results of the registration with this added knowledge and shows that the robustness is much improved.

Chapter 8 goes one step further and suggests that physiological knowledge, as well as anatomical shape knowledge, can be interleaved with the registration at each iteration to improve the accuracy and robustness further. This physiological knowledge is represented as a Jacobian regulariser. The results of this simultaneous algorithm are presented.

Finally in Chapter 9 we summarise the main results and outline some suggestions for future work.

Chapter 2

MRI Technology

This chapter will provide a brief introduction to the basic principles of MRI physics, and will describe how images are formed. We then discuss some of the problems that challenge MR imaging, particularly focussing on artefacts including the bias field. We describe some of the proposed methods for reducing the bias field from MRI after the images have been acquired. Finally, we implement a bias field removal algorithm for the colorectal MR images.

2.1 Background

Magnetic Resonance Imaging (MRI) has been widely used in radiology since it was first introduced into clinical practice in the 1970s. The technology is based on knowledge about Nuclear Magnetic Resonance (NMR), that is the idea of exploiting the magnetic properties of an atom's nucleus. The hydrogen nucleus, 1H , is particularly susceptible to this phenomenon [44], and it is found in the large quantities of water in the human body. The idea of aligning protons within a powerful external magnetic field and then perturbing this alignment using an electromagnetic field was discovered in 1946 by Felix Bloch and Edward Mills Purcell, and forms the basis of NMR.

MRI, formerly referred to as magnetic resonance tomography (MRT) or nuclear magnetic resonance (NMR), is a method used to visualise the internal anatomy of living organisms. It is primarily used to demonstrate pathological or other physiological or anatomical alterations of living tissues and is an increasingly commonly used

form of medical imaging. It produces high resolution images, that can be obtained in different planes. These images have good tissue contrast, and are produced without subjecting the patients to ionising radiation. The basics of MRI can be interpreted using both classical physics and quantum mechanics.

2.2 MRI Basics

When a patient goes for an MRI scan they are first placed into a large, often cylindrical magnet, which has a typical field strength of 1.5T. We denote this as the B_0 field. This field causes all the spinning charged particles (for example Hydrogen nuclei) within the patient's body to align with the direction of this field. When a spinning charged nuclei is placed into such a magnetic field it will also 'wobble' or precess about the axis of the field, Fig. 2.1. The rate at which it precesses is given by the Larmor equation:

$$\omega_0 = \gamma B_0, \quad (2.1)$$

where ω_0 is the angular precessional frequency, γ is the gyromagnetic ratio and B_0 is the field strength. The gyromagnetic ratio γ is a proportionality constant, fixed for the nucleus we are dealing with. For Hydrogen protons, $\gamma(\text{H}) = 42.6\text{MHz/Tesla}$.

The next stage is to transmit a radio frequency (RF) pulse into the patient, which will generate a much weaker magnetic field, denoted as B_1 . Since the RF pulse is an oscillating electromagnetic wave that is orthogonal to B_0 , the resulting field will also oscillate, rather like an AC voltage. Again, the charged spinning nuclei will precess about this field at a rate of $\omega_1 = \gamma B_1$. Since $B_1 \ll B_0$, then $\omega_1 \ll \omega_0$ also. The overall result is a spiral motion caused by the two precessions. This is called nutation.

If the B_1 field oscillates with a frequency ω_2 , and if $\omega_2 = \omega_0$ then resonance will occur. That is, energy will be added to the system. The result is that the protons will 'flip' and now precess about a different axis, Fig. 2.2, at the frequency ω_1 . The angle of the flip, θ , can be determined from the strength of the B_1 field and the duration, τ , that the field is applied for:




Figure 2.1: Precession of a vector in a static magnetic field.

$$\theta = \gamma B_1 \tau. \quad (2.2)$$

In order to gain information about the tissues in the patient, a series of measurements need to be taken. These are based on the relaxation times, that is the time taken for the spins to relax back to their lowest energy state when the B_1 field is removed. So, when the RF pulse is turned off, the charged nuclei will realign with the B_0 field and give up their excess energy, Fig. 2.2. The longitudinal relaxation time, i.e. the parameter that characterises the time for the spin to be realigned to the original B_0 field, is given by T_1 . That is, the M_z component of magnetisation slowly recovers along the z axis to M_0 , the original magnetisation, where

$$M_z(t) = M_0(1 - e^{-t/T_1}). \quad (2.3)$$

Similarly in the transverse direction, the relaxation time is given by T_2 , as the M_{xy} component of the magnetisation decreases rapidly,

$$M_{xy}(t) = M_0 e^{-t/T_2}. \quad (2.4)$$




Figure 2.2: Flipping and relaxation to equilibrium.

T_2 is independent of the external field. In reality however, we cannot measure the quantity T_2 as independent from the B_0 field, so instead we measure T_2^* , which is the transverse relaxation, also depending on B_0 . Whenever we refer to T_2 -weighted images throughout this thesis, the measurements will in fact, more strictly, be T_2^* . The rate of decay in the transverse direction is several times faster than the rate of recovery in the longitudinal direction as can be seen in Fig. 2.3.

This procedure needs to be repeated a number of times in order to actually produce an image, and the time between each RF signal is denoted by the repetition time, TR . The time between the RF pulse being switched off and the measurements being taken is known as the echo time, TE . Different tissue contrast properties can be obtained depending on the choice of TR and TE , and they are chosen differently depending on whether T_1 or T_2 is being measured, since T_2 relaxation is much shorter than T_1 . Usually for T_1 weighted imaging, a short TR is used, whereas for T_2 weighting, longer values of TE are selected to give maximum contrast, as seen in Fig. 2.4. A




Figure 2.3: Decay and Recovery of B_1 and B_0 fields giving T2 and T1 relaxation times.

proton density (PD) weighted image can also be acquired using a short TE, and a long TR.

2.3 Image Formation

In order to gain spatial information about the signals read out, magnetic field gradients are applied. There is usually one gradient in each of the x, y and z directions, denoted by G_x , G_y and G_z . A gradient is a magnetic field that changes from point to point, usually in a linear manner, for example in the z direction the magnetic field at z is then given by $B(z) = B_0 + zG_z$.

The first gradient that is used is a slice select gradient, G_z . This is a gradient in the external magnetic field. The bandwidth of the RF signal is then chosen such that resonance occurs in the slice of interest. As a slice is not infinitely thin, a bandwidth of frequencies of RF pulses is applied in order to cause resonance in all the frequencies within a thin slice, 3mm say. These frequencies are selected according to the Larmor equation described in Equation 2.1.

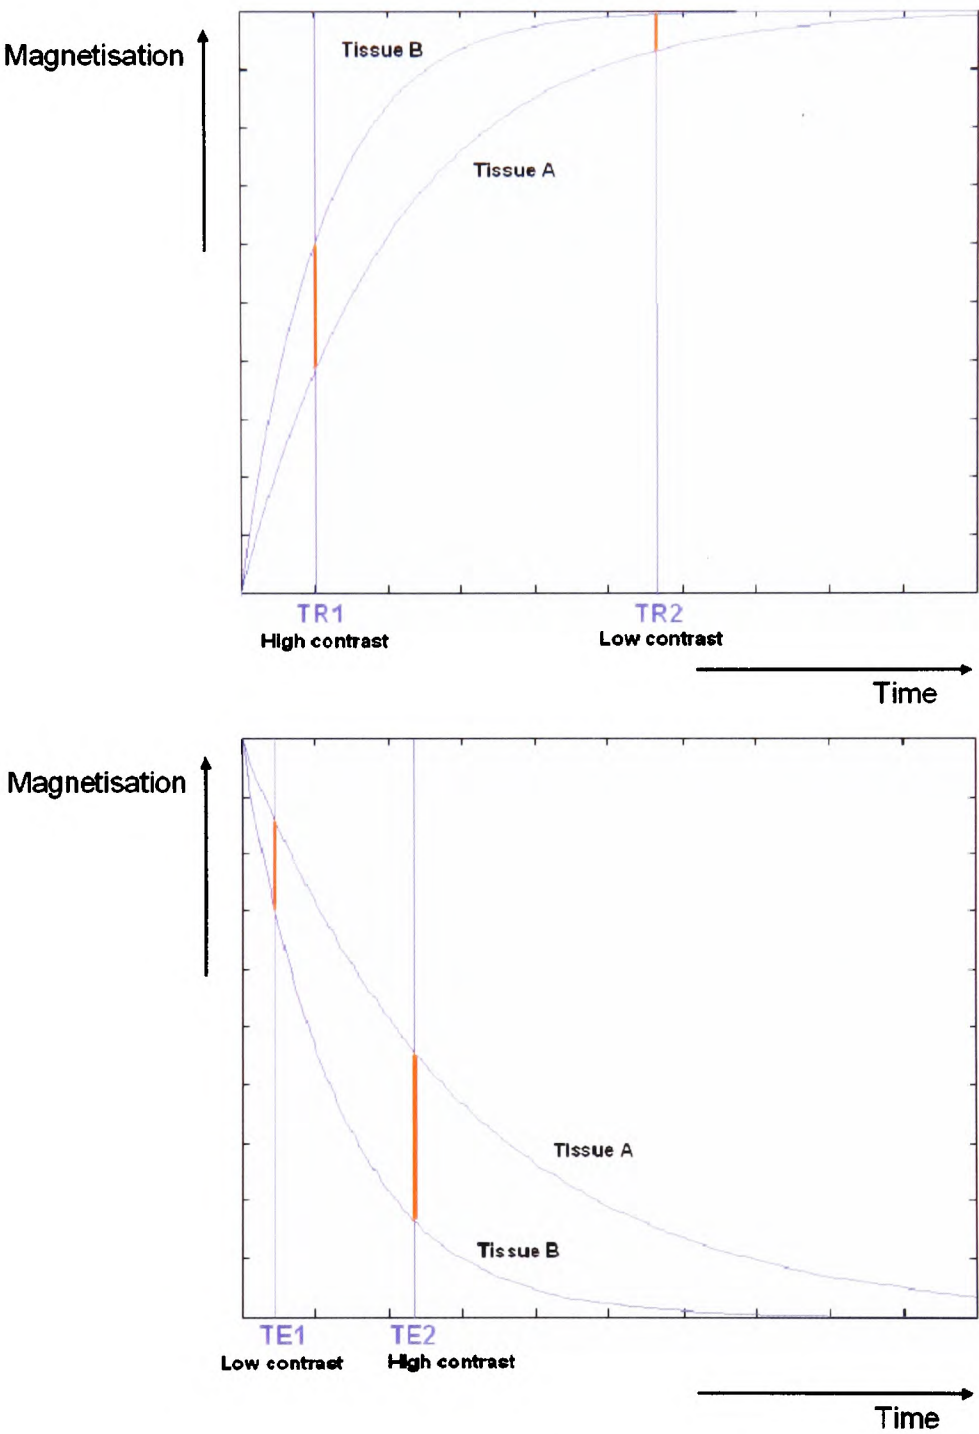


Figure 2.4: TR and TE times for T1 and T2 relaxations, indicating how a good contrast is achieved. The top graph shows how a short TR enhances the T1 contrast. The bottom graph shows that a longer TE can enhance the T2 contrast.




Figure 2.5: An example of a signal timing diagram for a spin echo sequence.

Having identified a single slice we then need to derive spatial information from within that slice. We first select a column within the slice using a frequency-encoding gradient, G_x . The G_z is turned on during the time when the RF pulse is applied and then turned off after. The G_x pulse is then turned on during the readout, i.e. after the echo time TE. The columns can then be identified by the frequencies that are read out.

Finally, to identify pixels on each column, a phase-encoding gradient is used, G_y . This is turned on straight after the RF pulse or just before the frequency-encoding gradient and adds a phase shift to the signal. Hence, the protons in each pixel can be identified by their distinct frequency and phase, which are unique and encode the x and y-coordinates for that pixel.

An example signal timing diagram is shown in Fig. 2.5 and shows a typical spin echo sequence. In this example a 90° gradient is used, followed by a refocussing 180° gradient to ensure maximum signal, and the relative timings of the three gradients can be seen in Fig. 2.5. More details of different sequences can be found in [44].

Each signal that is read out then fills a line in ‘k-space’, which is a digitised version




Figure 2.6: An example of k-space and its corresponding MR image.

of the dataspace. That is, G_z is kept constant for each slice, and G_y varies after each TR to fill the k-space line by line. The final image is then formed by taking the inverse fourier transform of the information in k-space. An example of k-space and corresponding image is shown in Fig. 2.6. Note that the final image only uses the intensity information. Phase information can also be useful but would need to be unwrapped [93].




Figure 2.7: Effect of field inhomogeneities on a colorectal MR image.

2.4 Image Artefacts

A number of artefacts can be present in MR images. In a few cases these can be useful for diagnostics, for example, chemical shift artefact, but in most cases they degrade the image quality and cause post-processing algorithms to fail. These artefacts are caused by a variety of different reasons, relating to the scanner, the patient, and the final image processing stages.

2.4.1 Gradient-Related Artefacts

Although the gradient fields are assumed to be linear, in many cases they often are not, and this results in large field inhomogeneity effects within the images, as can be seen in Fig. 2.7. This is a major problem in MR images for colorectal cancer. The intensity information is deficient and the details within the image cannot be seen. Consequently, it also causes problems when it comes to running any further image analysis algorithms on the images, and the algorithms will often fail. It is discussed in detail in Section 2.5 below.



Figure 2.8: Effect of aliasing on a colorectal MR image.

2.4.2 Patient-Related Artefacts

Patient-related artefacts are caused by voluntary or involuntary patient movements when they are in the scanner. Some of these are due to blood flow through vessels, which results in a pulsating motion. Other motion tends to be more random. Often the best solution to the latter problem is to tell the patient to remain still, or to sedate them. Compensation methods are also available to account for motion due to breathing.

2.4.3 Image Processing Related Artefacts

Aliasing is one of the most common image artefacts caused at the image processing stage. It is caused by undersampling during the read-out, hence it appears that the signal has a lower frequency than its real value. This in turn causes a wrap-around effect on the image, as seen in Fig. 2.8. The solution to this is to increase the sampling rate.

Chemical shift artefacts are caused by the fact that protons from different molecules precess at slightly different frequencies. This problem arises in the body as both fat




Figure 2.9: Effect of field chemical shift on a colorectal MR image.

and water are imaged. This causes an error in the frequency encoding direction and the result is a registration error between the two regions. This can be seen in Fig. 2.9. This effect can be reduced by using a ‘fat sat’, or fat saturation (suppression) imaging sequence, however it can be useful in identifying fat/water boundaries.

Gibbs ringing is an artefact that occurs at high contrast interfaces and causes alternating light and dark bands along the interface. It is caused by a limited number of samples or sampling time, which means a step-change in intensity cannot be found exactly. This is solved by increasing the sampling time and lowering the bandwidth, or by decreasing the pixel size.

Partial volume artefacts arise when an interface occurs within the voxel space. As such the resulting voxel intensity becomes a weighted average of the two tissues present within that voxel. This appears in the images as blurred edges, as can be seen in Fig. 2.10.




Figure 2.10: Effect of partial volume effect on a colorectal MR image.

2.4.4 RF Related Artefacts

A number of artefacts also arise due to the RF field. Crosstalk is caused by overlap in the Fourier transform between adjacent RF pulses, and appears on the images as interference from adjacent images. Zipper artefacts, often caused by imperfections in the RF pulse, appear as a line of bright and dark spots across the images, and also RF noise causes artefacts in the image that appear as static, and so can be misdiagnosed as pathology.

2.5 Bias Field Correction

There is increasing demand from clinicians for quantitative analysis of images in a wide range of medical applications. This is of particular importance when staging colorectal cancer because the size of the cancer needs to be known as well as the proximity to other organs. In order to analyse images quantitatively, a variety of segmentation and registration algorithms are often used. These algorithms require the images to be piecewise constant, that is, all areas of a particular tissue class have

the same intensity and texture pattern on the image [3]. Hence, in an ‘ideal’ image the intensity of each tissue class is constant, or perhaps has a Gaussian distribution with a small standard deviation. However, in MR images, a slowly varying field known as the bias field corrupts the image values.

The bias field is caused by imperfections in the RF transmitter/receiver (Tx/Rx) system. This, in turn, causes lowpass filtering of the k-space produced by the B_1 magnetic field, which can be represented as a convolution with some impulse response function $B(\bar{\omega})$. After inverse Fourier transformation, the result is a multiplicative bias field $b(x)$. Note that since $B(\bar{\omega})$ is real, $b(x)$ will be complex, so in practice the modulus of this signal is taken, $|b(x)|$, and so the noise is approximately Rayleigh distributed. That is, if the noise in the real and imaginary parts of $b(x)$ is Gaussian. The resulting effect is that the image is no longer piecewise constant and appears to have an undesirable ‘highlight’. Often, for qualitative analysis, this does not cause a problem (though it does adversely affect brightness and contrast depending on the magnitude of the effect), but quantitatively it produces large errors, e.g. in segmentation. Mathematically, the bias field can be described by:

$$s(x) = o(x) * b(x) + n(x), \quad (2.5)$$

where $s(x)$ is the signal that is reconstructed and is seen as the resultant image; $o(x)$ is the original or ‘ideal’ image that we are trying to extract; $b(x)$ is the bias field and $n(x)$ is noise. Initially, the noise is assumed to be negligible, compared to the effect of the bias field, and hence it is ignored in this part of the analysis. Taking logarithms, the equation becomes an addition rather than a multiplication and is therefore simpler to solve once the bias field has been estimated:

$$\log s(x) = \log o(x) + \log b(x). \quad (2.6)$$

Early techniques for estimating the bias field involved phantom imaging where a known object is scanned and the bias field estimated from the resultant image [85]. However, the bias field is also influenced by the interaction of the field inhomogeneities with the patient, and hence varies between patients. The best approaches therefore

calculate a new bias field for each individual dataset. The validation of whether a bias estimation and correction algorithm work can be both visual and quantitative. If successful the correction should be clearly visible by just looking at the image and seeing whether tissue areas are constant in intensity. The variation of intensity in the image can also be measured to give a quantitative comparison, using the standard deviations of those intensities across a particular tissue class. Two leading methods for estimating the bias field are described below.

2.5.1 Guillemaud and Brady method

Guillemaud and Brady [42] propose a statistical approach to estimate the bias field and provide a modification to the method previously developed by Wells *et al.* [90]. This method is used for bias field estimation and segmentation and is based on the Expectation Maximisation (EM) algorithm. This is an iterative method that interleaves the bias estimation based on current values, and their associated probable tissue classes, and is followed by an update of tissue class probabilities based on the current bias field estimate. A set of classes is chosen to represent the tissue types Γ_j based on the logarithm of their intensities Y_i . This is modelled by a Gaussian with mean μ_j and variance ψ_j .

$$p(Y_i|\Gamma_j) = G_{\psi_j}(Y_i - \mu_j), \quad (2.7)$$

where G_{ψ_j} is a scalar Gaussian distribution with variance ψ_j and Y is the logarithm of the intensity of the i^{th} voxel in the image volume. This leads to a likelihood equation for each voxel:

$$p(Y_i) = \sum_{\Gamma_j} p(Y_i|\Gamma_j)p(\Gamma_j). \quad (2.8)$$

Wells *et al.* [90] modified this approach in order to use the EM algorithm. After correction for the bias field β_i (and ignoring the noise),

$$p(Y_i|\Gamma_j, \beta_i) = G_{\psi_j}(Y_i - \mu_j - \beta_i). \quad (2.9)$$

This in turn gives the equation,

$$p(Y|\beta) = \prod_i [\sum_{\Gamma_i} p(Y_i|\Gamma_i)p(\Gamma_i)]. \quad (2.10)$$

The bias field is modelled by an n-dimensional zero-mean Gaussian prior probability density $p(\beta) = G_{\psi_i}(\beta)$. Bayes' rule is applied to give:

$$p(\beta|Y) = \frac{p(Y|\beta)p(\beta)}{p(Y)}, \quad (2.11)$$

to obtain the posterior probability of the bias field. Now let W_{ij} denote the probability that the bias-corrected voxel i belongs to the class j . Consider the mean residual $R = (R_1, R_2, \dots, R_n)^T$, where for voxel i ,

$$R_i = \sum_{\Gamma_i} W_{ij} \psi_j^{-1} (Y_i - \mu_j). \quad (2.12)$$

Denote the mean inverse covariance between voxels:

$$\psi_{ij}^{-1} = \begin{cases} \sum W_{ij} \psi_j^{-1} & \text{for } i = k \\ 0 & \text{otherwise} \end{cases} \quad (2.13)$$

Therefore,

$$W_{ij} = \frac{p(Y_i|\Gamma_j, \beta_i)p(\Gamma_j)}{\sum p(Y_i|\Gamma_j, \beta_i)p(\Gamma_j)}, \quad (2.14)$$

$$\beta = HR, \quad (2.15)$$

where, $H = [\psi^{-1} + \psi_{\beta}^{-1}]^{-1}$. Let F be a lowpass filter, then the bias field at the i^{th} voxel is updated by:

$$\beta_i = \frac{[FR]_i}{[F\psi^{-1}1]_i}, \quad (2.16)$$

with $1 = (1, 1, \dots, 1)^T$.

The EM algorithm is then applied. The expectation step calculates W_{ij} and the maximisation step then calculates β_i using the equations above. Once an estimate for the bias has been found, it is removed using the logarithms, as described in Eqn. 2.6.

It was found that Wells' method worked well for classes whose intensity distribution is Gaussian with low variance; but that it was unstable when variance is

large. Large variance is typical of Cerebrospinal Fluid (CSF) and pathological tissue (other). The modification that Guillemaud and Brady suggested was to introduce a new tissue class Γ_{other} with a non-Gaussian probability distribution. This class absorbs all the measurements that are outside the significance region of the other modelled categories. Now $p(Y_i|\beta_i) = \sum_{\Gamma_i(gaussian)} p(Y_i|\Gamma_j)p(\Gamma_j) + \lambda p(\Gamma_{other})$ where λ is a constant.

The approach used to get the bias estimate is very similar, but now

$$R_i = \sum_{\Gamma_i - \Gamma_{other}} W_{ij} \psi_j^{-1}(Y_i - \mu_j), \quad (2.17)$$

i.e. the bias is only estimated with respect to the Gaussian classes. The same EM algorithm is used. The paper shows results for simulated data, as well as brain scans and breast MRI. Good results are shown for brain MRI, but the bias effect is much less in that case because of optimised RF coils. The breast images show a 53% change in variance correction and this is a clearly visible effect in the images. The modification means that the algorithm is more robust and versatile than the original Wells *et al.* algorithm.

In order to remove noise from the images, (see equation 2.5), the method of Anisotropic Diffusion is used. This is described in detail by Perona and Malik [68] and then applied specifically to MRI data by Gerig *et al.* [40]. This method provides noise reduction and contrast enhancement, while preserving the edges. Guillemaud and Brady observed that it is equally effective to apply this filtering after the bias has been removed as it has negligible effect on the bias removal itself. It therefore proves to be effective as a post-processing filter for this particular method.

The limitations of the Guillemaud and Brady method are that the algorithm can converge to incorrect local maxima, and it is especially sensitive to the choice of classes. They propose a minimum entropy approach for choosing the class values using a histogram, but the algorithm is still sensitive to the starting values.

The major advantages of the method are its increased robustness, and requiring a less specific model. This is especially useful when looking at abnormalities in images.




Figure 2.11: Diagram of Legendre polynomials for two dimensions (0th to 4th degree).

2.5.2 Styner method

Styner *et al.* [80] propose a different approach to that of Guillemaud and Brady, which models the bias field implicitly as a low frequency artefact. They use a method called Parametric Bias Field Correction (PABIC) which models the bias field explicitly using Legendre polynomials (Kreyzig [53]). In two dimensions these are given by:

$$\hat{b}(\underline{x}, \underline{p}) = \sum_{i=0}^{m-1} p_i f_i(\underline{x}) = \sum_{i=0}^l \sum_{j=0}^{l-i} p_{ij} P_i(x) P_j(y). \quad (2.18)$$

$f_i(\underline{x})$ represents m smooth basis functions (Legendre polynomials) and these can be visualised in Fig. 2.11. $P_i()$ is a Legendre polynomial of degree i .

The bias estimate can be constrained by selecting the order of Legendre polynomials to use. This allows simple bias fields to be found quickly using a low order of polynomials while more complex fields can be found using a higher order of polynomials. This makes the method flexible. Prior knowledge about the nature of the field inhomogeneity can therefore be incorporated into the method. In Styner's method the original (or 'ideal') image $o(x)$ is defined by

$$o(x) = \mu_{k(x)}, \quad (2.19)$$

where $k(x)$ is the tissue class found at location x . There are a discrete, finite number of tissue classes found in the MR images of the colorectal region that can be identified. The PABIC algorithm finds the bias estimate directly from the image data. It does this by assuming that each pixel belongs to a certain image class k with an intensity μ_k . It then calculates an energy associated with each pixel according to whether or not it appears to fit into a class k . This energy function $e_0(\hat{d})$, where d is the difference between the signal data and the model, and this has a minimum at the class mean μ_k . For each class k a ‘valley’-function is defined:

$$e_0(\hat{d}) = \prod_k valley(\hat{d} - \mu_k). \quad (2.20)$$

The total energy for each pixel is defined as:

$$e_{tot} = \sum e_0(d_i) = \sum_i \prod_k valley(\hat{d} - \mu_k). \quad (2.21)$$

The total energy for the whole image is then:

$$e_{tot_img}(\mathbf{p}) = \sum_{x \in img} \prod_k \{valley(\tilde{s}(\mathbf{x}) - \tilde{b}(\mathbf{x}, \mathbf{p}) - \mu_k)\}. \quad (2.22)$$

An optimisation technique is then used to find the combination of Legendre polynomials that give the minimum energy. The strategy that is used is the (1+1) Evolution Strategy. It works by randomly selecting a new position in parameter space controlled by a probability function. The algorithm is described in more detail in Section 2.5.3.

The algorithm was first tested on simulated data and a mean squared distance measure (MSD) was used to validate the algorithm. All the corrections were successful. It was then applied to real two dimensional images of the brain with a moderate bias field and breast images with up to 40% bias field. The algorithm was seen to be successful and the corrected images were visually improved.

Although an effective method for removing bias, this method is limited by the computation time. The Evolution Strategy takes a long time to compute, and if a good first estimate of the parameters is not used, then the method can prove

unsuccessful. Also, the larger the order of Legendre polynomials used, the longer the computation time. This method is less sensitive to class parameters than that of Guillemaud and Brady, and a coarse estimate of classes still leads to successful estimation of the bias field. There are many other papers on the subject of bias removal, but most concentrate on brain images only [30], [76], or breast images only [41]. A further method was also proposed by Ahmed *et al.* [1]. The papers by Guillemaud and Brady, and Styner *et al.* have the advantage of using a general model that can potentially be applied to all MRI scans, containing different types of tissues and abnormalities.

2.5.3 Implementation

The method proposed by Styner *et al.* [80] was chosen to remove the bias field in the colorectal images. This is because it is less sensitive to the choice of a specific set of classes in the images. There is a diversity of anatomy visible in colorectal images, so it would be time consuming to estimate these classes individually. It will be seen later, Fig. 2.15, that the bias field is well modelled by 2nd order Legendre polynomials.

The first stage in implementing this method is to remove some of the intensity noise from the images such that we can ignore it before taking logarithms. This was done by anisotropically diffusing the images [68]. Noise removal made the bias estimation more accurate. Note that diffusing the image removes much of the texture. Such texture information is actually useful in determining different types of tissue and abnormalities. Hence once the bias field has been estimated using the diffused image, the bias field was then removed from the original (i.e. not diffused) image.

Anisotropic diffusion is a filtering process that improves the overall quality of the image [40], [68]. It has 3 three important properties:

1. It preserves the image boundaries and detailed structures.
2. It removes noise in homogeneous regions.
3. It enhances definition by sharpening any discontinuities.

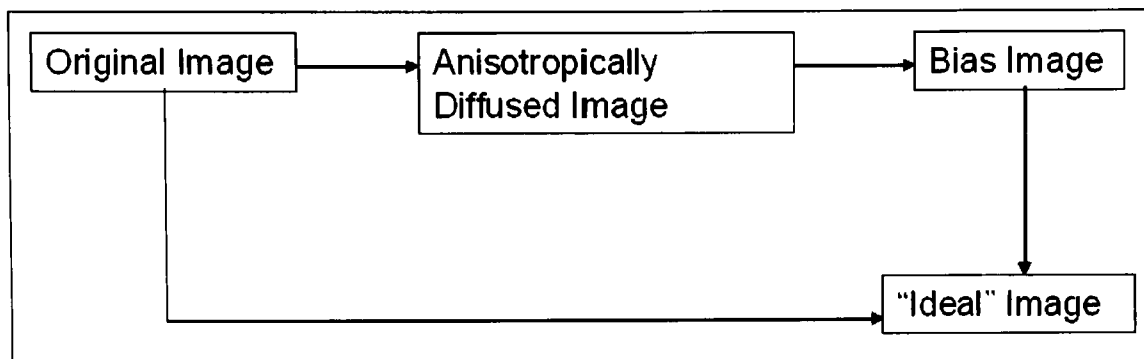


Figure 2.12: Bias Removal.

The latter is important because, in practice, the convergence of the bias field process is driven by the edges. The filtering method is modelled as a diffusive process,

$$\frac{\partial}{\partial t}(I(\bar{x}, t)) = \nabla(c(\bar{x}, t)\nabla I(\bar{x}, t)), \quad (2.23)$$

where $I(\bar{x}, t)$ is the image intensity. The diffusion strength is controlled by $c(\bar{x}, t)$, where \bar{x} are the spatial coordinates, such that $c(\bar{x}, t) = f(|\nabla I(\bar{x}, t)|)$. We used $c(\bar{x}, t) = -(\frac{|\nabla I(\bar{x}, t)|}{\kappa})^2$, where κ is a constant. We found the value of $\kappa = 20$ to work well for the images we tested. The ‘flow’ is defined as $\phi * (\nabla I) = c * \nabla I$. In order to filter a two dimensional image, a discrete implementation is used:

$$\begin{aligned} \frac{\partial}{\partial t}(I(\bar{x}, t)) &= \nabla^T[c(\bar{x}, t) * \nabla I(\bar{x}, t)] \\ &= \frac{\partial}{\partial x}[c(\bar{x}, t) * \frac{\partial}{\partial x}I(\bar{x}, t)] + \frac{\partial}{\partial y}[c(\bar{x}, t) * \frac{\partial}{\partial y}I(\bar{x}, t)] \\ &= \phi_{east} - \phi_{west} + \phi_{north} - \phi_{south}, \end{aligned} \quad (2.24)$$

where e.g. $\phi_{east} = \frac{1}{\Delta x^2}[c(x + \frac{\Delta x}{2}, y, t) * (I(x + \Delta x, y, t) - I(x, y, t))]$. The image is then updated:

$$\begin{aligned} I(t + \Delta t) &= I(t) + \Delta t * \frac{\partial}{\partial t}I(t) \\ &= I(t) + \Delta t * (\phi_{east} - \phi_{west} + \phi_{north} - \phi_{south}). \end{aligned} \quad (2.25)$$

The flow can be visualised in the Fig. 2.13




Figure 2.13: Flow of anisotropic diffusion.

Approximately 20 iterations of anisotropic diffusion, with $\kappa = 20$, were performed on the MR images. Once the images had been anisotropically diffused, the Styner algorithm could be applied. The (1+1) Evolution Strategy optimisation algorithm is given by:

$$\underline{\mathbf{r}}_t \sim N(\underline{\mathbf{0}}, I) \text{ computes a random vector } \underline{\mathbf{r}}_t \text{ at each iteration} \quad (2.26)$$

$$\underline{\mathbf{x}}_{t+1} = \underline{\mathbf{x}}_t + A_t \cdot \underline{\mathbf{r}}_t \quad (2.27)$$

$$A_{t+1} = \begin{cases} A_t(I + (c_g - 1) \frac{\underline{\mathbf{r}}_t \underline{\mathbf{r}}_t'}{\underline{\mathbf{r}}_t' \underline{\mathbf{r}}_t}) & \text{if } f(\underline{\mathbf{x}}_{t+1}) < f(\underline{\mathbf{x}}_{opt}) \\ A_t(I + (c_s - 1) \frac{\underline{\mathbf{r}}_t \underline{\mathbf{r}}_t'}{\underline{\mathbf{r}}_t' \underline{\mathbf{r}}_t}) & \text{otherwise} \end{cases} \quad (2.28)$$

$$\underline{\mathbf{x}}_{opt} = \begin{cases} \underline{\mathbf{x}}_{t+1} & \text{if } f(\underline{\mathbf{x}}_{t+1}) < f(\underline{\mathbf{x}}_{opt}) \\ \underline{\mathbf{x}}_{opt} & \text{otherwise.} \end{cases} \quad (2.29)$$

$\underline{\mathbf{x}}_{t=0}$ is used as an initial guess and $f()$ is an energy function. The values of c_g and c_s are weights, and $\underline{\mathbf{x}}_0$ and A_0 are given as input at the start of the algorithm. The values $c_g = 1.05$, $c_s = c_g^{-1/4}$ were used, as proposed by Styner *et al.* [80].

2.5.4 Results and Discussion

The effect of anisotropic diffusion can be seen by zooming in on a small section of the image as shown in Fig. 2.14. The texture is removed and the edges are much better




Figure 2.14: Before and after anisotropic diffusion.

defined.

When the optimisation algorithm was applied, it was found that the initial estimate of the bias field was important, a fact which was not discussed in Styner's paper [80]. Without a good first estimate, the algorithm became stuck in local minima. We found that a good first estimate was relatively straightforward to compute, however, by looking at the bias field in the image and predicting which part of the Legendre polynomial was having the greatest effect. The initialisation is therefore application specific. Being predominantly a function of the RF coil, the bias field in all the MR colorectal images tends to be in the same direction. The initialisation only needs to be estimated once for all the datasets, and could be done using a phantom, or numerically. We estimated the field numerically by taking the mean of typical fields. The ES algorithm is intrinsically slow, so a good initialisation helped considerably to speed it up. Our implementation of the algorithm limited the number of iterations, so that as the algorithm converged, it did not keep trying to find an improvement, when little or no further improvement is possible. This could also have been done using $|f(x_{t+n}) - f(x_t)| < threshold$. We also limited the order of the Legendre polynomials to two, so that the algorithm only needs to optimise six variables. This was found to be adequate to model the bias field and the visual results were much improved. Increasing the order increases the number of variables significantly, hence the computation time and numerical stability of the algorithm increase. Also, the more variables




Figure 2.15: Estimation of the Bias Field using the Styner method visualised in two dimensions and three dimensions.




Figure 2.16: Axial image before and after bias removal.

there are, the more chance that the algorithm becomes stuck in local minima. The estimated bias field for one of the colorectal MR images can be seen in Fig. 2.15 and it can be seen that the bias is therefore well modelled by the Legendre polynomials.

A numerical error was sometimes found to occur when the log of the bias function was calculated, due to proximity to zero. This is a problem that did not arise in Styner's work. We solved this by using a cut-off value for the $\log(bias)$. The bias field was then removed giving results typified by that in Fig. 2.16.

Table 2.1: Comparison of standard deviation/mean, averaged over all tissue classes.

Standard deviation/mean before bias removal	Standard deviation/mean after bias removal
0.2449	0.0977

The level of accuracy required is such that the image can be segmented. The accuracy was analysed by looking at the standard deviations of the image with respect to the mean values, averaged over all tissue classes. The results of this can be found in Table 2.1, and it can be seen that the (standard deviation/mean) is much lower for bias removed images.

The images used in diagnosing colorectal cancer usually take the form of a volume dataset. Styner had extended the method successfully to three dimensions, however this requires many more parameters, and much increased computation time. We noted, however, that the bias estimate is very similar on all the images running through the volume. Hence the bias field found for the first image can be used as the initial estimate for the second image and so on. That is, the two dimensional method is run on each slice consecutively. Alternatively, information from a phantom study could be used as a first volumetric estimate, but this would still take a long time to compute. Our implementation ensures that the computation time is kept to a minimum. It takes approximately 3 minutes running in Matlab, on an Intel XEON CPU 2GHz, 1.5GB RAM. Note that all algorithms throughout this thesis were run on this same machine. It was also noted that if the same bias field was used for all the images in the volume, the visual difference was negligible. This means that the bias can be calculated for the first image and then that bias removed on all the successive images. This is possible because the slices are thin (3mm) and do not vary a lot in content between slices. The results of a set of images after their bias field has been removed can be in Fig. 2.17.

Finally the algorithm that was implemented was tested briefly on breast images. It was seen to remove bias successfully on these types of images as well, and the images were visually improved, Fig. 2.18.




Figure 2.17: Bias removed images.




Figure 2.18: Left hand images are the original breast MR images, and right hand images are the corresponding bias corrected images.

2.5.5 Summary

This method works effectively and the quality of images is much improved without any overall filtering of the final data. This means that none of the useful information is lost in the bias removal process. The algorithm was seen to work on a variety of MR images proving that it is a flexible method. On careful inspection of the images it can be seen that there is still a small amount of bias visible in the images towards the edges. In order to solve this, higher order Legendre polynomials would be required, and the process would take much longer to run. However it will be demonstrated below that the bias removal process shown here is sufficient to enable segmentation

of the images, and hence increasing the order of the polynomials for a small effect is unnecessary. Further investigations will be carried out using a phantom to provide the first estimate for bias removal in these images. Clinically, the bias-removed images have better contrast and more consistent brightness, which aids decision making. From a computer vision point of view, this allows segmentation to be performed along with further post-processing methods.

Chapter 3

Anatomy Segmentation

Segmenting the key regions of the anatomy is an important first step in analysing the images. Extracting object boundaries in medical imaging is an important but challenging task because the signal to noise ratio is low. This boundary extraction enables the proximity of cancers to various anatomical structures to be found as well as enabling abnormal shapes and sizes of objects to be identified. Segmentation is the process of finding such object boundaries automatically. It is labour intensive to extract the boundaries manually, so there have been many attempts to come up with efficient and automatic methods. The method used is highly dependent on the nature of the object being segmented, and how much prior information is available.

For the application of staging and analysing colorectal cancers, the objects that ideally need to be segmented are the colorectum (where the primary cancer is usually situated), and the mesorectum (where any involved lymph nodes are most likely to be found). The mesorectal fascia is often taken as the circumferential resection margin, CRM, (the edge of the region that is resected in surgery), and the location of this edge gives a likely indication of the success of surgery. That is, if the tumour or any involved nodes are located near to this margin then the probability of recurrence is much higher. The MERCURY project [72] is a collaborative project that has been set up to investigate these margins more closely and compare the MR images with the resected specimens to see how well the CRMs can be estimated from the MRI.

MRI data for staging colorectal cancer uses a three dimensional dataset (i.e. a

volume); but it is usually viewed as individual two dimensional slices through that volume. However, it can be very useful to have a three dimensional representation of that data in order to see areas of particular interest (such as lymph nodes) in relation to the rest of the anatomy.

3.1 Active Contours

One object of particular interest is the colorectum, and this needs be segmented out of the images and viewed in three dimensions. There are many segmentation methods available in image processing, serving different purposes, and useful on different types of images. Active contours (or snakes) are used for this application since the boundary of the rectum is clear and well defined, as it is surrounded by the fatty mesorectal fascia. Since the colorectum is essentially a long tube, a tracking algorithm can be used to locate it through a series of images. Alternatively, methods such as level sets [65], [17], or graph cuts [12] could also be used.

3.1.1 Theory

The idea of deformable contour models, originally proposed by Kass, Witkin and Terzopoulos [48], overcomes many of the problems of manual slice editing such as noise, small gaps and other irregularities in the object boundaries. Snakes are used extensively in medical images and have been studied in detail by Cohen [22]. Cohen *et al.* improved upon the original snakes by introducing an ‘inflation’ or balloon force [23]. A snake is a dynamic contour that changes over time, such that it minimises the total energy of the contour. If $v(s, t)$ represents points on a dynamic curve at a discrete time t , then at time $t + 1$, $v(s, t + 1)$ is given by:

$$v(s, t + 1) = \arg \min_{v(s, t) \in V} (E(v(s, t)) + |\underline{b}(v(s, t))|), \quad (3.1)$$

where v is the set of all possible curve points $v(s, t)$, $E(v(s, t))$ is an energy function and $\underline{b}(v(s, t))$ is a balloon force. A dynamic curve is a two parameter function v :

$[0, 1] \times [0, \infty) \rightarrow \Omega$, where Ω is a two dimensional image domain. It parameterises image coordinates (x, y) by parameter s and time $t \in [0, \infty)$ such that,

$$v(s, t) = (x(s, t), y(s, t))^T. \quad (3.2)$$

The energy $E(v(s, t)) = E_{int}(v(s)) + E_{ext}(v(s))$ is defined by:

1. The internal energy $E_{int}(v(s)) = \alpha|v'(s)|^2 + \beta|v''(s)|^2$ where $\alpha, \beta \in \mathbb{R}^+$, $v'(s) = \frac{\partial v(s)}{\partial s}$ and $v''(s) = \frac{\partial^2 v(s)}{\partial s^2}$.
2. The external energy $E_{ext}(v(s)) = -\gamma|\nabla I(v(s))|^2$ where $\gamma \in \mathbb{R}^+$, $I(v(s))$ denotes the intensity at $v(s)$, and $|\cdot|$ denotes the Euclidean norm.

The balloon force $\underline{b}(v(s, t))$ is defined by:

$$\underline{b}(v(s, t)) = \delta(1 - \frac{|\nabla I(v(s))|^2}{(\max_s |\nabla I(v(s))|)^2})n(v(s)), \quad (3.3)$$

where $\delta \in \mathbb{R}$, $n(v(s))$ is the unit normal vector at $v(s)$. A further modification to the dynamic contours was proposed by Meegama and Rajapakse [63] whereby a NURBS (Non-Uniform Rational B-Spline) snake is used. This includes an extra weighting parameter to control the flexibility of the curve at each control point and is good for objects that have regions of high curvature.

The Meegama and Rajapakse approach [63] consists of a parametric curve that can be defined as a linear combination of control points, for example $v(s) = \sum_{i=0}^{n-1} f_i(s)p_i$ where p_i are the control points and $f_i(s)$ are a set of polynomial functions. $f_i(s)$ can be defined such that the points form a B-spline curve:

$$f_i(s) = N_{i,k}(s) \quad (3.4)$$

$$N_{i,0}(s) = \begin{cases} 1 & s \in [s_i, s_{i+1}) \\ 0 & \text{otherwise} \end{cases} \quad (3.5)$$

$$N_{i,k}(s) = \frac{(s - s_i)N_{i,k-1}(s)}{s_{i+k} - s_i} + \frac{(s_{i+k+1} - s)N_{i+1,k-1}(s)}{s_{i+k+1} - s_{i+1}} \quad (3.6)$$

$$(3.7)$$

Hence the curve is now:

$$v(s) = \sum_{i=0}^{n-1} N_{i,k}(s) p_i. \quad (3.8)$$

For a NURBS snake

$$v(s, t) = \frac{\sum_{i=0}^{n-1} w_i^t(t) N_{i,k}(s) p_i^t}{\sum_{i=0}^{n-1} w_i^t(t) N_{i,k}(s)}, \quad (3.9)$$

or, $v(s, t) = \sum_{i=0}^{n-1} R_i^t(s) p_i(t)$ where $R_i^t(s) = \frac{w_i^t(t) N_{i,k}(s)}{\sum_{j=0}^{n-1} w_j^t(t) N_{j,k}(s)}$, where w_i^t are the weights of each control point at time t . The internal and external energies and balloon forces are then given by:

$$E_{int}(v(s)) = \alpha |p^T R'(s)|^2 + \beta |p^T R''(s)|, \quad (3.10)$$

$$E_{ext}(v(s)) = -\gamma |\nabla I(p^T R(s))|^2, \quad (3.11)$$

$$b(v(s)) = \delta \left(1 - \frac{|\nabla I(p^T R(s))|^2}{(\max_s |\nabla I(p^T R(s))|)^2}\right) n(p^T R(s)). \quad (3.12)$$

A NURBS snake is a dynamic NURBS curve $v(s, t)$ whose energy is minimised with respect to points on the curve such that the curve points at time $t + 1$ are given by:

$$v(s, t + 1) = \arg \min_{v(s, t)} (E_{int}(v(s, t)) + E_{ext}(v(s, t))) + b(v(s, t)). \quad (3.13)$$

NURBS uses a weight-based control method. Let $v(s; w_i)$ be the set of curves obtained by changing the weight w_i at a control point p_i . Also, let $v(s; w_i = 0)$ be the set of curves where w_i is zero. Then:

$$v(s; w_i) = v(s; w_i = 0) + \tau(p_i - v(s; w_i = 0)), \quad (3.14)$$

where $\tau = \frac{N_{i,k}(s) w_i}{\sum_{j=0}^{n-1} N_{j,k}(s) w_j}$. New weights are then updated according to $w_i^{t+1} = w_i^t + \eta \frac{(\kappa^t)}{\|\kappa(s, t)\|}$, where $\eta \in \mathbb{R}$ controls the amount of attraction of the curve towards the control point, and $\|\kappa(s, t)\|$ is the maximum curvature of the closed NURBS curve.

The algorithm is implemented using dynamic programming to evolve the control points, and a diagram of this algorithm can be seen in Fig. 3.1.

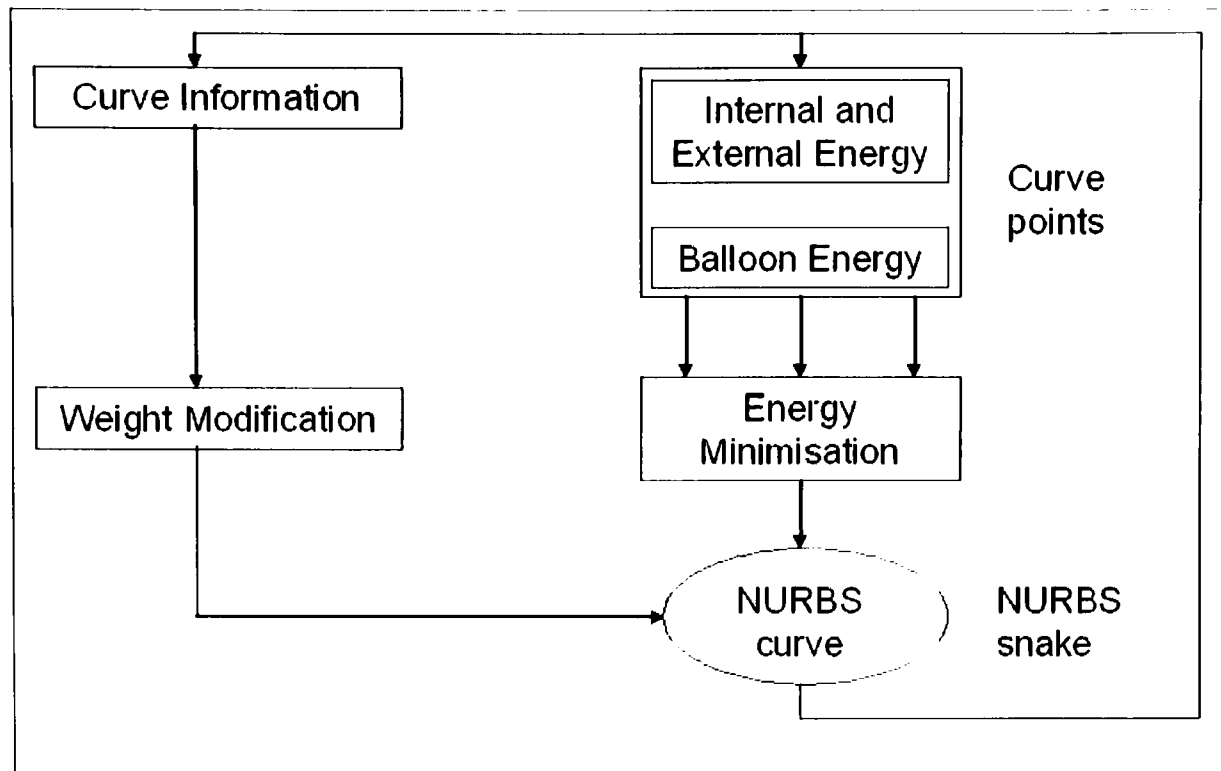


Figure 3.1: NURBS algorithm.

Experiments were performed by Meegama and Rajapakse using this algorithm on a range of different types of images: synthetic images, MR headscans, solar images, and aerial images. The NURBS approach was compared to other previously used snake methods (e.g. B-splines) and was found to give the best estimate. The snake accurately followed the boundaries even in regions of very high curvature. This is due to the extra weighting parameter that it includes, and there is therefore no need to duplicate or move relevant control points. The paper only demonstrates the method in two dimensions, but indicates that it can be easily extended to three dimensions. The major disadvantage is that it relies on clear and complete boundary edges, and hence is liable to fail where the edge is less obvious. It also doesn't include any prior information about the shape.

A review of the use of deformable models in medical image analysis was written by McInerney and Terzopoulos [62]. They indicated some of the main points to be considered when choosing a deformable model to use:

1. Automation vs. Control: There are some semiautomatic techniques but it is often useful to have a certain amount of user input.
2. Generality vs. Specificity: There are a number of ‘tailored’ models available for specific uses that are modified from the simpler general model.
3. Compactness vs. Geometric coverage vs. Topological flexibility: This indicates how much of the detailed ‘features’ are detected by the contour.
4. Curve vs. Surface vs. Solid models: This is very dependent on the shape and application of the object being segmented.
5. Accuracy and Quantitative power: The contour should be able to measure and control the parameter descriptiveness.
6. Robustness: The contour should be insensitive to the initial conditions and noisy data.

The method we develop here only requires user input to initialise the model. It is a general model that is relatively topologically inflexible, because the colorectum generally has a smooth edge. Since we begin by implementing the model in two dimensions, we use a curve. The distinctiveness of the edge means that the algorithm will be accurate and robust for this application.

3.1.2 Implementation

A snake algorithm, but with a few of the adornments suggested by numerous authors [63], has been used for segmenting the colorectum, since it is reasonably clear in the images (especially once the bias field has been removed). The purpose of the three dimensional representation of the colorectum at this stage is so that it can be visualised and relative positions can be observed. For our application, it is not necessary to be able to see every bump and crease for which the vastly superior spatial resolution of CT is necessary. Much more complex segmentations and reconstructions can be achieved on CT images, if required, [81], [82].

Since the object being represented is (locally) tube-like, it would be a waste of time to compute the snake on each slice. It is more economical, and numerically better conditioned, to track the object across slices using a Kalman filter. The overall method is given below:

1. A point was selected on the first image to indicate the approximate centre of the colon. This is currently done by hand and is not required to be accurate. The aim of this is to identify which object in the image is in fact the colorectum. This stage of the algorithm can be performed easily by a clinician or trained observer.
2. From the selected point, a snake was used to inflate to the outline of the rectum in the first image. The method is very precise in defining the edge of the object. This stage is performed on an anisotropically diffused, bias corrected image. It makes the edges even more defined as well as removing noise and field inhomogeneities.
3. Instead of using a snake on each consecutive image, the fact that the centre of the colorectum follows a three dimensional space curve was used. Snakes, though accurate, are relatively slow to calculate, so a Kalman filter was instead used to track the colorectum through the series of images. This provides sufficient accuracy to define the colorectum in three dimensional space and can be used on any dataset quickly and easily.

The Kalman Filter

The elastic model for shapes in motion that is most frequently used is the Kalman filter [89]. The idea of using Kalman filters in robotics to track an object through a series of images (either in time or space) has been used by Blake, Curwen and Zisserman [5], who track the bounding contours of objects. Tracking is an autonomous process and is a consequence of natural assumptions about geometry and uncertainty. The Kalman filter uses a predict/correct approach that can be seen in the Fig. 3.2.

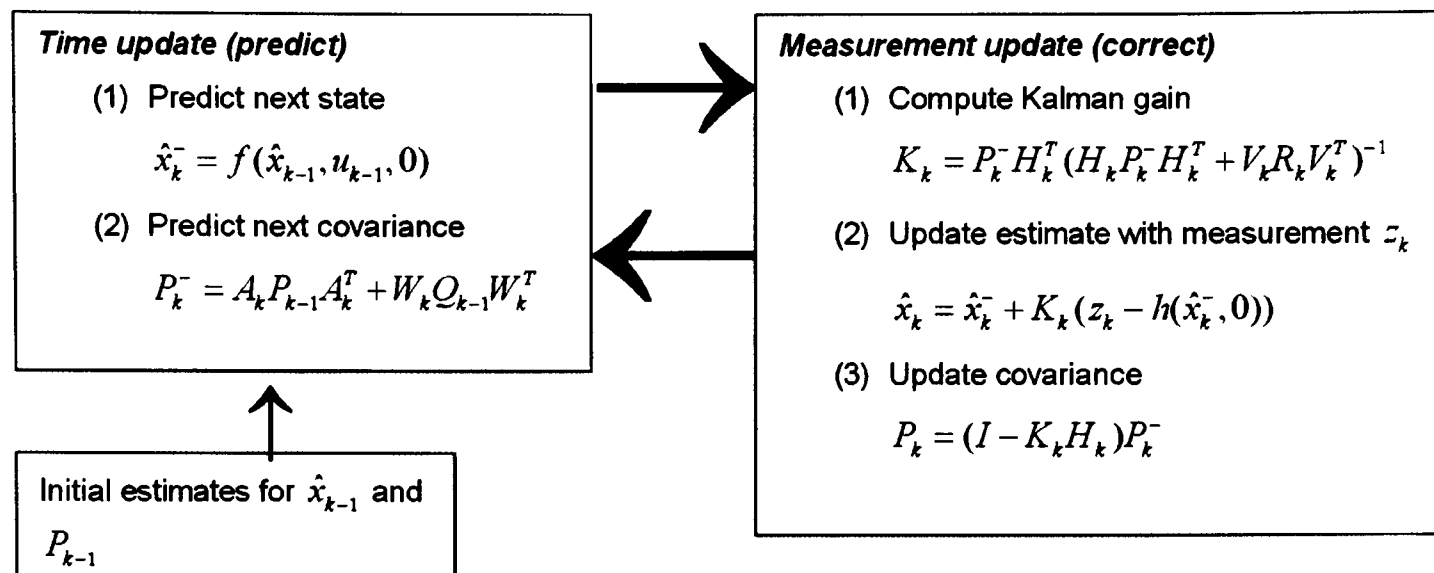


Figure 3.2: Kalman Filter Algorithm.

The moving feature is clearly the contour $v(s, t) = (X(s, t), Y(s, t))$, defined by the control points, $(X_n(t), Y_n(t))$, for $n = 1, \dots, M$, which we denote in terms of state vectors, $\mathbf{X} = (X_1, \dots, X_m)^T$ and $\mathbf{Y} = (Y_1, \dots, Y_m)^T$. The aim is to use the Kalman filter to develop an estimate for the contour state $(\hat{X}(s, t), \hat{Y}(s, t))$. We have assumed the measurement and estimation noise to be negligible (i.e. a perfect model) in the actual implementation.

The predict step is based on the assumption that a spline connects the corresponding control points in the consecutive images, and it is then updated using the measurement, which is by reference to a visual feature, $(X_f(s, t), Y_f(s, t))$. This feature is found by looking for the maximum gradient (i.e. edge) in the vicinity of $(\hat{X}(s, t), \hat{Y}(s, t))$, by searching along the line normal to the curve in that particular image.

The Kalman filter can then be described by:

$$\frac{d}{dt} \begin{pmatrix} \hat{\mathbf{X}} \\ \hat{\mathbf{X}} \end{pmatrix} = \begin{pmatrix} \dot{\hat{\mathbf{X}}} \\ 0 \end{pmatrix} + K(\mathbf{X}_f - \hat{\mathbf{X}}), \quad (3.15)$$

where $K(t)$ is the Kalman gain matrix, defined in terms of the measurement process $P(t) = E[(\mathbf{X} - \hat{\mathbf{X}})(\mathbf{X} - \hat{\mathbf{X}})^T]$ of the current estimate:




Figure 3.3: Active contour outlining the rectum on the first image in a series of axial images.

$$K = PH^T R^{-1}. \quad (3.16)$$

H is the measurement matrix $H = (\mathbf{I} \ 0)$, and R is the predicted covariance based on what has happened previously.

3.1.3 Results and Discussion

The initialisation of the snake algorithm is relatively unimportant, so long as it is well within the colorectum region. However, the snake converges on the boundary at a faster rate if the starting point is towards the centre of the region. Nevertheless, if the start point is placed towards the edge, the algorithm does still converge, but more slowly. Naturally, when observing the image and choosing a start point, most people would put the start position near the centre of the region, so this doesn't cause a problem. The contour that is found by the snake can be seen in Fig. 3.3.

Having converged on the first image, the Kalman snake filter was then applied to the remaining images in turn, and resulting contours can be seen in Fig. 3.4,




Figure 3.4: Contours propagated through the set of axial images using a Kalman filter.

which shows a typical result. This takes approximately 4 minutes to run in Matlab, depending on the number of slices and resolution of the images.

It can be seen that this method gives a good estimate of the position of the colorectum. However, there are a couple of potential sources of error that can occur. The first is due to the partial volume effect, which arises when a voxel is at a boundary and comprises two substantive tissue classes. If there are two different types of tissue in a voxel, then the intensity of that pixel is the volume weighted sum of the two tissue type intensities. The net effect is to give a blurred edge. This is reduced to some extent by the anisotropic diffusion, but not entirely. Hence some of the images may not have a perfect contour as shown in Fig. 3.5. However, for this application, the contour is usually corrected in the next image. If a more exact representation of




Figure 3.5: A slice from the data-set shown in Fig. 3.4 showing the partial volume effect.

the rectum were required, then this effect would need to be accounted for in a more rigorous way.

The second source of error in the tracking arises when there are other parts of the anatomy with well defined edges close to the region being tracked. This occurs just below the mesorectal fascia, and the rectum is no longer surrounded by a layer of fat, as shown in Fig. 3.6. Fortunately, this region is not of particular interest for the staging as it does not contain lymph nodes. The error could be eliminated by restricting or predicting the motion of the snake as it is tracked, or by using a three dimensional snake model.

The tracked images that form the colorectum can be assembled into a three dimensional representation, shown in Fig. 3.7. The colorectum could also have been segmented using a three dimensional method, such as a three dimensional snake or an active shape model. This would perhaps have resulted in a smoother shape, but the result for each slice would have been less accurate, as the slices are 3mm apart, compared to the pixel size of approximately 0.4mm. Hence, smoothing across slices could introduce additional errors. Also, three dimensional methods are often slow due to an increased numbers of parameters that need to be optimised.




Figure 3.6: A slice from the data-set shown in Fig. 3.4 showing the tracking error.




Figure 3.7: Three dimensional visualisation of the lower colon/rectum using the contours found automatically.

Motion effects do not appear to be problematic in the datasets provided to us by the Churchill Hospital, and the colorectum stays fixed with little movement for the duration of the scan. There are no apparent artefacts due to respiratory motion. If motion were to become a problem, then registration algorithms would need to be applied as well.

One advantage of the snakes method is that it tracks the outside of the tumour assuming this is part of the colorectum. This is due to the relatively similar intensities of the tumour and colorectum. Although this is not entirely accurate, anatomically speaking, it is useful when it come to segmenting the mesorectum and searching for lymph nodes, as it will exclude the tumour region from the search.

If we were required to segment the tumour as a separate section of the image then a method that incorporated texture information would be required. The texture of a region gives a lot of information about what type of tissue it is, so incorporating this into a tracking algorithm would be helpful. Cootes and Taylor use this idea to incorporate texture information into their active appearance models [24]. This would be particularly helpful when looking at the primary cancer, as, although its boundary is not always clear, it does have a different texture to both the colorectum and the mesorectal fascia. Fig. 3.8 shows a tumour in the colorectum, which is hand segmented in the second image. It can be seen how indistinct the primary tumour is. It can sometimes require some previous knowledge from the endoscopy or barium enema in order for the clinicians to know where to look for the tumour.

3.1.4 Summary

The three dimensional representation, if only as a visualisation technique, was deemed to be very useful by both by the radiologists and the surgeon who are involved in the project. The accuracy of this representation is sufficient to enable them to visualise the colorectum, tumour and nodes in three dimensional space. It is particularly useful to be able to see the nodes in relation to the rest of the anatomy and these can be placed on the three dimensional image as shown in Fig. 3.9. The method used is fast,




Figure 3.8: A typical axial slice showing the tumour outlined on the right hand image.




Figure 3.9: Three dimensional model showing a lymph node.

efficient and effective, suiting the requirements of the application.

3.2 Shape Models

The mesorectum is the layer of fat surrounding the rectum, and the mesorectal fascia is the edge of this band of fat. Lymph nodes are usually found within this region

and these are prime candidates to be cancerous if they are proximal to the primary tumour. Shape models can be used to find the shape of the mesorectal fascia in the colorectal images. The mesorectal fascia has a distinctive shape but its boundary is not always clear.

There is a lot of medical literature on the subject. Martling *et al.* [60] for example, highlight the issue of using MRI to look at the mesorectum, which investigates the prognostic value of pre-operative MRI of the pelvis. They state that if the circumferential resection margin (mesorectal fascia) is involved (i.e. cancer has spread to this margin either by extension of the main tumour, or involved cancerous lymph nodes) then the patient is at a higher risk of recurrence. Whether the cancer has spread this far is usually indicated by the pathological findings after a resection, but Martling states that pre-operative radiotherapy is in fact more effective than post-operative treatment, hence there is a need to know this information before surgery. Finally, Martling concludes that current pre-operative MRI techniques can indicate the extent of tumour growth with a reasonable accuracy (75% were correct compared with histology). This indicates how useful it would be to be able to segment the mesorectum and accurately locate the lymph nodes within it.

3.2.1 Theory

Statistical shape models are a different method for aiding the process of segmenting objects of interest from medical images. The objects of interest usually have a characteristic shape and location relative to other objects in the vicinity. Although the same clinical structures in different patients do vary they still have the same key features and location and hence can be identified manually. This process can be automated using active shape models. The ‘typical’ shape variation is learnt and then fit to new variations of the same object. There are a number of different methods of finding shape:

1. Point distribution, active shape, active appearance models (Cootes and Taylor [24]).

2. Fourier snakes (Szekely [83]).
3. Active contours (described above).
4. Parametrically deformable models (Staib and Duncan [78],[79]).

The Cootes and Taylor method involves fitting a statistical shape model to the training data. The training data is aligned using the Procrustes algorithm, which is explained in depth in [24]. Principal component analysis (PCA) is also used to investigate the shape variation across the training set. The shape model can then be applied to unseen data, by initially aligning the model with the image (position, orientation and scale). A search algorithm is used to find the best-fit points in the region of each point. This is done by looking at the local structure of the image, and then finding a best fit of the model. One disadvantage of this method is that it is point based. It looks at the local structure, and may pick out a local jagged edge, when in actual fact the overall smooth edge is desired (e.g. in the case of segmenting the mesorectal fascia described here).

Fourier surface models, introduced by Szekely *et al.* [83], are described by spherical harmonics. These surface models are elastically deformable and their flexibility is described by a set of parameters that can be varied to find the best fit on new image data. The eigenmodes of the parameters are used to define the shape variation across a number of shapes so that their flexibility can be modelled.

Staib and Duncan [79] use a torus model to define their shape model in three dimensions by decomposing the surface into sinusoidal basis functions. The surface finding is then performed as an optimisation using a gradient descent method. The object being modelled is assumed to have a smooth boundary that is continuously deformable, and hence a uniform representation of the entire shape is needed. Prior information is needed to form a flexible bias towards more likely shapes. Staib and Duncan choose to use a Fourier representation to express their surface that they are finding. This expresses any object as a weighted sum of known functions, as described for the active contours above. To express a function $X(t)$, in the interval (a, b) , in

terms of the basis $\phi_k(t)$, $X(t) = \sum_{k=1}^{\infty} p_k \phi_k$ where $p_k = \int_a^b X(t) \phi_k(t) dt$. Fourier representations are the most suitable because they are concise and not limited to a particular class of objects. A surface in three dimensions can be described by two surface parameters (u, v) :

$$\underline{x}(u, v) = (x(u, v), y(u, v), z(u, v)). \quad (3.17)$$

Therefore, a basis function of two variables is needed:

$$\begin{aligned} \phi = & \{1, \cos(m.u), \sin(m.u), \cos(l.v), \sin(l.v), \cos(m.u) \cos(l.v), \\ & \sin(m.u) \cos(l.v), \cos(m.u) \sin(l.v), \sin(m.u) \sin(l.v), \dots \\ & (m, l = 1, 2, \dots)\}. \end{aligned} \quad (3.18)$$

Each function is then represented by:

$$\begin{aligned} X(u, v) = \sum_{m=0}^{K_1-1} \sum_{l=0}^{K_2-1} \lambda_{m,l} [a_{m,l} \cos mu \cos lv + b_{m,l} \sin mu \cos lv + \\ + c_{m,l} \cos mu \sin lv + d_{m,l} \sin mu \sin lv], \end{aligned} \quad (3.19)$$

where

$$\lambda_{m,l} = \begin{cases} 1 & \text{for } m = 0, l = 0 \\ 2 & \text{for } m > 0, l = 0 \text{ or } m = 0, l > 0 \\ 4 & \text{for } m > 0, l > 0 \end{cases} \quad (3.20)$$

There are then three sets of parameters corresponding to the three coordinate functions

$$p_x = \{a_x, b_x, c_x, d_x\}, \quad (3.21)$$

$$p_y = \{a_y, b_y, c_y, d_y\}, \quad (3.22)$$

and

$$p_z = \{a_z, b_z, c_z, d_z\}. \quad (3.23)$$

These parameters can be used to describe open surfaces, tubes, and closed surfaces from the torus representation. An objective function is then used to find the boundary of a new occurrence of the surface. This is done using a boundary measure, b , such as

the gradient function of the three dimensional image. An energy function is defined based on the previous data (training sets) and the boundary measure. It is given by:

$$M(b, \mathbf{p}) = \ln Pr(\mathbf{p}) + \frac{k}{2\sigma_n^2} \int \int_A [b(x(\mathbf{p}, u, v), y(\mathbf{p}, u, v), z(\mathbf{p}, u, v))] dA, \quad (3.24)$$

where $dA = \left| \frac{\partial \mathbf{x}}{\partial \theta} \times \frac{\partial \mathbf{x}}{\partial \phi} \right| d\theta d\phi$. The method of conjugate gradient descent is then used to optimise this energy function [79].

This method uses the idea of incorporating the current knowledge of the dataset, from the data being segmented, along with prior knowledge of similar datasets of similar anatomy. The method is also less sensitive to broken boundaries, compared with using snakes. The paper displays good results for segmenting medical image data, specifically in the case of cardiac MRI. The main disadvantage however is that to make the method as effective as possible, a large training set is required.

3.2.2 Implementation

The mesorectal fascia can be seen in Fig. 3.10, segmented manually using the commercial software RevealMVSTM provided to the project by Siemens Molecular Imaging (formerly Mirada Solutions Ltd.).

The size and shape of the mesorectal fascia varies widely between patients, hence is difficult to define. Shape models use datasets in order to learn the probable shape of the image. They then provide a first estimate of the shape and use this along with new image data to find an accurate estimate. Shape models therefore need a training set of one or more sets of images and then unseen data on which to try the algorithm.

The method for finding the mesorectal fascia can be broken down into three steps. These are: finding a coordinate frame of reference for the dataset in which to place the shape model; defining the model itself; and finally optimizing the parameters.

Coordinate frame of reference

Since the colorectum can be modelled in a straightforward way by extracting information directly from the image, so too can other distinct features in the image. The




Figure 3.10: Outlined mesorectal fascia.

hips are particularly easy to identify and they can also be tracked and modelled. The colorectum together, with the hips and coccyx, enable a three dimensional coordinate frame to be constructed for the individual patient, and then points of interest that are not easy to track using standard techniques can be estimated. The edge of the mesorectal fascia is important to both radiologists and surgeons during diagnosis, but is not always clearly defined on the MR images.

In order to define a coordinate system, the following three items need to be determined:

1. An origin.
2. Coordinates - two or more directions.
3. A scale.

These can be found by knowing the locations of key points on the images. The key points that are searched for in these images are the hips and the coccyx. These points are then connected with a triangle. It was observed that the mesorectal fascia could be approximated by the inscribed circle contained within the triangle, Fig. 3.11.




Figure 3.11: Frame of reference of MR image.

This circle is then used as a first estimate of the mesorectal fascia in adjacent images, by finding its centre and radius. Two dimensional methods such as snakes that could be used to refine the estimate do not take into account the predicted shape of the mesorectal fascia and hence end up being jagged and inaccurate, Fig. 3.12. A method that takes into account both the predicted position and the shape is needed, hence the use of shape models.

Finding the model parameters

Since the mesorectal fascia is a relatively simple tubular shape, and fits in between the hip joints (which are clearly visible on the MR images), a simple model could be tested initially. From the circular first estimate, an elliptical model could be formed using two parameters $p = [a, b]$ in the x and y directions, which could be optimised to find the best fit.

However, a more complex model would better describe the shape of the mesorectum, and so the method using Fourier surfaces was also used to model the data. This Fourier surface could then be optimised to fit the data.




Figure 3.12: Tracking mesorectal fascia using a snake.

Fourier surfaces can be described in parametric form by the variables θ and ϕ .

$$r(\theta, \phi) = \begin{pmatrix} x(\theta, \phi) \\ y(\theta, \phi) \\ z(\theta, \phi) \end{pmatrix}, 0 \leq \theta \leq \pi, 0 \leq \phi \leq 2\pi. \quad (3.25)$$

Spherical harmonics [13] are used to describe the surface, along with the coefficients:

$$\underline{c}_l^m = \begin{pmatrix} c_{xl}^m \\ c_{yl}^m \\ c_{zl}^m \end{pmatrix}. \quad (3.26)$$

The surface then takes the following form:

$$r(\theta, \phi) = \sum_{l=0}^{\infty} \sum_{m=-l}^l c_l^m Y_l^m(\theta, \phi), \quad (3.27)$$

where $Y_l^m(\theta, \phi)$ are the spherical harmonics given by

$$Y_l^m(\theta, \phi) = \sqrt{\frac{2l+1}{4} \frac{(l-m)!}{(l+m)!}} P_l^m(\cos \theta) e^{jm\phi}, \quad (3.28)$$

$$Y_l^{-m}(\theta, \phi) = (-1)^m Y_l^{m*}(\theta, \phi), \quad (3.29)$$

and $P_l^m(x)$ are the associated Legendre polynomials:

$$P_l^m(x) = (-1)^m (1 - x^2)^{m/2} \frac{d^m}{dx^m} P_l(x) = \frac{(-1)^m}{2^l l!} (1 - x^2)^{m/2} \frac{d^{m+l}}{dx^{m+l}} (x^2 - 1)^l, \quad (3.30)$$

since $P_l(x) = \frac{1}{2^l l!} \frac{d^l}{dx^l} (x^2 - 1)^l$.

The coefficients are calculated using a least squares method:

$$C = (B^T B)^{-1} B^T X, \quad (3.31)$$

where $C = (c_0^0, c_1^{-1}, c_1^0, \dots)^T$, $X = (\mathbf{r}_0, \mathbf{r}_1, \dots, \mathbf{r}_{n_{vert}-1})^T$ and $b_{i,j(l,m)} = Y_l^m(\theta_i, \phi_i)$.

These equations can be used to find the coefficients of a Fourier surface when a number of points on the surface are already known. To find a model based on the colorectal cancer data a ‘leave-one-out’ strategy was used in order to train the model. That is, a model was derived from all but one of the datasets, and then tested on the dataset that was left out. The algorithm was run ten times. Each time it was trained on 9 datasets, and then tested on the tenth.

The model was produced by fitting a Fourier surface to the points on each training set, having normalised them based on position (from the first estimate calculated using the hip and coccyx positions), and scale (calculated from the radius of the inscribed circle of this first estimate). The mean of the parameters is then found in order to describe the model shape based on all the training sets.

This model can be translated and scaled onto the dataset being tested to give a ‘first guess’ of the surface. This is based on the proposal by Szekely *et al.* [83], that uses prior shape information from new image using geometry of the anatomy. An optimisation method can then be used to improve the surface description by altering the coefficient values.

Optimisation: Conjugate Gradient Method

A gradient descent optimisation method was implemented as proposed by Staib and Duncan. This method minimises the objective function from Eqn. 3.24, such that

$$\frac{\partial M}{\partial \mathbf{p}_x} = \frac{\partial \ln Pr(\mathbf{p})}{\partial \mathbf{p}_x} + \frac{k}{2\sigma_n^2} \int \int_A \left[|b(x, y, z)| \frac{\partial}{\partial \mathbf{p}_x} \left| \frac{\partial \mathbf{x}}{\partial \theta} \times \frac{\partial \mathbf{x}}{\partial \phi} \right| + \right. \\ \left. + \frac{\partial |b(x, y, z)|}{\partial x} \frac{\partial x(\mathbf{p}, \theta, \phi)}{\partial \mathbf{p}_x} \left| \frac{\partial \mathbf{x}}{\partial \theta} \times \frac{\partial \mathbf{x}}{\partial \phi} \right| \right] d\theta d\phi, \quad (3.32)$$

and similarly for y and z . The parameters \mathbf{p} in this case are the coefficients of the spherical harmonics.

The gradient descent method works using the following steps where $f(x)$ is the objective function M . Consecutive steps following the conjugate directions are found i.e. $\underline{\mathbf{d}}_i^T \underline{\mathbf{H}} \underline{\mathbf{d}}_j = 0, \forall i \neq j$, where $\underline{\mathbf{H}}$ is the Hessian.

1. Choose $\underline{\mathbf{x}}_0$, which is the first guess of the coefficients found using the patient-centric coordinate system. Compute:

$$\underline{\mathbf{d}}_0 = -\nabla f(\underline{\mathbf{x}}_0),$$

$\underline{\mathbf{H}}(\underline{\mathbf{x}}_0)$ (Hessian), and

$$\alpha_0 = \frac{\underline{\mathbf{d}}_0^T \underline{\mathbf{d}}_0}{\underline{\mathbf{d}}_0^T \underline{\mathbf{H}}(\underline{\mathbf{x}}_0) \underline{\mathbf{d}}_0}.$$

2. Determine $\underline{\mathbf{x}}_{k+1} = \underline{\mathbf{x}}_k + \alpha_k \underline{\mathbf{d}}_k$, $\nabla f(\underline{\mathbf{x}}_{k+1})$ and $\underline{\mathbf{H}}(\underline{\mathbf{x}}_{k+1})$.

3. Evaluate:

$$\beta_k = \left[\frac{|\nabla f(\underline{\mathbf{x}}_{k+1})|}{|\nabla f(\underline{\mathbf{x}}_k)|} \right]^2,$$

$$\underline{\mathbf{d}}_{k+1} = -\nabla f(\underline{\mathbf{x}}_{k+1}) + \beta_k \underline{\mathbf{d}}_k,$$

$$\alpha_{k+1} = \frac{\underline{\mathbf{d}}_{k+1}^T \nabla f(\underline{\mathbf{x}}_{k+1})}{\underline{\mathbf{d}}_{k+1}^T \underline{\mathbf{H}}(\underline{\mathbf{x}}_{k+1}) \underline{\mathbf{d}}_{k+1}}.$$

4. Repeat from step 2 until convergence.

This method converges linearly: $\|\underline{\mathbf{x}}_{k+1} - \underline{\mathbf{x}}^*\| = \beta \|\underline{\mathbf{x}}_k - \underline{\mathbf{x}}^*\|$.

Final Algorithm

The implementation of both the simple and complex methods took the following form: Points on the mesorectal fascia were selected and these made the basis of the training shape. A completely new, unseen data set was then used, and the mesorectal fascia position was predicted in the following way:

1. The geometric centre of new images was found using the hips and coccyx in the same way as described above. The radius of the inscribed circle was also found.
2. The shape model was calculated using the normalised training data.
3. The shape model was then fit to the new dataset, using the new radius (scale factor) and centre point, and also by varying the parameters (as described previously) to find the optimum.

3.2.3 Results and Discussion

The result from the initial simple method using just two parameters was almost identical to the first estimate, indicating that the first circular estimate is in fact very accurate. However, using just two parameters does not take into account the whole three dimensional shape of the mesorectum, and is therefore only useful as a rough estimate. This estimate however is quite robust.

In order to more accurately segment the outline of the mesorectum, the complex shape model was required, with more varying parameters, using Fourier surfaces. A Fourier surface fitted to one of the sets of data can be seen in Fig. 3.13. It can be seen that the Fourier surface (2nd order) gives a smooth and accurate estimate of the shape in 3 dimensions.

Ten datasets were tested using the 'leave-one-out' strategy. The pre-treatment scans were used. A Fourier surface was fitted to each dataset, and the mean of the parameters was calculated to give the shape model. This shape model can be seen in Fig. 3.14. This model can then be seen having been scaled and translated onto the test dataset in Fig. 3.15 to give the first estimate of the mesorectal fascia.

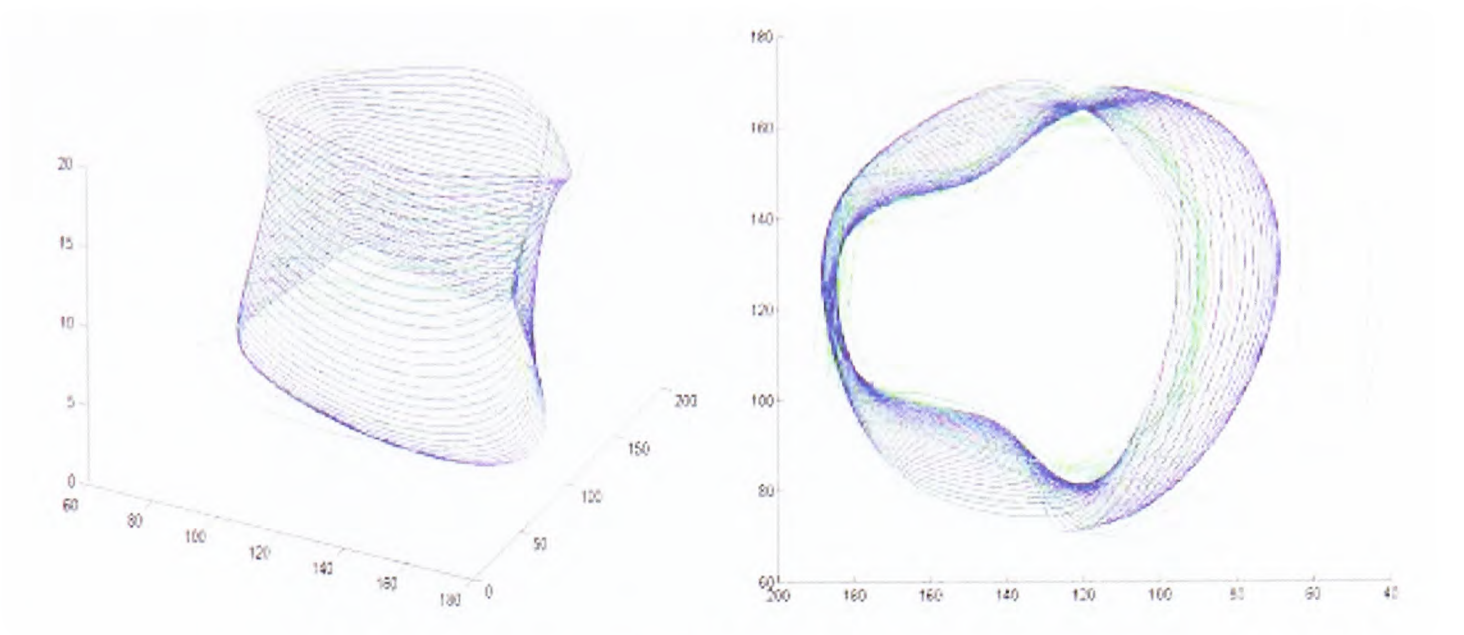


Figure 3.13: Fourier surface (blue) fit to training set (green).




Figure 3.14: Mean shape of the mesorectum calculated from 10 images, showing the Fourier surface in blue, and the mean of the training sets in green.

The red line shows the points on the actual mesorectum, and the yellow line shows the first estimate of the predicted shape. The two lines are very close with a mean maximum error of 5mm across the ten datasets tested. This is of the same order as the differences between the manual outlines drawn by two experienced clinicians, which varied between 0.5mm and 7mm.




Figure 3.15: First estimate of the mesorectum (mean shape).

Thirty iterations of the gradient descent algorithm were performed on each dataset in order to refine this estimate and give a more accurate prediction of the position of the mesorectal fascia. The time taken was approximately 5 minutes using Matlab. This was successful for six of the datasets, and some of these successful results can be seen in Figs. 3.16, 3.17, 3.18 and 3.19. The yellow line indicates the estimated shape, and the red line is the shape plotted by hand.

The volumes of the mesorectal fascia estimated for datasets ten and eight can be seen in the Fig. 3.20 and Fig. 3.21. These volumes look different due to the large variation in shape of the mesorectum. A set of consecutive, two dimensional images showing the mesorectal fascia plotted onto the MR images for dataset eight (Fig. 3.22). The maximum error where the algorithm runs correctly is 3mm.

There were four out of ten datasets where the active shape algorithm failed, and the errors can be seen in Table 3.1. Failures occurred due to two reasons:

1. False minima can occur when there are other edges that the shape model attaches to in the optimisation stage of the algorithm. They occur when the edges of the mesorectum are particularly indistinct in comparison to the surround-




Figure 3.16: Estimated mesorectum (yellow), and the manual segmentation (red) for dataset 1.




Figure 3.17: Estimated mesorectum (yellow), and the manual segmentation (red) for dataset 5.




Figure 3.18: Estimated mesorectum (yellow), and the manual segmentation (red) for dataset 8.




Figure 3.19: Estimated mesorectum (yellow), and the manual segmentation (red) for dataset 10.




Figure 3.20: Volume plot of the estimated mesorectum (yellow) for dataset 10.




Figure 3.21: Volume plot of the estimated mesorectum (yellow) for dataset 8.




Figure 3.22: Dataset 8 showing the raw data (top row), and the estimated mesorectal fascia (bottom row).

Table 3.1: Results of finding the mesorectum with the shape model for each dataset.

Patient	Pass/Fail	Maximum error (mm)
1	PASS	2.0
2	INSTABILITY FAILURE	
3	FALSE MINIMA FAILURE	
4	INSTABILITY FAILURE	
5	PASS	2.8
6	INSTABILITY FAILURE	
7	PASS	2.9
8	PASS	2.5
9	PASS	3.0
10	PASS	2.9

ings, Fig. 3.23, or when the colorectum is inflated due to disease, and the model instead optimises to attach to the colorectum, Fig. 3.24.

2. Instabilities are due to the fact that the spherical harmonics describing the surface are sufficiently complex that changing from $r(\theta, \phi) = \sum_{l=0}^{\infty} \sum_{m=-l}^l c_l^m Y_l^m(\theta, \phi)$ to $r(\theta, \phi) = \sum_{l=0}^{\infty} \sum_{m=-l}^l (c_l^m + \delta c_l^m) Y_l^m(\theta, \phi)$ can significantly and suddenly change the surface, that is, it is very numerically sensitive. That is, if a set of plausible constants are varied by 5% for example, the outline will no longer represent a likely surface shape. An example of the effect of this instability can be seen in Fig. 3.25 which is clearly not a plausible estimate of the shape.

The above results all refer to tests using 2nd order Fourier surfaces. Higher orders could be used to give a more exact representation of the curvature and intricacies of




Figure 3.23: Error due to unclear boundary.




Figure 3.24: Error in fitting to colorectum.




Figure 3.25: Instability due to brightness of bladder.

the shape of the mesorectum, however at higher order the results would be even more unstable. Robustness is therefore achieved using low order surfaces, and defaulting to the circular representation of the mesorectum in cases of instability.

3.2.4 Summary

The first circular estimate of the mesorectum using the coordinate frame of reference is an accurate and robust method of predicting the location and scale of the mesorectum. As such, the first estimate using a Fourier shape model, calculated from a training set, gives an accurate representation of the position of the mesorectum. This estimate can be refined by optimising the parameters, however this optimisation is not robust, and risks failing due to instabilities and false minima. Other shape models such as active appearance models that are more robust to noise, [4] could also be used here. However, such methods do require a large training set, and are more suited to anatomies that deform in predictable ways.

When the estimate is refined correctly the reduction in error of the outline is only reduced by a small amount (2mm), while the error of the first estimate (5mm) is still

of the same order as that of comparing outlines by two different radiologists. Hence this first estimate of the position of the mesorectum gives a robust outline such that we now know where to look for the lymph nodes in the image. Lymph node detection is described in detail in the next chapter.

Chapter 4

Node Detection and Analysis

The diagnosis of lymph nodes for colorectal cancer is crucial in making patient management decisions. The number of involved lymph nodes affects the staging, and hence the prognosis of the patient, and the location of the nodes is also important to know: if there are lymph nodes close to the circumferential resection margin then the likelihood of recurrence of the disease after surgery is greatly increased. Hence, the number and locations of lymph nodes affects the decision as to whether to give the patient downstaging chemo-radiotherapy prior to surgery.

Lymph nodes can be very small and difficult to detect on MRI scans. They can also be confused with blood vessels that can be of a similar size and shape in two dimensional slice images. The aim is to differentiate between the nodes and the blood vessels, and this will be done by utilising the entire three dimensional MRI dataset. In principle, when viewed in three dimensions, the lymph nodes are approximately spherical, whereas the blood vessels are tube-like structures, hence should be easily distinguished. However, in practice, this is rather more difficult due to the non-isotropic sampling of the MRI data. The lymph nodes are usually situated relatively close to the colorectum, and are within the mesorectal fascia (the area of fat surrounding the rectum). Since the boundaries of the colorectum and mesorectal fascia have already been determined by segmentation, the search region for the nodes is defined. A region of interest model is therefore incorporated into the search for the lymph nodes, as well as a probabilistic model based on the clarity of nodes in the




Figure 4.1: The area of the mesorectum surrounding the colorectum, where the lymph nodes are found.

images. Currently lymph nodes are detected manually by radiologists. As such, there is an associated risk of human error, particularly not spotting smaller nodes that may be cancerous. Automation enables such missed nodes to be picked out. The three dimensional representation also allows clinicians to spot any concentrations of lymph nodes within the mesorectal fascia which may be suspicious. The region of interest that constrains the search area for lymph nodes can be seen plotted in three dimensions in Fig. 4.1.

In order to automatically detect the nodes and blood vessels within the images, an appropriate segmentation method needs to be implemented. The method of snakes originally proposed by Terzopoulos *et al.* [62] is unsuitable for finding such small areas within the images since they are often only a few pixels in size. Hence the method of anti-geometric diffusion, proposed by Manay *et al.* [58] was considered as a possible solution. Allegedly, it enables the regions within the image to be segmented to varying scales depending on the regions of interest. It is a robust method and has been seen

to work on many types of images as a segmentation method [58]. The method is described in this chapter, and applied for general segmentation of colorectal cancer images. A refinement is then described for detecting lymph nodes and finally, these lymph nodes are analysed according to radiological criteria.

4.1 Anti-Geometric Diffusion for Image Segmentation

A fast and noise robust segmentation of medical images can be acquired using anti-geometric diffusion. If followed by a merging algorithm, it can be used to find areas of constant intensity within an image, as well as to find particular features of interest. The method is described in detail in this section

4.1.1 Anti-Geometric Diffusion and the Heat Flow Equation

Anisotropic diffusion is a common method for smoothing along image boundaries and was used as a pre-processing step for the bias-field removal in colorectal MR images, described in Chapter 2. It is particularly useful for de-noising images, but less effective for segmentation. However, by diffusing the image in the opposite manner, that is normal to the edges, then regions on opposite sides of the edges can be easily discriminated.

Edge directions are related to tangents of the iso-intensity contours on an image, Fig. 4.2. The direction normal to the iso-contour is defined by η , and ξ is the tangential direction. These directions can be written in terms of the first derivatives of the image I_x and I_y :

$$\eta = \frac{(I_x, I_y)}{\sqrt{I_x^2 + I_y^2}}, \quad (4.1)$$

$$\xi = \frac{(-I_y, I_x)}{\sqrt{I_x^2 + I_y^2}}. \quad (4.2)$$

Since η and ξ are orthogonal,




Figure 4.2: Normal and tangential directions of iso-intensity contour.

$$\frac{\partial I}{\partial t} = \nabla \cdot (\nabla I) = I_{\xi\xi} + I_{\eta\eta}, \quad (4.3)$$

which is the linear heat equation for $I(\xi, \eta, t)$.

If we omit the normal component and diffuse only in the tangential direction, then we have the geometric heat flow, as described previously. However, if we omit the tangential component and diffuse in the normal direction, then we have the anti-geometric heat flow, which is given by:

$$\frac{\partial I}{\partial t} = \frac{I_x^2 I_{xx} + 2I_x I_y I_{xy} + I_y^2 I_{yy}}{I_x^2 I_y^2}. \quad (4.4)$$

This can be extended for three dimensional volumes:

$$\frac{\partial I}{\partial t} = \left(\frac{\nabla I}{\|\nabla I\|} \right)^T \nabla^2 I \left(\frac{\nabla I}{\|\nabla I\|} \right), \quad (4.5)$$

where $\nabla^2 I$ is the Hessian.

4.1.2 Classification

In order to use this diffusion to perform a segmentation, we can classify the pixels during the diffusion process. In the first few iterations of the diffusion only those pixels near the object boundaries change significantly, yielding a confident classification as locally light or dark, depending on their diffusion behaviour. As diffusion continues, pixels further away can be classified. The more the image is diffused, the more likely that smaller features are lost, and hence the diffusion allows for a fine-to-coarse scale classification.

A short diffusion finds the well defined edges of most interest quickly, but leaves many pixels unclassified. The diffusion is run in this manner and then all the pixels are classified into three groups:

1. Pixels that increase significantly in intensity.
2. Pixels that decrease significantly in intensity.
3. Pixels that do not change significantly.

Each class of pixels is considered equally, and pixels are grouped into regions of adjacent, similarly classed pixels. Each of these connected regions R_i is given a unique class label i . This results in an oversegmentation, and hence the next stage for the segmentation is to merge similar regions.

4.1.3 Region Merging

Each region R_i has a mean intensity μ_i , and from this we can measure the segmentation error, $E = \sum_i E_i$, where $E_i = \sum_{R_i} (I - \mu_i)^2$. The error will normally increase if we merge two regions R_i and R_j to give $R_{ij} = R_i \cup R_j$, and replace the means by μ_{ij} . The change in error is $\Delta E_{ij} = E_{ij} - E_i - E_j$, where $E_{ij} = \sum_{R_{ij}} (I - \mu_{ij})^2$. Therefore a sensible criterion for merging R_i and R_j yields the smallest increase in ΔE_{ij} .

This merging can be repeated as many times as required or until there is a significant jump in ΔE_{ij} . A stopping criteria is required to suggest a number of iterations,




Figure 4.3: Initial image and oversegmentation using 3 classes.

or a particular minimum jump in ΔE_{ij} . This is fast to implement as we only ever need to consider combinations of adjacent regions.

An effective way to implement the region merging is to store all the details about each region R_i in a data structure, so for each R_i , we have:

1. A_i is the number of pixels within that region,
2. S_i is the sum of the intensities of those pixels,
3. Q_i is the sum of the squared intensities,
4. $\mu_i = S_i/A_i$ is the mean intensity of the region, and
5. $E_i = Q_i - 2\mu_i S_i + \mu_i^2 A_i$ is the sum squared error for the the region R_i .

We can then calculate the parameters for the proposed new merged region $R_{ij} = R_i \cup R_j$:

1. $A_{ij} = A_i + A_j$,
2. $S_{ij} = S_i + S_j$,
3. $Q_{ij} = Q_i + Q_j$,




Figure 4.4: Region merging to give 70 and 500 regions.

$$4. \mu_{ij} = S_{ij}/A_{ij},$$

$$5. E_{ij} = Q_{ij} - 2\mu_{ij}S_{ij} + \mu_{ij}^2 A_{ij},$$

$$6. \Delta E_{ij} = S_i\mu_i + S_j\mu_j - S_{ij}\mu_{ij}.$$

Results of the merging can be seen in Fig. 4.4, giving examples of merging to give 70 and 500 regions respectively. The former gives a coarse segmentation, in which only the large anatomical features can be seen, such as the bones and colorectum. The latter gives a finer segmentation, showing more of the details.

4.1.4 Shape Criteria

When searching for lymph nodes within images, we can add an extra term into the merging criterion in order to preserve regions of a certain shape. This can be done, for example, using eccentricity, or the second moment of area. The eccentricity of the i th region is given by:

$$\epsilon_i = \frac{\text{Second y moment}}{\text{Second x moment}} = \frac{\sum_{j=1}^p (y_j - y_{i\mu})^2}{\sum_{j=1}^p (x_j - x_{i\mu})^2} \quad (4.6)$$

where p is the number of pixels in the region, y_p is the y -value of the p th pixel, and $y_{i\mu}$ is the mean y -value of the pixels in region i .




Figure 4.5: Initial image showing nodes, and the segmentation with the nodes preserved, indicated by the arrows.

The error in the segmentation and merging that we are seeking to minimise then becomes:

$$\Delta E'_{ij} = (1 + e^{-c|\frac{\epsilon_i - \epsilon_d}{\epsilon_d}|})(1 + e^{-c|\frac{\epsilon_j - \epsilon_d}{\epsilon_d}|})(1 - e^{-c|\frac{\epsilon_{ij} - \epsilon_d}{\epsilon_d}|})\Delta E_{ij} \quad (4.7)$$

where c is an empirically chosen factor, ϵ_d is the desired eccentricity, which for lymph nodes is one, since we require them to be spherical, and ΔE_{ij} is the same as that given previously [57].

An example of using this criterion can be seen in Fig. 4.5. The nodes that were previously merged into the mesorectum are now preserved as features. The maximum runtime for this is 5 minutes and is dependent on the number of iterations performed.

4.1.5 Two Dimensional versus Three Dimensional Implementation

Although this method is easily extensible to three dimensions, we ran the region merging individually for two dimensional slices. This is due to the fact that the slice thickness is between 3mm and 3.3mm and, as such, can be larger than the node size. Nodes that occur close to the colorectum are often merged into the colorectum region at the first diffusion stage of the algorithm and so are not detected as separate regions.




Figure 4.6: Problem with three dimensional implementation.

This effect can be seen in Fig. 4.6.

4.2 Node detection

Since lymph nodes are only found within the mesorectal fat, we only need to search for them in the region inside the mesorectal fascia, and outside the colorectum. We can hence find all the regions that are within this region and then begin to classify them, as to whether or not they are in fact lymph nodes and whether they are likely to be involved (cancerous).

4.2.1 Classification Criteria

MRI has been shown to be a good imaging technique for lymph node analysis, as it is easier to tell the difference between lymph nodes and blood vessels than on CT [55]. Abnormal lymph nodes are almost always enlarged; however, enlarged lymph nodes are not always malignant. Therefore, there is a risk of false positives in the staging process.

Brown *et al.* [15] noted that different criteria for calling involved lymph nodes are used at different centres. They investigated 284 lymph nodes from 42 patients, using both MRI and the resected pathological specimens. They showed that as well as being enlarged, involved nodes are also characterised by heterogeneous signal intensity

and an irregular border. Using all of these criteria together would therefore reduce the number of false positives. Of course, it is also important not to diagnose false negatives, so we have developed a criterion that gives us an estimate of the probability of involved nodes, whilst still noting the positions of all the lymph nodes we find, even if they are thought not to be involved.

This classification is carried out using a series of different criteria, based on what we have learned from the collaborating clinicians about lymph nodes:

1. Size - Lymph nodes are small.
2. Intensity - Lymph nodes have a predominantly low intensity on T2-weighted MRI.
3. Heterogeneity - Involved lymph nodes often have a heterogeneous signal, that is visible as bright spots within the node.
4. Boundaries - Involved lymph nodes usually have irregular and indistinct boundaries.
5. Connectivity - Lymph nodes are usually only visible on a maximum of 3 slices, whereas blood vessels trace through a larger number of slices.
6. Shape - Lymph nodes are spherical.

4.2.2 Training and Likelihood graphs

Radiology reports from 10 datasets, pre- and post-chemo/radiotherapy were taken and the criteria for diagnosing the nodes were analysed. These reports contain both the conclusions that the radiologists arrived at in their diagnoses and their reasoning. A series of probability graphs, based on Gaussian distributions were drawn from this information so that we could numerically estimate the likelihood that any given region was in fact an involved lymph node.




Figure 4.7: Probability that a node is involved given its size.

Clinicians' analysis of likely nodes is much more qualitative than quantitative, hence obtaining statistical measures of the likelihood of involvement of lymph nodes is not straightforward. A typical analysis of lymph nodes in an image is as follows:

"There is a 5 mm mesorectal lymph node laterally on the left and other smaller posteriorly placed mesorectal nodes. At this size they are indeterminate."

The first criterion we analysed was the size. Lee *et al.* [55] and Brown *et al.* [15] both acknowledge that large lymph nodes can be a sign of malignancy and over 60% of involved nodes are 4mm diameter or greater, that is they are made up of 70 pixels or more.

Using the size of lymph node data from Brown *et al.* [15], we plotted the distribution of size of malignant lymph nodes and found that it fit a normal distribution, using the Jarque-Bera test [45] from the Matlab Statistics toolbox. Directly from this we allocated a Gaussian probability graph relating the size of lymph nodes to the likelihood of malignancy, and this can be seen in Fig. 4.7.

However, having noted that size does not by itself provide a sufficient criterion,




Figure 4.8: Probability that a node is involved given its mean intensity.

we then looked at the mean intensity of the node. Lymph nodes appear with low intensity on MRI and are distinctive against the high intensity fat where they are located. Using mean intensity as a criterion helps to root out any small regions that are perhaps due to the tumour itself, or just texture within the fat. We investigated the mean intensities of each of the regions picked out from the radiology reports as being lymph nodes. The shape of the histogram followed that of a normal cumulative distribution, and this can be seen in Fig. 4.8.

We looked very closely at the heterogeneity of the lymph nodes, since cancerous nodes often have bright spots within them due to the chemical composition of the tumour compared to that of normal lymph nodes, [52]. The boundaries of involved lymph nodes are often irregular and not distinct in shape, so this adds to the heterogeneity of the region, and as such the edge pixels are often different intensities to the rest of the region. These effects can be seen in Fig. 4.9. We can look at the overall heterogeneity of the signal by taking the standard deviations of the mean intensity within the node and the likelihood due to this criteria is shown in Fig. 4.10.

The connectivity of regions across a series of slices was analysed, as it is easy to

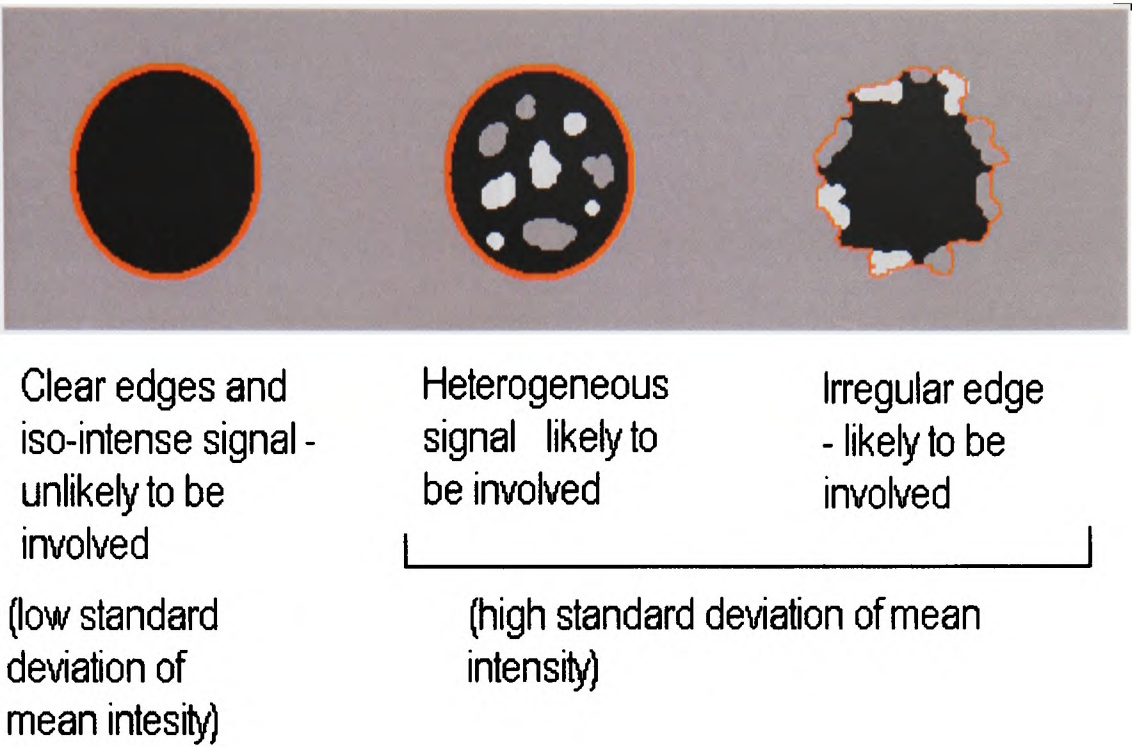


Figure 4.9: Nodes with heterogeneous signals and irregular edges.




Figure 4.10: Probability that a node is involved given its overall heterogeneity.




Figure 4.11: Visibility of nodes and blood vessels on a series of slices.

mistake lymph nodes for blood vessels. However, blood vessels that appear spherical in a two dimensional image, trace through a series of consecutive images. Lymph nodes, however, are usually only present on one or two images. This effect can be visualised in Fig. 4.11. In order to predict whether a region is a vessel therefore we look at whether or not regions overlap on consecutive images. The results of the distribution of lymph node size [52] are again useful in building a probability distribution of how likely the region is to be a lymph node based on how many images it is visible in. Lymph nodes that are 3mm or less, will only be visible in one slice, those between 3-6mm will be visible on one or two slices, and those between 6-9mm will be visible on two or three slices, and so on. Again plotting this distribution show that it fits the normal cumulative distribution seen in Fig. 4.12.

Finally, we looked at the shape of the region. Although lymph nodes can have irregular edges, their overall shape is usually spherical. In order to use this criterion, we take the second moment of area of the region, or eccentricity; and if it is near to one then the region is deemed to be spherical in two dimensions, whereas if the eccentricity is nearer to zero then the region is elongated and therefore much less likely to be a lymph node. The probability based on this is seen in Fig. 4.13.




Figure 4.12: Probability that a region is a node given the number of slices it is visible on.




Figure 4.13: Probability that a region is a node given its eccentricity.

4.2.3 Node Analysis

The size and intensity criteria are used initially to filter out larger and high intensity regions, so only regions that are of a low intensity and have a size of 2 - 12mm are considered.

The data structure for each region is then used to derive a probability that any given node is involved. As described above, the data structure contains information on the region's intensity and size. Eccentricity and position can be easily added to this structure and from these we can also calculate the heterogeneity and the connectivity.

Given a probability based on each criterion, an overall probability that any given region is an infected lymph node is then given by:

$$Pr(node) = Pr(node/I) * Pr(node/S) * Pr(node/H) * Pr(node/E) * Pr(node/C), \quad (4.8)$$

where $Pr(node)$ is the probability that any given region is an involved node, $Pr(node/I)$ is the probability of a node given the intensity, $Pr(node/S)$ is the probability of a node given the size, $Pr(node/H)$ is the probability of a node given the signal heterogeneity, $Pr(node/E)$ is the probability of a node given the eccentricity, and $Pr(node/C)$ is the probability of a node given the connectivity across the axial slices [6].

4.2.4 Visualisation

The overall probability is depicted for the small, low intensity regions using a red to blue colourmap. That is, those regions that have a high probability of being involved are coloured in red; and those that have a low probability are coloured in blue, Fig. 4.14. We indicate these probabilities on the MRI scans by circling the regions in their respective colour. Two examples of this can be seen in Figs. 4.15 and 4.16.

One of the most commonly encountered problems in detecting lymph nodes is to be able to differentiate them from blood vessels, that often look very similar on a two dimensional scan. By just taking the connectivity criterion into account, we can




Figure 4.14: Colourmap of likelihoods of involved lymph nodes.

segment out the regions that we suspect to be vessels, and this is seen in Fig. 4.17 where the blood vessels are colored in red and the other regions are indicated in green.

This example also provides a useful three dimensional visualisation of the information, and this same visualisation can be extended to use the colormap approach. Fig. 4.18 shows an example with all regions, whereas Fig. 4.19 just plots those regions with a 90% probability of being involved.

4.2.5 Results and Discussion

Validation of this algorithm is not straightforward as we do not currently have a gold standard to relate the MRI scans to. Therefore, we resorted to ‘validating’ the algorithm by comparing detection and classification results by the program against those of the radiologists. Two experienced radiologists independently assessed the images, and indicated on 10 datasets where all the nodes were and whether or not they were likely to be involved.

The algorithm picked out all the radiologists’ mesorectal nodes as regions to be analysed, and successfully discounted all the blood vessels as having a low probability of being involved. The radiologists only picked out the larger nodes ($>4\text{mm}$) within the images as it is almost impossible to estimate the likelihood of involvement of




Figure 4.15: Top left - Original image; Top right - Bias removed image; Bottom left - all regions plotted according to their probability of being involved nodes; Bottom right - Regions with a 90% probability of being involved.




Figure 4.16: Left - all regions plotted according to their probability of being involved nodes; Right - Regions with a 90% probability of being involved.




Figure 4.17: A typical three dimensional representation of the blood vessels in the mesorectum.




Figure 4.18: A three dimensional representation of the small regions in the mesorectum, coloured according to their probabilities of being an involved lymph node.

smaller nodes by just looking at the images due to the relatively low spatial resolution.

We found that our algorithm picked out 20% more regions that were likely to be involved nodes than the radiologists, and hence it is possible that our algorithm provides more false positives. Most of the extra nodes that our algorithm found were small, and, after consultation with the radiologists about these extra possible nodes, they commented that they were too small to analyse with any confidence, but they could not rule out the possibility that they were involved.

When the pathologists dissect the specimen, they usually find many more nodes than the radiologists report, due to the limited resolution of the images, and the fact that radiologists tend only to report very large heterogeneous nodes that are close to the circumferential resection margin. That is, they only report the nodes that they can confidently say they think are involved.

A full validation would require us to compare our algorithm with the results from the pathology; but, in the interim, it provides a good second estimate for radiologists




Figure 4.19: A three dimensional representation of the small regions in the mesorectum that have a 90% probability of being an involved lymph node.

to pick out any nodes that they may miss. Carrying out this full validation is not straightforward: It would first require the pathologist to label the position of every lymph node they find within the specimen, so that it can subsequently be related back to the MRI, by registering the MR data with the histology. This pathology information of position of infected lymph nodes is not currently available, but is a large area for future work.

4.2.6 Summary

The algorithm to detect and estimate the likelihood of involvement was thought to be a useful tool by the clinicians, providing further nodes for the radiologist to look at that they may have missed. In the long term however this does need to be validated against ‘gold standard’ histology data.

Chapter 5

Registration

As part of the treatment of colorectal cancer, currently 65% of patients are given courses of chemotherapy or radiotherapy. This is particularly likely to be the case if the patient has suspected involved lymph nodes within the mesorectum. Chemotherapy and radiotherapy serve the purpose of reducing the size of the tumour, as well as shrinking any involved lymph nodes by killing the cancer cells within them. This, in turn, serves the purpose of reducing the risk of metastatic disease. Since both nodal involvement and tumour size often decrease, so does the assessment of tumour stage, so that pre-resection chemotherapy is often referred to as down-staging chemotherapy. Chemotherapy and radiotherapy always cause dramatic changes in the anatomy. The body's defence mechanism against tumours include swelling by aggregation of fluid in the area surrounding the tumour (adenosis), and this is much reduced by the treatment. It is very difficult, therefore, to match up the images pre- and post-treatment, motivating the need for registration. After treatment, further patient management decisions need to be made as to whether or not to proceed to surgery. Surgery would be unnecessary if the tumour had all been killed, and since many of the patients are elderly, surgery is not always a desirable option. High Intensity Focussed Ultrasound (HIFU) is increasingly an option for treatment of cancers, especially in Oxford. This involves the use of a machine that gives off high frequency sound waves. These waves are focussed to deliver a strong beam to a specific part of a cancer, in order to kill the cancer cells [29]. Many trials are currently underway into the use of this treatment




Figure 5.1: Equivalent slices pre- and post-therapy.

for different types of tumour.

In order to make this patient management decision, a further set of MRI scans are performed to look at the changes that occurred over the course of the treatment. The size and shape of the tumour and lymph nodes are carefully analysed. There are often many lymph nodes in the images and it is important that the corresponding lymph nodes are matched up correctly. However, since the scans are taken up to six months apart and the patients have had treatment in between, the images before and after typically vary dramatically, as shown in Fig. 5.1. Therefore, registration is essential. The purpose of this review is to analyse the effectiveness and accuracy of using affine, rigid and non-rigid registration for this purpose to see whether it is sufficiently accurate. We start by discussing how best to validate the registration, then describe the affine and non-rigid methods for performing registration, followed by the results of using these methods. Finally we discuss what the current methods appear to be lacking.

5.1 Validation

Reviewing the success of non-rigid registration algorithms is not a straightforward process, and many of the papers in the literature use differing methods. A review of some of these methods is provided by Crum *et al.* [26]. Often, there is no right or wrong answer, since the result depends greatly on exactly what we are registering and the required outcome of the registration. For example, if we are registering two healthy patients to each other we may require a perfect match across the entire image; but if we are registering a diseased patient to a healthy one, we will want to be able to see and quantify what the differences are. As such, we need to take into account what it is we are registering before we can validate whether or not the algorithm ‘works’. Some common validation measures are described below.

5.1.1 Statistical Measures

A number of statistical measures have been proposed which measure the correspondence between the registered source image and the target image. These can be based on intensity if the images being registered are both from the same modality [75], or can use other measures such as local phase [64]. These statistical measures vary in complexity, from using the maximum difference errors, or sum squared errors (SSD), to more complex measures such as correlation coefficients (CC). The SSD is defined as:

$$SSD = \frac{1}{n} \sqrt{\sum (I(x', y', z', t_0) - I(\mathbf{T}(x, y, z), t_1))^2} \quad (5.1)$$

while the CC is defined as:

$$CC = \frac{\sum (I(x', y', z', t_0) - \bar{I}(t_0))(I(\mathbf{T}(x, y, z), t_1) - \bar{I}(t_1))}{\sqrt{\sum (I(x', y', z', t_0) - \bar{I}(t_0))^2 \sum (I(\mathbf{T}(x, y, z), t_1) - \bar{I}(t_1))^2}} \quad (5.2)$$

In this equation, $\bar{I}(t_0)$ and $\bar{I}(t_1)$ denote the average intensities of the image before and after the treatment, respectively.

Inverse consistency measure, and transitivity error [39] can also be used as statistical measures of validation to ensure that the registration results are both plausible and sensible.

These validation methods are useful when there is a definite correspondence between the two images, but are considerably less useful when we are looking at images that may have changed significantly due to patient differences, differences in disease, or, in our case, differences over time due to treatment.

5.1.2 Simulations and Phantoms

Simulations provide a very helpful way to begin to validate registration results. A ‘typical’ warp field is often applied to the dataset, such that the original or ‘ideal’ image result is known. This is then a ‘gold standard’ correct answer that can be used to compare to. By comparing the calculated registration result to the ‘gold standard’, the statistical methods can be used to provide an indication of performance.

Phantoms provide a similar method of validation in that the ‘gold standard’ is known. A phantom is an object that is scanned and looks similar in some way to real data. However, the structure of the phantom is obviously known, so this information can be used as a comparison in order to assess the registration algorithm.

Although very useful, such methods are not sufficient, since the registration also needs to be tested on real data for which the gold standard is not known. In some cases, simulations and phantoms can be overly simplistic. For example, modeling tumors or warps as circles that expand outwards in a consistent manner [66] is very unlikely in real clinical data. Finding an appropriate model for building simulations and phantom experiments is not trivial and the characteristics of the medical problem need to be modelled in an intelligent and well thought out manner.

5.1.3 Fiducials

Fiducials are often used to provide markers by which the results of registration can be monitored. Such fiducials can be physical implants which will show up clearly

on the imaging modalities being used. Differences and errors in their locations after registration can then be calculated.

However, in practice this is rarely possible, and so distinctive point features within the images, or on the surface of the object being imaged are used. These are often anatomically significant features and, as such, provide a good method of validation. They also have the advantage of being specific to the individual being treated. However, such features do need to be chosen carefully as those features with a high entropy (Section 5.3.1) are more likely to be successfully aligned by an algorithm, whereas less prominent features may remain misaligned. Therefore, it would be easy to show an algorithm performing well with this measure when in fact only a subset of very distinctive points on the images are correctly registered.

5.1.4 Validation for pre- and post-treatment image alignment

Taking into account the factors discussed above, we must choose carefully how we are to validate the registration of the pre- and post-treatment images, bearing in mind the medical problem we are studying.

A statistical method such as SSD or CC will not be appropriate, as we expect, indeed hope, that there will be intensity and shape changes within the images. Textures of some of the tissue types are also expected to change after treatment, in particular the fat in the mesorectum is likely to become ‘streaky’ in appearance. We also expect there to be shape changes in terms of the lymph nodes appearing, disappearing or changing size due to treatment, and we do not want this fact to count against the validation of success of the algorithm.

Matching features within the images was therefore seen as the best method by which to validate the registration. Our shape model uses the most distinctive objects in the images, hence it would not be a good test of accuracy to use the alignment of these features. For these reasons we used the alignment of the lymph nodes as a measure of accuracy. These are features that are very useful to the clinicians in




Figure 5.2: Correspondance of lymph nodes in pre- and post-treatment images.

their patient management decisions, but they do vary between patients. They also vary as to whether they appear in both pre- and post-treatment images. We have used the radiologists' reports to identify corresponding lymph nodes in pre- and post-treatment images and then we measure how well they are aligned, giving a mean and standard deviation for each dataset. A typical dataset will contain at least five, and possibly up to thirty, corresponding lymph nodes of varying sizes (although they may not all be involved in the cancer). An example of corresponding nodes can be seen in equivalent pre- and post-treatment images in Fig. 5.2.

5.2 Affine and Rigid Registration Methods

5.2.1 Theory

Registration is the process of finding a mapping from coordinates in one space to those in another, to achieve biological, anatomical or functional correspondence. Such mappings, or transformations, are defined by the number of degrees of freedom (DOF) they have. A rigid transformation has 6 DOF (translation (3) and rotation (3)), whereas an affine transformation has 12 DOF (translation (3), rotation (3), scaling (3) and shear (3)). Transformations are best described in a matrix formation. For

rigid registration, the points (x, y, z) will map onto (x', y', z') by the following transformation:

$$T_{rigid}(x, y, z) = \begin{pmatrix} x' \\ y' \\ z' \\ 1 \end{pmatrix} = \begin{pmatrix} r_{11} & r_{12} & r_{13} & t_x \\ r_{21} & r_{22} & r_{23} & t_y \\ r_{31} & r_{32} & r_{33} & t_z \\ 0 & 0 & 0 & 1 \end{pmatrix} \begin{pmatrix} x \\ y \\ z \\ 1 \end{pmatrix} \quad (5.3)$$

where t_x , t_y and t_z are translation vector components and r_{11}, \dots, r_{33} are orthogonal components.

Similarly for an affine transformation:

$$T_{affine}(x, y, z) = \begin{pmatrix} x' \\ y' \\ z' \\ 1 \end{pmatrix} = T_{shear} \cdot T_{scale} \cdot T_{rigid} \cdot \begin{pmatrix} x \\ y \\ z \\ 1 \end{pmatrix} = \begin{pmatrix} a_{11} & a_{12} & a_{13} & t_x \\ a_{21} & a_{22} & a_{23} & t_y \\ a_{31} & a_{32} & a_{33} & t_z \\ 0 & 0 & 0 & 1 \end{pmatrix} \begin{pmatrix} x \\ y \\ z \\ 1 \end{pmatrix} \quad (5.4)$$

where t_x , t_y and t_z are translation vector components and a_{11}, \dots, a_{33} are affine matrix components.

One method of finding these transformations is to use feature-based registration. This is the method that appears to be most appropriate for registering the colorectal images, since there are a number of features within the images that may be assumed to remain in fixed positions despite the chemotherapy and radiotherapy. It is worth noting at this point that even if rigid/affine registration proves insufficient, all non-rigid registration algorithms (e.g. Rueckert's B-spline algorithm [75]) usually start with an affine registration to ensure that the ensuing non-rigid method begins in a good starting position.

Feature-based registration requires two stages, feature extraction and feature alignment. The rigid features within the images can be easily identified as point features (more details in Chapter 6) and then point-based registration can be used to align them. The algorithm used for this is the Procrustes algorithm [24], which performs optimal fitting using least squares, defined as follows.

Given two sets of N corresponding points $\{x_i\}$ and $\{y_i\}$, the transformation matrix

that minimises the mean squared distance between the points is found from:

$$FRE^2 = \frac{1}{N} \sum_{i=1}^N \|Rx_i + t + y_i\|^2 \quad (5.5)$$

where FRE is the fiducial registration error, R is the rotation matrix and t is the translation vector.

The algorithm then proceeds as follows:

1. Centre the points

- (a) Compute the mean: $\bar{x} = \frac{1}{N} \sum_{i=1}^N x_i$, $\bar{y} = \frac{1}{N} \sum_{i=1}^N y_i$

- (b) Centre with respect to the mean: $\tilde{x}_i = x_i - \bar{x}$, $\tilde{y}_i = y_i - \bar{y}$

2. Determine the rotation matrix R by calculating the singular value decomposition (SVD) of the correlation matrix H .

$$H \equiv \sum_{i=1}^N \tilde{x}_i \tilde{x}_i^T$$

$$H = U\Lambda V^T, \text{ where } U^T U = VV^T = I, \Lambda = \text{diag}(\lambda_1, \lambda_2, \lambda_3), \lambda_1 \geq \lambda_2 \geq \lambda_3.$$

$$R = VDU^T, \text{ where } D = \text{diag}(1, 1, \det(VU^T)).$$

3. Determine the translation vector, t .

$$t = \bar{y} - R\bar{x}$$

This will give the rigid transformation matrix:

$$\begin{pmatrix} R & t \\ 0 & 0 & 0 & 1 \end{pmatrix} \quad (5.6)$$

Alternatively, the affine transformation matrix can be found by direct evaluation of three landmark points (features), by direct substitution. However, the points that are expected to move rigidly within the image are unlikely to be scaled or sheared, and hence rigid registration will have the same overall effect.

Three landmark points are the minimum required for the Procrustes algorithm. However, it is worth validating these 3 points as parts of the anatomy that will move rigidly. This is done by selecting more than the minimum number of landmark points and then performing the RANSAC (Random Sample Consensus) algorithm, developed by Fischler and Bolles [37]. This algorithm will eliminate any outliers and give an estimate of the error margins in the landmark points. The algorithm works as follows:

The minimum number of landmark points is n ($n = 3$ in this case) and there is a dataset of m landmark points, where $m > n$.

1. Select n data items at random.
2. Calculate the parameters R and t as described using the Procrustes algorithm above.
3. Find how many of the other data points, K , fit into this transformation model within a given tolerance.
4. If K is sufficiently large, accept fit and exit with success.
5. Repeat steps 1-4 L times, where L is related to probability of failure.
6. Fail at this point.

It is unlikely for this application that the RANSAC algorithm will fail, due to the reasonably accurate placement of landmarks on a small number of rigidly moving parts of the anatomy. Instead, it will find a good rigid fit within a certain tolerance and eliminate any outliers in the landmark data points. It can also be used to loop through all the successful iterations to find the best rigid transformation. From this, it can be determined how good the best fit actually is, by calculating the number of outliers for the given tolerance.



Figure 5.3: Rigidly Moving Anatomical Features in Colorectal Images

5.2.2 Applying Affine Registration to Colorectal Cancer Images

The main features that move rigidly in the colorectal images, that is they remain in constant positions despite the chemotherapy and radiotherapy, are the bone structures: the hips and the coccyx. These can be translated into three point features by taking the centre point of the hips and the lowest point of the coccyx. These can be found very easily on the axial small field of view images, using the Mirada software package, RevealMVSTM, as shown in Fig. 5.3.

Ten patients were compared pre- and post-chemo/radiotherapy and the axial small field of view datasets were aligned using rigid registration transformations, found by direct evaluation using Procrustes algorithm. In order to validate the transforma-




Figure 5.4: Rigidly Moving Landmarks in Colorectal Images

tions, a number of lymph nodes from each dataset were observed. Their positions on the pre-therapy image were recorded, along with their registered post-therapy positions. These positions were then compared in order to analyse the effectiveness of the registration. The misalignment of the lymph nodes was up to 22.5mm, with a mean standard deviation of 6mm across the ten datasets. This is unacceptable in terms of clinical application, since it will still take the clinicians a long time to find corresponding nodes with any certainty.

The above was repeated using a larger set of landmarks in the images. The landmarks are other easily identifiable points on the rigid structures of the hips and coccyx, making 17 landmarks in total on each dataset drawn in three dimensions in Fig. 5.4.

A RANSAC analysis was then performed to identify any outliers among these

Table 5.1: Node Misalignment (mm) using Affine Registration.

Pa- tient	Misalignment of Lymph Nodes before registration(Mean \pm Std)	Misalignment of Lymph Nodes after affine registration(Mean \pm Std)
1	49.85 \pm 4.71	26.85 \pm 7.00
2	89.80 \pm 9.03	21.65 \pm 9.00
3	25.90 \pm 5.23	15.00 \pm 3.17
4	13.70 \pm 4.42	16.31 \pm 4.41
5	38.01 \pm 11.79	17.50 \pm 4.92
6	18.30 \pm 5.27	17.66 \pm 4.03
7	69.00 \pm 6.05	36.21 \pm 6.14
8	110.90 \pm 13.29	30.00 \pm 3.65
9	25.53 \pm 11.50	17.34 \pm 4.55
10	20.37 \pm 7.12	15.63 \pm 6.00

landmarks and hence find the best rigid registration. The transformations were found, and the number of outliers eliminated and the typical misalignment of the lymph nodes was 21mm, with a mean standard deviation of 5.5mm across the ten datasets, Table 5.1. The mean number of outliers was 3, and ranged from 0 to 7.

The affine registrations were all performed on the images using the Siemens Molecular Imaging (formerly Mirada Solutions, Ltd) RevealMVSTM software. A typical result of this can be seen in the Fig. 5.5. Although the rigid points, such as the hips and coccyx now line up, the rest of the image does not. It can be clearly seen that the deformation of the bladder, colorectum and mesorectum are too great to be aligned by an affine transformation.

5.2.3 Summary

The results above show that, in general, the lymph nodes do not align well as a result of affine registration. They can move as much as 2 cm with the mesorectal fat on successfully affine registered images (indicated by having 25% (4/17) or less outliers). Much of this movement is due to changes in the bowel and bladder, as well as the shrinking of the tumour. This movement is non-rigid, hence its extent cannot be predicted with sufficient accuracy. Radiotherapy, in particular, causes many changes in the structure of the area it is focussed on. Finding the best possible




Figure 5.5: Affine Registration using RevealMVS, showing the target image at the top, and the target image with the overlaid registered source image at the bottom.

rigid registration using RANSAC makes little difference to the movement of the nodes. Although the landmarks align well, this does not cause the nodes to align in the same way.

The rigidity of the ‘rigidly moving’ landmarks can also be questioned: The coccyx, for example, can in fact move relative to the hips, depending on the positioning of the patient lying in the scanner. In practice, the likelihood of this is minimal as the posture of any given patient within the scanner will not vary that much. However in cases where this does occur, the nodes obviously will not align well.

Since, despite affine registration, the nodes still do not align well, a non-rigid registration is required in order to align the nodes in a more effective way, so that the radiologist can quickly and easily find corresponding lymph nodes in post-treatment images. There are many methods that claim to perform non-rigid registration. However, it is important to ensure that as much anatomical knowledge as possible can be integrated into the method to aid robustness, efficiency and accuracy. An affine registration using landmarks, though not sufficient itself, does perform a global alignment of the images, which is a useful step to perform prior to non-rigid registration.

5.3 Non-Rigid Registration Methods

Rigid and affine registration methods have proved to be insufficient to register pre- and post-chemotherapy and radiotherapy images. Although they may lead to (poor) local minima [35], they do provide a reasonable global alignment that can be used as a first step in non-rigid registration algorithms. A number of non-rigid methods are now reviewed in order to see how effective they are on this dataset. The theory behind each method is first summarised, then the results on the colorectal cancer datasets are presented. A prior global affine alignment is assumed for each algorithm.

5.3.1 Theory

Non-rigid registration methods comprise four elements:

1. Image Descriptor: This is usually intensity, but need not be. For example,

Mellor and Brady [64] have recently advocated the use of local phase as a descriptor that is well suited to multi-modal image registration. Though Mellor and Brady did not claim substantial advantages for using local phase for mono-modal registration (the case here), its use does mean that intensity artifacts have less effect on the registration, and this can reduce error in the algorithms.

2. Transformation Space: These are based on basis functions and can be rigid, affine, elastic, fluid, etc.
3. Similarity Metric: This measures the similarity between the apparently corresponding points in the two images. Several different similarity metrics have been proposed, including cross-correlation, mean squared distance, mutual information, and the correlation ratio. Mutual Information has been the most used similarity metric. It maximises a probabilistic relationship between the image intensities. The relationship is not explicitly defined but is determined at least partly by the transformation model and optimisation strategy chosen in conjunction with the data. It can therefore allow for changes in intensity across pairs of multi-modal and contrast images [43]. In our case, the intensities change primarily due to chemotherapy and radiotherapy.

It maximises a probabilistic relationship between the image intensities. The relationship is not explicitly defined but is determined at least partly by the transformation model and optimisation strategy chosen in conjunction with the data.

4. Optimization Method: This searches for the best transformation to minimize the similarity metric; methods range from straightforward numerical descent procedures to highly complex optimization procedures.

Given the intrinsic complexity of medical images and the often substantial non-linear nature of non-rigid registration, it is not surprising that most algorithms iterate to find the combination of these four elements that gives the best possible result in

any given application. Some of the elements of non-rigid registration are described in more detail below, along with some complete algorithms that seek to combine the four elements to give an optimum registration result.

Image Descriptor: Phase-based methods

Mellor *et al.* [64] use local phase for a feature-based registration method. This method is therefore much more robust to image artefacts, as it is independent of intensity. The local phase and energy of the image are estimated using the monogenic signal, which consists of local energy, local phase and orientation. It can be computed as the output of three filters. The first is a rotationally symmetric, zero mean filter that gives the bandpass image I_b (the even component of the signal). The odd component is the response of two anti-symmetric filters to the even part. These two filters are described in the Fourier domain as:

$$H_1(u, v) = \frac{u}{\sqrt{u^2 + v^2}}, \quad (5.7)$$

$$H_2(u, v) = \frac{v}{\sqrt{u^2 + v^2}}, \quad (5.8)$$

where u and v are the Fourier domain coordinates. Local phase, ϕ , local orientation, θ , and local energy, E , can then be calculated:

$$\phi(x, y) = \tan^{-1}\left(\frac{I_b}{\sqrt{(h_1 \otimes I_b)^2 + (h_2 \otimes I_b)^2}}\right), \quad (5.9)$$

$$\theta(x, y) = \tan^{-1}\left(\frac{h_1 \otimes I_b}{h_2 \otimes I_b}\right), \quad (5.10)$$

$$E(x, y) = \sqrt{I_b^2 + (h_1 \otimes I_b)^2 + (h_2 \otimes I_b)^2}. \quad (5.11)$$

Any non-rigid registration method that maximises the mutual information can then be used to maximise the phase mutual information.

Similarity Metric: Mutual Information

When registering two sets of images, intensity values are often not the best measures to use as a basis for comparison. Often, multi-modal images, or images formed

using contrast-agents are compared and these will have different intensity variations. Mutual information, originally proposed by Viola and Wells [88], requires no prior information about the changes between the two datasets, but is based on the idea that there exists a functional relation between the intensities that is not known.

Therefore, this similarity measure can be useful in the case of colorectal images, particularly when considering the mesorectal region, where the density changes due to treatment will cause intensity changes. Registration can be performed by estimating a transform based on the maximisation of the mutual information:

$$\hat{T} = \arg \max_T MI(u(x), v(T(x))). \quad (5.12)$$

Mutual information itself is defined in terms of entropy in the following way:

$$MI(u(x), v(T(x))) \equiv H(u(x)) + H(v(T(x))) - H(u(x), v(T(x))). \quad (5.13)$$

H is the entropy of a random variable and is defined as $H(x) = - \int p(x) \ln p(x) dx$, while the joint entropy of two random variables, x and y , is $H(x, y) = - \int p(x, y) \ln p(x, y) dx dy$. This may be interpreted as a measure of uncertainty, variability or complexity. Entropy cannot be found directly using the above expression but it can be expressed as $H(z) = E_z(-\ln p(z))$. Then to estimate the entropies, the Parzen windowing technique (for example) can be used such that $p(z) \approx \frac{1}{N_A} \sum_{z \in A} G_\Psi(z - z_j)$, where $G_\Psi(x) = (2\pi)^{-\pi/2} |\Psi|^{-1/2} \exp(-\frac{1}{2} x^T \Psi^{-1} x)$, and N_A is the number of trials. However, it should be noted here that Joshi *et al.* [46] found that this is, in fact, not a good approximation to the real histogram, largely because of the partial volume effect, discussed in Chapter 2.

Next, the statistical expectation is approximated with the sample average:

$$E_z(f(z)) \approx \frac{1}{N_B} \sum_{z_i \in B} f(z_i). \quad (5.14)$$

An approximation for the entropy of a random variable can now be written as:

$$H(z) \approx \frac{-1}{N_B} \sum_{z_i \in B} \ln \frac{1}{N_A} \sum_{z_j \in A} G_\Psi(z_i - z_j). \quad (5.15)$$

In order to find the maximum of mutual information, the derivative of the entropy is calculated with respect to the transformation parameters T .

$$\frac{d}{dT} H(z(T)) \approx \frac{1}{N_B} \sum_{z_i \in B} \sum_{z_j \in A} W_z(z_i, z_j) (z_i - z_j)^T \Psi^{-1} \frac{d}{dT} (z_i - z_j), \quad (5.16)$$

where $W_z(z_i, z_j) = \frac{G_\Psi(z_i - z_j)}{\sum_{z_k \in A} G_\Psi(z_i - z_k)}$ is between 0 and 1.

The derivative of the mutual information is also required to find the maximum of the MI , and this is defined as:

$$\frac{d}{dT} I(u(x), v(T(x))) = \frac{d}{dT} H(v(T(x))) - \frac{d}{dT} H(u(x), v(T(x))). \quad (5.17)$$

An estimate for this can be obtained as:

$$\frac{d\hat{I}}{dT} = \frac{1}{N_B} \sum_{x_i \in A} \sum_{x_j \in B} (v_i - v_j)^T [W_v(v_i, v_j) \Psi_v^{-1} - W_{uv}(w_i, w_j) \Psi_{vv}^{-1}] \frac{d}{dT} (v_i - v_j), \quad (5.18)$$

where the weighting factors are defined as $W_v(v_i, v_j) = \frac{G_{\Psi_v}(v_i - v_j)}{\sum_{x_k \in A} G_{\Psi_v}(v_i - v_k)}$ and $W_{uv}(w_i, w_j) = \frac{G_{\Psi_{uv}}(w_i - w_j)}{\sum_{x_k \in A} G_{\Psi_{uv}}(w_i - w_k)}$, and $u_i = u(x_i)$, $v_i = v(T(x_i))$, $w_i = [u_i, v_i]^T$.

A stochastic analog of the gradient descent algorithm [56] is then used to find this maximum of the mutual information, and this will converge at the optimum registration.

Algorithm: B-Spline Algorithm

Rueckert *et al.* proposed a non-rigid registration method whereby the global motion is modelled by an affine transformation and the local deformation is modelled using free-form deformations (FFDs) based on B-splines. FFDs are a common technique used in computer graphics for modelling three dimensional deformable objects, and can be described by an $n_x \times n_y \times n_z$ mesh of control points Φ with spacing $\delta_x \times \delta_y \times \delta_z$.

The combined transformation is given by $T(x, y, z) = T_{global}(x, y, z) + T_{local}(x, y, z)$. The global transformation is the affine part, described previously, and the local transformation is then given by:

$$T_{local}(x, y, z) = \sum_{l=0}^3 \sum_{m=0}^3 \sum_{n=0}^3 B_l(u) B_m(v) B_n(w) \phi_{i+l, j+m, k+n}, \quad (5.19)$$

where $i = \lfloor \frac{x}{n_x} \rfloor - 1$, $j = \lfloor \frac{y}{n_y} \rfloor - 1$, $k = \lfloor \frac{z}{n_z} \rfloor - 1$, $u = \frac{x}{n_x} - \lfloor \frac{x}{n_x} \rfloor$, $v = \frac{y}{n_y} - \lfloor \frac{y}{n_y} \rfloor$, $w = \frac{z}{n_z} - \lfloor \frac{z}{n_z} \rfloor$, and B_l is the l^{th} basis function of a B-spline,

$$\begin{aligned} B_0(u) &= (1 - u)^3/6 \\ B_1(u) &= (3u^3 - 6u^2 + 4)/6 \\ B_2(u) &= (-3u^3 + 3u^2 + 3u + 1)/6 \\ B_3(u) &= u^3/6 \end{aligned} \quad (5.20)$$

The control points Φ act as the parameters of the B-splines and the degree of non-rigid registration depends on the resolution of the mesh of control points. Rueckert proposes a hierarchical multi-resolution approach where Φ^1, \dots, Φ^L are a hierarchy of control point meshes. Therefore, each control point mesh and its associated FFD defines a local transformation at that resolution such that:

$$T_{local}(x, y, z) = \sum_{l=1}^L T_{local}^l(x, y, z). \quad (5.21)$$

To constrain the spline based FFDs to be smooth, a penalty term is introduced that regularises the transformation. In three dimensions this is:

$$\begin{aligned} C_{smooth} = \frac{1}{V} \int_0^X \int_0^Y \int_0^Z & \left[\left(\frac{\partial^2 T}{\partial x^2} \right)^2 + \left(\frac{\partial^2 T}{\partial y^2} \right)^2 + \left(\frac{\partial^2 T}{\partial z^2} \right)^2 + \right. \\ & \left. + 2 \left(\frac{\partial^2 T}{\partial x \partial y} \right)^2 + 2 \left(\frac{\partial^2 T}{\partial x \partial z} \right)^2 + 2 \left(\frac{\partial^2 T}{\partial y \partial z} \right)^2 \right] dx dy dz \end{aligned} \quad (5.22)$$

where V is the volume in the image domain. Note that C_{smooth} is zero for an affine transformation, therefore only penalises non-affine transformations. Rueckert *et al.*

use normalised mutual information as the similarity criterion instead of intensities, as intensities can change depending on circumstances (for example if a contrast agent is being used). If the images are aligned, then the normalised mutual information between the two images A and B is maximised:

$$C_{similarity}(A, B) = \frac{H(A) + H(B)}{H(A, B)}. \quad (5.23)$$

An optimisation is used that minimises a cost function associated with the global transformation parameters (Θ) and the local transformation parameters (Φ),

$$C(\Theta, \Phi) = -C_{similarity}(I(t_0), T(I(t))) + \lambda C_{smooth}(T), \quad (5.24)$$

where λ is a weighting parameter that acts as a trade-off between alignment and smoothness of the transformation. This weighting is usually determined experimentally. An iterative technique such as the gradient descent algorithm is then used to find the optimum non-rigid transformation.

Algorithm: Adaptive Registration using Local Information Measures

Park and Meyer [66] use an adaptive registration method whereby control points are only added where they appear to be needed. They use Thin Plate Splines (TPS) to model the deformation field. These splines are placed on ordered pairs of control points on corresponding loci of the two datasets that are being registered. Assuming that x denotes the reference coordinates, and x' denotes the homologous coordinates, and that there are N control point pairs, then the formulation of the TPS is:

$$T(x) = a_0 + a_1 x + \sum_{i=1}^N w_i U(r_i), \quad (5.25)$$

where r_i is the Euclidean distance between x and x_i , a_0 and a_1 are the affine parameters, w_i is the warp coefficient and $U(r)$ is the basis function,

$$U(r) = \begin{cases} r^2 \log(r^2) & \text{in 2D,} \\ |r| & \text{in 3D.} \end{cases} \quad (5.26)$$

Mutual information (MI) is used as the similarity measure, so that for the two datasets A and B ,

$$MI(A, B) = \sum \sum p(a, b) \log_2(p(a, b)/p(a)p(b)) \quad (5.27)$$

where $p(a)$ and $p(b)$ are marginal densities of A and B , and $p(a, b)$ is the joint density.

The goal of the Park-Meyer algorithm is to increase the DOF selectively in the required regions, these regions being identified using the local mismatch measure. Local areas of interest for grid refinement are those with a high local entropy and low local mutual information. The normalised measure of mismatch, M , is defined:

$$M = 1 - \frac{MI(a, b)}{\min(H(a), H(b))}. \quad (5.28)$$

All the local measures of MI and H are computed over a finite sub-block and not the whole dataset.

The algorithm solves for the registration transformation iteratively by adding one control point at a time to the previous geometric solution using the local mismatch measure M . An affine registration is first applied to globally align the images, then a sparse set of control points is used to initialise the registration. The location with the maximum mismatch is identified and a control point is added at this location. The control points are optimised to maximise the global MI, and then the process is repeated iteratively until the global MI increase is greater than some threshold. The mismatch measure can detect deformations reliably for those whose regions of effect are similar to the size of the sub-block. Hence the sub-block size has to be chosen carefully.

One of the main features of this algorithm, compared to that of Rueckert *et al.* [75] is that the number of control points remains small. In Rueckert's algorithm, the number of control points can increase by 2^n (where n is the number of dimensions) at each scale. The net effect for Park-Meyer's algorithm is a deformation grid in which several different scales are presented in a single grid.

Algorithm: Adaptive Bases Algorithm for Intensity-Based Non-Rigid Registration

Rohde *et al.* [74] use spatial adaptation in order to reduce the number of degrees of freedom by introducing a number of different concepts. First, they use radially symmetric basis functions rather than B-splines. Second, they use a new metric to identify which regions are poorly registered. Third, they partition the global registration problem into several smaller ones and finally introduce a constraint that ensures the transformation is topologically correct. They perform registration by building the deformation field incrementally, region by region, and hence can focus on the misregistered regions. The total deformation field $v(x)$ is modelled as a linear combination of basis functions irregularly spaced over the image domain:

$$v(x) = \sum_{i=1}^N c_i \Phi(x - x_i). \quad (5.29)$$

Wu's compactly supported positive definite radial basis functions [91] are used:

$$\Phi(x) = \phi\left(\frac{\|x\|_2}{s}\right), \quad x \in \mathbf{R}^d, \quad (5.30)$$

where d is the dimension, $\phi(r) = (1 - r)_+^4 (3r^3 + 12r^2 + 16r + 4)$ for $r \geq 0$, where $(1 - r)_+ = \max(1 - r, 0)$, and s is a scale parameter.

An iterative multi-scale and multi-resolution approach is used. The resolution is the spatial resolution, and the scale is related to the transformation, and hence to the bases' region of support. The registration is initialised on a low resolution and then iteratively progresses to finer resolution and smaller scales. The final deformation is given by:

$$v(x) = v_1(x) + v_2(x) + \dots + v_M(x), \quad (5.31)$$

where M is the total number of levels.

The next stage is to develop the local measure of misregistration. When the algorithm moves from one level to another it places basis functions on a regular grid

and models the deformation field:

$$d(x) = x + \sum_{k=1}^{m-1} v_k(x) + \sum_{i=1}^N c_i \Phi\left(\frac{x}{s} - k_i\right), \quad (5.32)$$

with $k_i = \frac{x_i}{s}$ where x_i is the position of the basis function.

Then the gradient \mathbf{G} of the cost function of Normalised Mutual Information, $NMI(A(x), B(d(x)))$, is calculated and used to determine which regions in the images $A(x)$ and $B(x + \sum_{k=1}^{m-1} v_k(x))$ are most likely to be misregistered at the current level. \mathbf{G} is evaluated and then the individual components are ordered according to their magnitude $\|G_i\|$. The centres of the regions of misregistration x_i are chosen as the locations of the basis functions for which G_i is above the selected threshold.

Local optimisation of the area surrounding each of these locations is performed to build the local deformation. The methods of gradient descent and quadratic interpolation are applied to the area of basis functions under the following cost function:

$$F(A(x), B(x'), x'), \quad (5.33)$$

in which $x' = x + \sum_{k=1}^{m-1} v_k(x) + v_m(x)$ and $v_m(x) = \sum_{j=1}^J \sum_{i=1}^8 c_i^j \Phi(x - x_i^j)$. J is the number of regions of mismatch. The optimisation is conducted one region at a time. Constraints are also included to ensure that there is no folding and the deformation is topologically correct. More details can be found in [74].

Algorithm: Fluid Registration

Crum *et al.* [27] implement Normalised Mutual Information (NMI) as the driving method for a fluid registration framework. They define derivatives for MI, NMI and Joint Entropy (JE) by considering point perturbations of the displacement vector field. These give the gradient vector at each voxel which indicates the direction which will give the greatest increase in similarity measure. These gradients are then used in fluid registration, [20]. Fluid registration allows large magnitude deformations in the vector field transformation, whilst still constraining it to be smooth. Smoothness is guaranteed by constraining the transformation based on a Stokesian limit of fluid-dynamical Navier-Stokes equations. The optimisation problem is formulated over

a space-time continuum that takes into account the path of the deformation. As such the relaxation of stress with time means that large magnitude deformations do not necessarily induce large penalties. Instead large penalties are the result of rapid distortions of the underlying continuum, and an infinite penalty is the result of any discontinuities in the deformation.

Commercial Software

A method using block matching is used in commercial software such as RevealMVSTM produced by Siemens Molecular Imaging, formerly Mirada Solutions, Ltd. Therefore, it is worth comparing the effectiveness of this software to perform non-rigid registration.

Open Source Methods

The Insight Registration and Segmentation ToolKit (ITK) is an open source software system designed to integrate different registration methods for different applications. It enables registration algorithms to be built easily using already tested routines, and as such enable the software system to be extended. It includes many different algorithms, including finite element methods. It also includes routines to calculate the similarity measures such as Mutual Information and Normalised Mutual Information.

Other Methods

The methods described above are a judicious selection of the leading non-rigid registration algorithms available. Many more algorithms have been proposed, for example those that deal with specific registration issues such as partial data [67], and so on. However those algorithms outlined do provide a good and varied cross-section of current state-of-the-art methods.

5.3.2 Applying Non-rigid Registration to Colorectal Cancer Images

Experiments were run on the colorectal datasets to give an overview of the effectiveness of the currently available non-rigid registration algorithms on this data:

1. Rueckert method, with B-splines and MI.
2. Park-Meyer method, adaptive registration.
3. Fluid method, using NMI.
4. RevealMVSTM.
5. ITK.

The results for each method will be shown visually, and also by comparing the error in node positions (in mm) between registered pre- and post-treatment. The original pre- and post- data was shown previously in Fig. 5.1. They are viewed in two dimensions, although they are three dimensional datasets and the actual registration is performed in three dimensions.

The results of a typical application of Rueckert's algorithm can be seen in Fig. 5.6 which shows the warped image computed by the algorithm. It seems that the key issue is that the algorithm analyses entropy at a single scale and equates signal complexity with 'interesting features' that drive the registration. However, the area that is of most clinical interest (the mesorectum and colorectum) are the image regions that are the least varying, most bland. The algorithm appears to find, at a coarse scale, points that are locally complex (high entropy) but which are of no clinical interest. These erroneous matches are then largely preserved as the grid of control points is refined. Hence, features of interest are misaligned by typically 30mm, and in four out of ten cases, the corresponding nodes are impossible to identify.

Park and Meyer suggested using a mis-match measure [66] in order to determine those areas of the image which (in the current iteration) seem to be maximally misaligned. The mismatch measure for two images in the data set was calculated and




Figure 5.6: Result of Rueckert's algorithm.

can be seen in Fig. 5.7. It is high across the whole image with an average value of 0.85 (on a scale of 0-1). As described above, control points are added in the maximally misaligned areas, and then registration is performed. However, in our datasets, it was seen that the control points were placed in areas surrounding, but never inside, the mesorectum. Again, the algorithm becomes fixated on the areas of the image that are of least clinical interest. A typical set of such control points can be seen in Fig. 5.8. However, we know from prior knowledge of the effects of therapy that, after affine registration has been performed, the areas of maximum misalignment are entirely inside the mesorectal region. The control points have been placed where the entropy is greatest, but where the interest is least. The misalignment of features of interest within the mesorectum is therefore the same as that in affine registration (typically 21mm). Since the colorectal region does not have the highest level of mismatch according to this measure, the nodes will not align any better than that of the affine registration that is performed first on the images. The algorithm can be improved by running over very many iterations, but this then loses the advantage of having a fast adaptive method, as control points are eventually placed across the entire image. A




Figure 5.7: Local mismatch measure.



Figure 5.8: Areas of maximum misregistration.

typical node misalignment of 9mm can be achieved this way, but there are three out of ten failures. Hence, the adaptive registration method proposed by Park and Meyer is very useful for looking at quite small localised changes in images, but is unsuitable for the large changes caused by radiotherapy in the colorectal region.

Crum *et al.* [27] proposed a fluid non-rigid registration algorithm. The results were analysed by looking at the difference image between the target image and source, and then the target image and the warped source image. In difference images the misregistered areas are identified by bright and dark areas, so that a perfectly aligned image would everywhere have the same gray-scale value. The original difference image between the target image and the affinely registered source image can be seen in Fig. 5.9. The difference image in Fig. 5.10 shows the difference when the source image is warped using Crum's method and an intensity-based similarity measure. It can be seen that this is better than the simple affine registration. Finally Crum's algorithm was run using Normalised Mutual Information as the similarity measure, and the difference image is shown in Fig. 5.11. This gives an improvement in the registration over just using the intensity-based similarity measure. However, although the peripheral areas of the dataset are well aligned, the central mesorectal area is still mis-registered, with features of interest misaligned by typically 7mm. This behaviour is similar to the effect of Park and Meyer's method which first attempts to register the area of greater entropy. Registration failures occurred on two out of the ten datasets. That is the registration was not improved at all from the affine initialisation, and in two cases the result was worse.

When the RevealMVSTM software is used to perform a nonrigid registration similar to that of block matching, the result in Fig. 5.12 is obtained. The algorithm fails in a similar way to that of Rueckert as it picks out false matches for areas of high entropy, causing an overall incorrect registration. Although the effect is not warped to the same extent, the failure mechanism is essentially the same. The difference image can also be seen in Fig. 5.13 which shows the extent of the mismatch. The large differences in the source and target image after it has been aligned means that the nodes do not align (this can be seen on the difference image). The misalignment distance of the nodes can be typically 7mm, with complete failures to register on three out of ten images.




Figure 5.9: Affinely Registered Intensity Difference Image.




Figure 5.10: Fluid Registered Intensity Difference Image.



Figure 5.11: Fluid Registered Intensity Difference Image Using Normalised Mutual Information.



Figure 5.12: Image registered using RevealMVS.




Figure 5.13: Intensity Difference Image Using RevealMVS.



Figure 5.14: Intensity difference image after local affine registration using ITK.

The method of local affine registration was also tested on one image of the dataset using the Insight Toolkit (ITK), but this similarly failed to register the data successfully, as can be seen in the difference image in Fig. 5.14, where the features of interest are mismatched by typically 10mm, and there are two out of ten failures.

Table 5.2: Node Misalignment Comparison.

Registration	Node Misalignment Distance (mm)	Failure rate (%)
Affine registration	21	0
B-splines	30	40
Adaptive registration	9	30
Fluid registration	7	20
RevealMVS TM	7	30
ITK	10	20

A summary of the node misalignments and robustness for all these methods can be seen in Table 5.2.

5.3.3 Summary

The first method by Rueckert *et al.* was seen to be unsuccessful on the colorectal images. Analysis of node positions was not required to validate this as it can be clearly seen that the image volume is mis-registered. Due to the large changes in this region from the treatment, Rueckert's method, despite using mutual information, fails to match up the correct areas in the volume. The changes are too large to be guessed by a simple coarse to fine grid scenario.

The Park-Meyer method fails to register the images, essentially for a similar reason. The areas it detects as those with the greatest mismatch are not the areas that we are interested in registering, and hence the wrong regions are matched up. This method is ideal for images where there are small localised changes, but in the colorectal images post-treatment the changes are large and span the whole image.

The Crum method worked well in registering the areas of the image volume that were not affected by the treatment. However the central mesorectum area was mis-registered due to the large changes caused by treatment and the absence of strong edges in this region. Running this method with both intensity-based similarity measures and NMI did indicate that NMI is a better measure to use for these images.

The non-rigid registration implemented in the RevealMVSTM software also fails to correctly register the images. Again, incorrect regions are 'locked' together and

hence the registration does not align the images. The difference image clearly shows the extent of mis-registration present even after the algorithm has been run. Likewise using ITK and the method of local affine registration also fails to register the images correctly.

Although the registered images acquired using the above methods look quite different visually, their failure mechanisms are very similar. They match regions together that are not equivalent due to the large changes in the images pre- and post-treatment, and as such are not robust.

5.4 Scale Saliency

None of the diverse non-rigid registration methods performed either well or reliably on our pre- and post-therapy data. This, despite considerable help from the authors of those systems. Is that because of certain special characteristics of our registration task, or is it symptomatic of a more generic problem? Is there some implicit assumption in all the methods that does not hold for the colorectal MR images, and possibly for other medical imaging tasks? We contend that there are two main reasons for the observed poor performance, and that they are quite generic.

The first concerns the saliency, or ‘feature-ness’, of the images. All of the registration algorithms align the images where there are sufficiently many ‘features’, at a certain scale. That is, prior to regularisation that propagates local matches, they preferentially align areas of high feature complexity, which is generally equated with high entropy, at the scale of interest. However, Kadir and Brady [47] have demonstrated that scale and saliency are implicitly related, both to each other and to image complexity. That is, features are associated with a particular scale within the image. They show this by considering the entropy of local image regions over a range of scales (though a similar analysis applies to other image representations such as local energy or local phase). They suggest that areas that have a consistently high entropy, that is, if they exist as features over a sufficiently wide range of scales, then they should not in fact be judged salient, since they can not be considered rare, which is the essence of

saliency. However, if there is a peak at a single scale (or spread over a small range of scales) in the entropies over a particular area then this implies that there is a feature at that position localised to that scale. Kadir and Brady demonstrate that saliency has to be analysed in terms of scale as well as space. However, all of the non-rigid registration algorithms assess image complexity at a single scale. Even those algorithms that implement a so-called multi-scale approach actually treat the entropy independently at each scale. None of them consider the notion of scale-saliency.

Does this matter in practice? To investigate this question, we applied a recent refinement of the algorithm proposed by Kadir and Brady [47] to the colorectal cancer images. We set the algorithm to find the 85 most salient regions across both scale and space (empirically this number produces good results and is insensitive to changes $\pm 5\%$). A typical result can be seen in Fig. 5.15, where the size of a circle indicates the scale of the feature that has been found. The first thing to note is that the algorithm detects mostly sensible features but that the associated circles are of substantially varying sizes, indicating that it is unlikely that all of the features sought can be detected at a single scale. It should also be noted that the more complex areas of the images, for example the muscle regions in the bottom left and bottom right hand corners of the image, do not yield any salient points. This is because the entropy in these areas is high across the entire range of scales, hence they are not considered salient.

Why has this not previously been raised as an issue? It seems that many, if not most, current algorithms are tested on images where the features of interest are all present at a particular scale. This is true in particular for brain images, and to an extent breast images. However, colorectal images exhibit structure (and complexity) over a wide range of spatial scales and so algorithms should not make this assumption. Of course, a scale-saliency analysis such as that of Kadir and Brady [47] could be incorporated into non-rigid registration algorithms, though the proper integration of registration of images across a stack of spatial scales is not entirely straightforward.

The second main reason for the observed poor performance of the registration




Figure 5.15: Illustration of saliency across scale and space.

methods is the fact that application specific knowledge is not used. We intend to mobilise prior knowledge of the effects of treatment along with an initial segmentation of the images, to find corresponding features, [7], [8]. This is discussed in detail in the next chapter.

Chapter 6

Knowledge Representation

We noted in the previous chapter that a lack of knowledge concerning the anatomy and physiology of the region that is of prime interest to be registered is one of the main reasons that the current ‘generic’ non-rigid registration algorithms fail, in the sense that they converge, far too often to be clinically useful, to a local, non-global minimum. Although scale saliency algorithms, such as that by Kadir *et al.* [47], are able to find features within the images, that correspond to regions of high entropy at specific scales, such features do not necessarily bear any resemblance to interesting anatomical features. Frantz *et al.* [38] also propose a method for finding ‘interesting’ features within images, such as points of high curvature, but again they are not necessarily anatomically meaningful. Also, in order to align the features found using such methods, we would have to go through a process of establishing correspondences between pairs of such features in the two image sets. To address this problem, in this chapter we describe how we search for interesting anatomical features that will always be present in the images, and that we want to align in the registration process since they are corresponding locations.

We start by reviewing a number of previous methods by which people have tried to incorporate shape knowledge into registration, including atlases and segmentations. We then outline the method developed in this thesis for finding the anatomical regions of interest of the specific patient, and finally review the performance of this algorithm and describe how it can be useful in aiding the registration process.

6.1 Incorporation of Knowledge

Lack of homology between the two datasets is one of the major factors affecting the performance of any registration algorithm. This is frequently noted in discussions of inter-person registration, for example, the different sulcal patterns of different individuals' brains, but it is equally true when treatment causes major changes to the anatomy. In our case, we know that the size, shape and appearance of the tumour and lymph nodes are certain to change post-treatment, but as described previously, there are also changes in texture, especially in the mesorectal region of the image.

A simulation experiment was conducted to re-emphasise the problems encountered in the previous registrations, as well as looking at how to enable possible solutions. Two simulation images were formed and these can be seen in Fig. 6.1. These images represent the colorectum and the mesorectum, and also include the high entropy regions outside the mesorectum that do not necessarily correspond pre- and post-treatment. The central circular section is similar to the colorectum, surrounded by the much less clearly defined mesorectum (a layer of fat). The smaller circles indicate nodes, and the line a blood vessel type structure. The 'node' below the colorectum in the images moves considerably between the images, so if this were a real image, it would be likely that they are in fact different features.

Rueckert's B-spline algorithm [75] was used throughout these simulations. Any of the algorithms we tested in Chapter 5 could have been used, and Rueckert's was chosen because the implementation was such that the warp field was straightforward to access and modify. The two images in Fig. 6.1 were aligned by directly applying the algorithm to them. The result can be seen in Fig. 6.2, along with the difference image. It can clearly be seen that the registration has been unsuccessful in that although the colorectum has been aligned, the other details have been incorrectly registered. Much of the noisy area surrounding the image has been registered, even though there is, in fact, no anatomical correspondence between these two areas.

A second simulation was performed, using the same two original images but without the noisy surroundings. The original images can be seen in Fig. 6.3. The resulting




Figure 6.1: Simulated images of before and after treatment, that need to be registered.




Figure 6.2: Registered image using B-splines and Mutual Information, showing registered source image (left) and difference image (right).




Figure 6.3: Simulated images of before and after treatment without the ‘noisy’ surrounds.




Figure 6.4: Registered central regions of the images using B-splines and Mutual Information, showing registered source image (left) and difference image (right).

registration, again using Rueckert’s B-spline algorithm, can be seen in Fig. 6.4. It gives a much better registration result. The only area that is not aligned correctly is the bottom node, but, as noted earlier, this would be classed as a change in homology.

Finally, the transformation found above for the correct registration was used on the noisy image, in order to provide a registration of just the area we are interested




Figure 6.5: Registration of whole images using just central information, showing registered source image (left) and difference image (right).

in. The resulting images can be seen in Fig. 6.5. The results show the central section is aligned whereas the noisy outer area is not at all, but since this outer area is not homologous, this is the required result.

A similar example can be shown with the real colorectal MR data, where some of the important shape features were segmented by hand. These features include the bone structures, colorectum, mesorectum, and edge of the body, and can be seen in Fig. 6.6.

The registration was performed on the two ‘mask’ shape images and the result and difference image can be seen in Fig. 6.7. The same transformation was then applied to the original image and the results of this can be seen in Fig. 6.8. Although much of the image is not correctly registered, the colorectal and mesorectal areas are well aligned, and by focussing on this area some changes can be observed. The disappearance of the lymph node and change in intensity of the tumour within the colorectum can be clearly seen in Fig. 6.9.

From these experiments, it is clear that providing shape information can drastically alter the robustness and effectiveness of the registration. However, the key question then becomes how to segment the key areas of interest within the image.




Figure 6.6: Real images before and after treatment, and their segmentations.

We have already seen that it is possible to find the colorectum and mesorectum, with a small amount of user interaction. Ideally, however, we would like to perform the registration with no user interaction at all. Many shape based methods are available to detect individual features, for example, Tsai *et al.* [87] introduced a shape based method using level sets. However, we are interested in finding a series of different anatomical features within the images.

Many authors have suggested using atlases as a way to segment images, since the atlas represents what one may expect where, within the image. Many examples in brain imaging use atlases that attempt to account for the presence of tumours. For




Figure 6.7: Registration of segmented images, showing registered source image (left) and difference image (right).




Figure 6.8: Registration of real data using same transform as that calculated for segmented images, showing registered source image (left) and difference image (right).




Figure 6.9: Mesorectal area showing disappearance of lymph node and change in intensity of tumour in the colorectum.

example, Kyriacou and Davatzikos [54] propose a bio-mechanical model of deformation. Dawant *et al.* [31] and Cuadra *et al.* [28] try to solve the same problem using optical flows. Using atlases usually involves fitting the test data to the given atlas, and assumptions that the tumour grows in a predictable way given its seed point. Although this is sensible for some applications, in our application this would entail a loss of information. We are only interested in registering images of the same patient, and, as such, do not want all our data to match to a ‘generic’ case. Also, the types, sizes and evolution of colorectal tumours varies vastly between patients.

Tsai *et al.* [86] proposed a method using coupled shape models, involving deformable active contours and signed distance functions in order to segment regions in the lower gastro-intestinal area of the body. Their method using trained shape models worked well for segmenting anatomies of healthy patients. Similarly, Yang *et al.* [92] proposed a method of neighbour constrained segmentation using a Maximum A Posteriori (MAP) framework. Again, their method relies on training data from healthy patients, such that the shapes are not in fact varying that much. However, in our case we have much greater variation in anatomies due the large and varying tumours within the colorectal region.

Christensen *et al.* [19] realised this problem and also noted that shape knowledge is crucial in registering such mobile areas of the body. They manually segmented the images and built three dimensional models which they then registered successfully. The main problem with this method is that it is time and labour intensive and requires a skilled clinician to segment the images accurately.

Instead, our method builds on that of Felzenszwalb and Huttenlocher [36]. Their work aims to find pictorial structures within images based on the appearance of the different aspects of the image, and the distances between these parts of the image. They represent objects as connected parts arranged in a deformable configuration. They use an appearance model for each individual part, as well as an edge deformation for linking the parts together. The algorithm then finds an optimum match for the objects by computing the global minimum energy function, calculated from the distance functions and the appearance models.

Felzenszwalb and Huttenlocher define their model as an undirected graph $G = (V, E)$ where $V = (v_1, \dots, v_n)$ corresponds to the vertices (that is, n = number of ‘features’), and then there is an edge between each vertex $(v_i, v_j) \in E$. The model is matched to the image using appearance models where the mismatch between the image and the model is given by $m_i(l_i)$ where l_i is the location of the vertex, v_i . A function to measure the degree of deformation is also given as $d_{ij}(l_i, l_j)$ for each pair of vertices, or edge. The best global match of the model to the image can then be given as:

$$L^* = \arg \min_L \left(\sum_{i=1}^n m_i(l_i) + \sum_{(v_i, v_j) \in E} d_{ij}(l_i, l_j) \right). \quad (6.1)$$

They show their method works successfully and robustly on face and human body detection within larger images.

6.2 Relational Representation of Colorectal Cancer MR Images

Certain features within the colorectal images are ‘connected’ in a well defined way, that only varies a small amount from patient to patient and is not affected to any large extent by tumours within the colorectum. Hence, we can use this information in order to find their positions accurately, reliably and quickly by building a model of the basic anatomical shapes and relative positions.

The method outlined above by Felzenszwalb and Huttenlocher [36] assumes that pairs of features are equally dependant on each other; that is, they model the structure as an undirected graph. This is not so in our case: we require a directed graph model. In fact, the only pair of anatomical features that are equally dependent on each other are the hips. That is, we assume that they are a certain distance apart, and that they are at similar superior and posterior locations, and that they have similar radii. We search for instances of features in the order, according to the directed graph, shown in Fig. 6.10. A similar hierarchical model idea was proposed by Marai *et al.* [59], to aid registration of bones in hands.

Following Felzenszwalb and Huttenlocher, to search for a particular set of features, we solve Eqn. 6.1. However, in our implementation, we replace this single equation by a series of consecutive minimizations to be consistent with our directed graph model. These consecutive steps are outlined below for each part of the anatomy.

6.2.1 Hips

First, we search for the hips by minimising the following equation:

$$L^* = \arg \min_L ((m_{hip1}(l_{hip1})) + (m_{hip2}(l_{hip2})) + (d_{hip1hip2}(l_{hip1}, l_{hip2})) + d_{rad1rad2}(l_{hip1}, l_{hip2})), \quad (6.2)$$

where l_{hip1} and l_{hip2} are defined by their coordinate positions x, y, z and their radii, r .

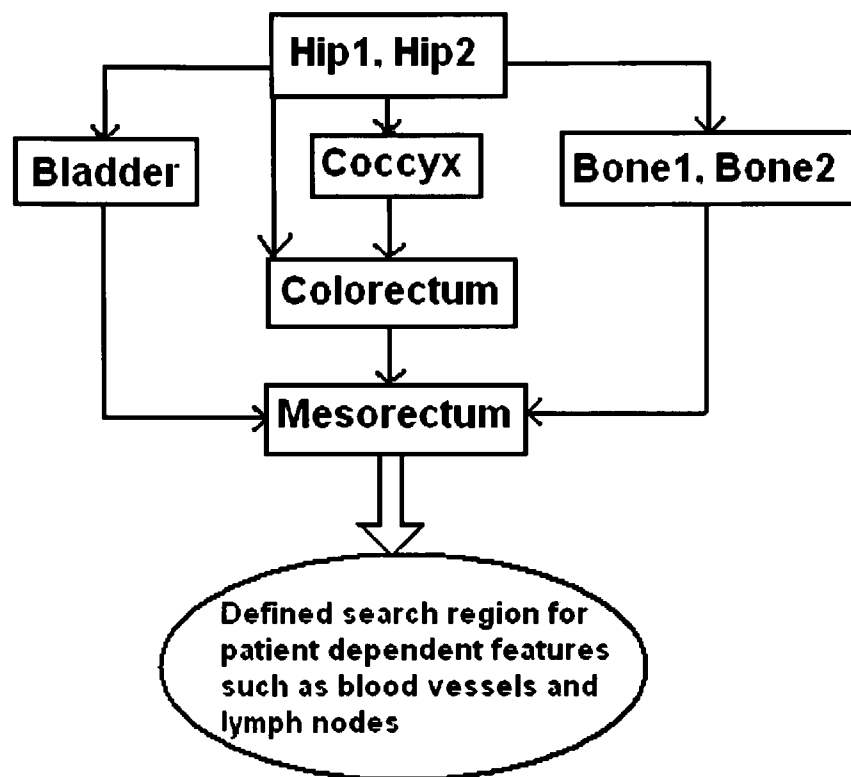


Figure 6.10: Directed Graph Model for Colorectal Anatomy.

The appearance model we use is that of a section of the edge of a sphere which we fit to the image data. Often, only half of the hip is visible in the image, so we model it as a semi-circular arc that will fit to the edges in the segmented images. For each appearance model, we develop a mismatch measure of how well any particular point fits the model. We then use our prior knowledge of where the hips are located in the image to define search areas in order to test what the mismatch measure is at every possible point within these areas. The search regions for each vertex, or feature, are shown in Fig. 6.11.

For each hip, we test the appearance model using 4 parameters: the x position, the y position, the z position, and the radius r . Having found the mismatch measures corresponding to the appearance models for the two hip regions, the edge deformations and energy functions for each combination of the eight parameters, $(x_{hip1}, y_{hip1}, z_{hip1}, r_{hip1}, x_{hip2}, y_{hip2}, z_{hip2}, r_{hip2})$, are found. The distance metrics are based on constraints that we know to be true: we use the fact that the hips are at the same level and only a small rotation is permitted, $d_{hip1hip2}$; also, the difference




Figure 6.11: Search regions for key features of interest within colorectal cancer MR images.

between the radii of the two hips must be small, $d_{rad1rad2}$. MRI radiographers plan the imaging such the hips are level on the images as much as possible, so we expect their radii to be similar with only a small difference between them. These constraints are modelled by a Gaussian cost function incorporated into $d_{hip1hip2}$ and $d_{rad1rad2}$.

Due to the number of iterations required to find the deformations and energies for every combination of the parameters, this is a very computationally expensive process. Hence, a multi-scale approach has been implemented in order to speed up the Felzenszwalb and Huttenlocher algorithm further. The mismatch measure is calculated for the pixels in the full size image and then downsized up to 4 times. During the downsizing the minimum mismatch measure is taken for each pixel as opposed to calculating the mean over a region. Often the point with the best match to the shape model is surrounded by points that have a poor match, due to the nature of the model.

The distance deformations are then calculated on this low resolution set, along with the corresponding energies. The minimum energy is found, and hence the points of best match are found at this low resolution. However, we are interested in accurate




Figure 6.12: Illustration of multiscale approach.

positions of the centre points of the hips, so we scale up the image and look for best match again but only in the region of lowest energy from the previous scale, and so on. This can be seen in Fig. 6.12 which shows how the ‘best match’ combination of points is found using the multi-resolution approach. The red square is equivalent to the ‘best match’ series of points.

6.2.2 Coccyx

Having found the optimal hip positions, we can repeat the minimisation with a new set of parameters in order to find the coccyx:

$$\begin{aligned}
 L^* = \arg \min_L & \left((m_{cx}(l_{cx})) + \right. \\
 & + (d_{hip1cx}(l_{hip1}^*, l_{cx})) + (d_{hip2cx}(l_{hip2}^*, l_{cx})) + \\
 & \left. + d_{hip1cx-hip2cx}(l_{hip1}^*, l_{hip2}^*, l_{cx}) \right),
 \end{aligned} \tag{6.3}$$

where the distance metric $d_{hip1cx-hip2cx}$ forces the coccyx to be equidistant from each hip.

The appearance model for the coccyx is also straightforward to define: We know that the coccyx is a small area of high intensity surrounded by an edge a few pixels

wide of low intensity, so this forms the basis of our model. We use the segmentation from the antigeometric diffusion to search for the coccyx region, and hence it is just based on one parameter, the region number. Since the number of regions that we are searching through to find the coccyx is small, we no longer need the multi-resolution approach for this section.

6.2.3 Bladder

The hips also constrain the search region for the bladder, and the bladder can be found in a similar way to the coccyx:

$$\begin{aligned} L^* = \arg \min_L & ((m_{bl}(l_{bl})) + \\ & + (d_{hip1bl}(l_{hip1}^*, l_{bl})) + (d_{hip2bl}(l_{hip2}^*, l_{bl})) + \\ & + d_{hip1bl-hip2bl}(l_{hip1}^*, l_{hip2}^*, l_{bl})). \end{aligned} \quad (6.4)$$

The antigeometric diffusion segmentation is again used to find the region corresponding to the bladder. The appearance model is based on the fact that the bladder is a large, high intensity region. Again, $d_{hip1bl-hip2bl}$ ensures that the centroid of the bladder is equidistant from the two hips.

6.2.4 Bone structures

The pelvic bone structures encasing the hips can also be found by knowing the hip positions:

$$\begin{aligned} L^* = \arg \min_L & ((m_{b1}(l_{b1})) + \\ & + (d_{hip1b1}(l_{hip1}^*, l_{b1})) + (d_{hip2b1}(l_{hip2}^*, l_{b1}))), \end{aligned} \quad (6.5)$$

$$\begin{aligned} L^* = \arg \min_L & ((m_{b2}(l_{b2})) + \\ & + (d_{hip1b2}(l_{hip1}^*, l_{b2})) + (d_{hip2b2}(l_{hip2}^*, l_{b2}))). \end{aligned} \quad (6.6)$$

Segmented regions are again found, and the appearance model is characterised by the large, relatively high intensity regions. The appearance model is also constrained

by the fact that the bone is present only in one side of the image, and does not cross the line that runs between the midpoint of the hips and the coccyx. The distance metric is dependent on the hips, such that each bone region is close to one hip, and further away from the other.

6.2.5 Colorectum

The position of the colorectum is relative to both the hips and coccyx positions:

$$\begin{aligned}
 L^* = \arg \min_L & \left((m_{cr}(l_{cr})) + \right. \\
 & + (d_{hip1cr}(l_{hip1}^*, l_{cr})) + (d_{hip2cr}(l_{hip2}^*, l_{cr})) + \\
 & \left. + d_{cxcr}(l_{cx}^*, l_{cr}) \right). \tag{6.7}
 \end{aligned}$$

That is, the colorectum is at approximately the centroid of the hips and the coccyx, so this specifies the distance metric based on the hip and coccyx positions. Its appearance model is a low intensity, approximately circular region.

6.2.6 Mesorectum

Finally, from these anatomical locations the mesorectum position can be estimated as the inscribed circle of the triangle that connects the hips and the coccyx, as previously described in Chapter 3. This segmentation can therefore be used to find points within the mesorectum for the purposes of registration, or constrain lymph node searches.

6.3 Results

The anatomy-guided search algorithm was run on the same series of 20 colorectal MR datasets, of which 10 were pre-therapy and 10 were post-therapy. The datasets spanned a wide range of anatomical variation, and, as such, constitutes a fair test of the search algorithm. The search was implemented using an adaptation of the method proposed by Felzenszwalb and Huttenlocher [36], to use directed graphs and a multiscale search, in order to speed up the search. We also implemented a multi-

scale search for the appearance matches, again to speed up the process. The total runtime was approximately 6 minutes, running in Matlab.

The search algorithm reliably finds the centres of the hips in every case to within 2 pixels. The actual positions were found using a measurement tool on an image viewer. If the centres were to be drawn on by hand, without this measuring tool, the error was seen to be up to 4mm. The algorithm also reliably finds regions that correspond to the coccyx, pelvic bones, bladder and colorectum, notwithstanding the large variability of some of these features. Results can be seen in Fig. 6.13 where the segmented image and selected regions can be seen alongside the original two dimensional slices. It can be seen that accurate regions have been found in all cases.

Note that in this first implementation the centres of the hips and coccyx have been found from the whole three dimensional dataset, and the regions corresponding to the other bones, the bladder and the colorectum are found individually for each slice. This could, however, be extended straightforwardly to a three dimensional search, if required.

This anatomy atlas is a first stage for registration, since it constrains a structure in the pre-therapy image to map to its correspondent in the post-therapy image. In particular, the bladder is present in both the pre- and post-therapy images even if it dramatically changes shape. As such, an initial anatomical model fit to the data should provide a good initial constraint for a non-rigid registration algorithm. However, we do not want this segmentation to be the only driver for the registration, as errors in the segmentation will follow through into the registration, providing an incorrect result. As such we use the anatomy model to find point correspondences on images, such as the extreme points of the pelvic bones, bladder, and colorectum. The registration methods and results using this shape knowledge are presented in the next chapter.

6.4 Summary

We have developed a method for automatically segmenting useful anatomical features of interest, by implementing a directed graph model for the process of object recognition developed by Felzenszwalb and Huttenlocher. We include the knowledge that not all features are equally dependent on each other, as well as knowledge of expected appearance and positions within the dataset. This enables an initialisation for the non-rigid registration of colorectal cancer images pre- and post-therapy and could also be used in other applications that require the registration of very large scale deformations.




Figure 6.13: The left hand column shows the original two dimensional slices and the right hand column show the corresponding anatomies.

Chapter 7

Mobilising Shape Knowledge

We defined an anatomical representation for colorectal cancer MR images in the previous chapter, and we now discuss how we can best use that representation to improve the results of non-rigid registration. The key anatomical feature locations can be incorporated into generic non-rigid registration algorithms. This is done in different ways, so as to be tailored to the requirements of the particular algorithm.

This chapter discusses exactly how this shape knowledge is incorporated into a number of the algorithms outlined previously, and the results of the subsequent registration. The algorithms are all run on the same ten datasets that were used previously for testing the non-rigid registration algorithms. Although the datasets do have large scale deformations between the pre- and post-treatment images, they are not atypical. That is, they represent a reasonable test of the likely performance of such algorithms on colorectal cancer MRI.

7.1 Shape Representation

The non-rigid registration algorithms we tested in Chapter 5 vary greatly in their implementation and level of user interaction. For example, RevealMVSTM is a highly interactive tool, and as such, the knowledge can be added by plotting the landmark features that we have found. The implementation of Crum's algorithm on the other hand is not at all interactive and runs by simply inputting two images to register. In fact, many 'generic' algorithms are designed to just include a source and a target



Figure 7.1: Key landmark points on the relational representation of the colorectal anatomy.

image, although in some cases they can incorporate an initial warp field estimate.

7.1.1 Finding Landmark Points

In order to incorporate the anatomical representation from Chapter 6 into a useful form for registration, we find a series of landmark points from the relational representation. We do this by finding the extreme, key points on the representation of the colorectum, bladder, and bones. This is done with a simple threshold, and the landmark points can be seen in Fig. 7.1. We take the centres of the hips and coccyx as three landmarks. We also take the extreme points on the pelvic bones, that is the top central point and the bottom central point. The centre point on the bottom of the bladder is found along with the two 'corner' points on the bladder. Finally four points are taken on the edge of the colorectum, at the top, bottom and two sides.

7.1.2 Thin Plate Splines

In order to find the initial warp field we can apply the Thin Plate Splines algorithm to the series of landmark points [11]. Thin Plate Splines are diffeomorphic deformations,

that do not fold. As such they are useful for defining deformations within biological structures that have a minimum bending energy. These transforms were mentioned briefly in Chapter 5, but we will now describe how they are used in more detail.

A set of points on an initial image are defined as $P_1 = (x_1, y_1)$, $P_2 = (x_2, y_2)$, ..., $P_n = (x_n, y_n)$. The distances between any two of these points is $r_{ij} = |P_i - P_j|$.

The matrices K , P and L are defined, based on the basis function:

$$z(x, y) = -U(r) = - \begin{cases} r^2 \log r^2 & \text{in 2D} \\ |r| & \text{in 3D} \end{cases} \quad (7.1)$$

$$K = \begin{pmatrix} 0 & U(r_{12}) & \dots & U(r_{1n}) \\ U(r_{21}) & 0 & \dots & U(r_{2n}) \\ \dots & \dots & \dots & \dots \\ U(r_{n1}) & U(r_{n2}) & \dots & 0 \end{pmatrix}, n \times n. \quad (7.2)$$

$$P = \begin{pmatrix} 1 & x_1 & y_1 \\ 1 & x_2 & y_2 \\ \dots & \dots & \dots \\ 1 & x_n & y_n \end{pmatrix}, 3 \times n. \quad (7.3)$$

$$L = \begin{pmatrix} K & P \\ P^T & 0 \end{pmatrix}, (n+3) \times (n+3). \quad (7.4)$$

Let $V = (x'_1, x'_2, \dots, x'_n)$ be the x -coordinates of the homologous (corresponding) points of P_1, P_2, \dots, P_n , and $Y = (V \mid 0 \ 0 \ 0)^T$. W is defined as $W = (w_1, w_2, \dots, w_n)$ and the coefficients a_1 , a_x , and a_y are also defined, such that:

$$L^{-1}Y = (W \mid a_1 \ a_x \ a_y)^T. \quad (7.5)$$

The elements of $L^{-1}Y$ are then used to give a function $f_x(x, y)$ everywhere in the plane:

$$f_x(x, y) = a_1 + a_x x + a_y y + \sum_{i=1}^n w_i U(|P_i - (x, y)|). \quad (7.6)$$

This is repeated for $V = (y'_1, y'_2, \dots, y'_n)$, the y -coordinates of the homologous points, to give $f_y(x, y)$. Then, $f(x, y) = [f_x(x, y), f_y(x, y)]$ maps the points (x_i, y_i) to its homologue (x'_i, y'_i) , and is the Thin Plate Spline mapping, or deformation.

The landmark points were described in the previous chapter, and are the extreme points on the bones, points on the colorectum, points on the bladder, and points within the mesorectum. These landmarks are therefore at key anatomical locations on the deformable areas. These points also include knowledge of how the anatomy will deform. For example, we can place four extreme points on the colorectum, assuming that it will not rotate. We know from prior medical knowledge that the main deformations are due to either shrinkage of the region due to the therapy, or expansion due to the tumour growing.

The Thin Plate Splines algorithm outputs a warp field that can be directly input into another algorithm. In the cases where the algorithm just requires two images to be entered, we input an image registered using this initial warp field. We are still interested in the ‘generic’ algorithms and do not at this stage want to alter the structure of the algorithms.

7.2 B-splines and Mutual Information

The B-spline and mutual information registration written by Rueckert [75] was run by incorporating the shape information as an initial registration. This initialisation therefore incorporated the large scale aspects of the deformation allowing the algorithm to then proceed normally by registering the small scale deformations that remained.

7.2.1 Method

A rigid registration was first performed on the image using the rigidly moving anatomical landmarks (hips and coccyx), and then a non-rigid registration was performed using the other non-rigidly moving anatomical landmarks. Thin Plate Splines and tri-linear interpolation were used to register the image. This ‘initialised’ image was then input into Rueckert’s algorithm along with the source image and the registration was run as it had been previously.

Table 7.1: Node Misalignment using B-Splines, Mutual Information and Shape Information.

Pa- tient	Affine misalignment (mm)	Without Shape Knowledge (mm)	With Shape Knowledge (mm)
1	26.85 ± 7.00	35.67 ± 3.66	5.96 ± 4.02
2	21.65 ± 9.00	35.64 ± 8.22	5.96 ± 4.02
3	15.00 ± 3.17	28.00 ± 8.54	16.46 ± 6.64
4	16.31 ± 4.41	19.33 ± 6.03	3.80 ± 1.26
5	17.50 ± 4.92	FAIL	5.17 ± 4.45
6	17.66 ± 4.03	39.03 ± 3.55	8.39 ± 7.44
7	36.21 ± 6.14	FAIL	5.41 ± 2.45
8	30.00 ± 3.65	FAIL	3.29 ± 0.97
9	17.34 ± 4.55	FAIL	15.54 ± 0.76
10	15.63 ± 6.00	22.33 ± 4.51	8.54 ± 2.37

7.2.2 Results

This was run on ten different patient datasets, registering pre- and post-chemotherapy. The node misalignments were measured to give an indication of the accuracy of the algorithm. They are shown in Table 7.1.

Overall, the mean misalignment was 8.31mm with a standard deviation of 6.25mm. Two typical results for patients 1 and 2 are shown in Figs. 7.2 and 7.3. It shows the target and source images that are input into the B-spline algorithm, and then the aligned source images. The improvement is qualitatively visible for Patient 1, but much less so for Patient 2. The node misalignment results also indicate that this method is not that accurate as large node misalignments can still result. However, it is more robust than the algorithm used without shape knowledge, as it does not fall into local minima at the larger scales.

Although the results indicate a large improvement from the results of this algorithm with no shape knowledge (up to 25mm), the nodes are still not that well aligned. The outer edges of the image are also poorly aligned still, and actually get worse after the B-spline algorithm has been applied since the problem of local minima remains at the finer scales of registration. This can be seen especially in Fig. 7.3. However, the robustness is much improved in that there are no longer any complete registration




Figure 7.2: Results using B-splines, Mutual Information and shape information for Patient 1: Top left – Target image, Top right – Initially aligned source image, Bottom left – Final warped source image




Figure 7.3: Results using B-splines, Mutual Information and shape information for Patient 2: Top left – Target image, Top right – Initially aligned source image, Bottom left – Final warped source image

Table 7.2: Node Misalignment using Adaptive registration and Shape Information.

Pa- tient	Affine misalignment (mm)	Without Shape Knowledge (mm)	With Shape Knowledge (mm)
1	26.85 ± 7.00	8.76 ± 3.93	3.22 ± 0.48
2	21.65 ± 9.00	8.77 ± 2.87	5.76 ± 1.68
3	15.00 ± 3.17	12.74 ± 4.25	3.88 ± 0.34
4	16.31 ± 4.41	8.33 ± 2.77	7.08 ± 2.45
5	17.50 ± 4.92	FAIL	3.58 ± 0.80
6	17.66 ± 4.03	8.76 ± 4.80	7.01 ± 3.85
7	36.21 ± 6.14	5.06 ± 2.13	4.05 ± 1.53
8	30.00 ± 3.65	FAIL	5.36 ± 1.94
9	17.34 ± 4.55	FAIL	6.36 ± 0.85
10	15.63 ± 6.00	7.59 ± 3.84	3.77 ± 0.50

failures where corresponding lymph nodes can no longer be identified.

7.3 Adaptive Registration

Park and Meyer [66] propose registration by finding the maximally misaligned regions within the images and then placing control points in these regions. The images are then aligned using these control points.

7.3.1 Method

The algorithm was initialised with the shape-aligned images in a similar way to the B-splines algorithm, and then control points were placed at the areas where maximum misalignment still occurred. These areas were then aligned using the mutual information and a simple gradient descent optimiser.

7.3.2 Results

The algorithm was run on the same ten patient datasets. The measured node misalignments are shown in the Table 7.2.

Overall, the mean misalignment was 5.01mm with a standard deviation of 2.06mm. An example result is shown in Fig. 7.4 and it demonstrates that control points are still mostly placed around the edges of the images, at the points of highest entropy.




Figure 7.4: Results using adaptive registration and shape information for Patient 7: Top left – Target image with points of maximal misalignment, Top right – Source image showing changes in points of maximal misalignment, Bottom left – Source image showing aligned points, Bottom right – Warped source image.

In fact, 92% of control points in the ten images tested are still outside the mesorectal fascia, which is our region of interest. As such, the algorithm does little to improve the images from the initial registration unless a large number of control points are used. However, the shape alignment does provide added robustness, and although not always accurate, the method ran on all images giving sensible results much improved from just running the registration without any initialisation. That is, the method does not fail in the incidence of large deformations.

Table 7.3: Node Misalignment using Fluid Flow registration and Shape Information.

Pa- tient	Affine misalignment (mm)	Without Shape Knowledge (mm)	With Shape Knowledge (mm)
1	26.85 ± 7.00	9.60 ± 3.50	7.92 ± 2.50
2	21.65 ± 9.00	7.72 ± 3.05	5.10 ± 0.94
3	15.00 ± 3.17	10.00 ± 8.95	3.93 ± 1.74
4	16.31 ± 4.41	2.10 ± 1.52	3.39 ± 1.41
5	17.50 ± 4.92	8.31 ± 3.50	8.02 ± 1.62
6	17.66 ± 4.03	3.87 ± 4.25	8.05 ± 2.43
7	36.21 ± 6.14	FAIL	7.37 ± 2.76
8	30.00 ± 3.65	15.11 ± 4.50	9.11 ± 0.98
9	17.34 ± 4.55	FAIL	3.87 ± 1.85
10	15.63 ± 6.00	2.19 ± 0.85	2.63 ± 0.62

7.4 Fluid Registration

The algorithm by Crum *et al.* [27] uses a fluid flow method for performing the registration using mutual information, or normalised mutual information.

7.4.1 Method

The fluid flow algorithm was applied in the same way as that of the B-splines, using normalised mutual information. The data was initially aligned using the anatomical control points and then these aligned datasets were input to the algorithm.

7.4.2 Results

The algorithm was run on the same ten patient datasets. The measured node misalignments are shown in Table 7.3.

Overall the mean misalignment was 5.94mm with a standard deviation of 2.76mm. Two typical results for patients 4 and 7 are shown in Figs. 7.5 and 7.6. They show the target and aligned source images. The results are visibly good. Although the nodes are not perfectly aligned, the algorithm does not fall into local minima and a sensible result is obtained for each dataset.




Figure 7.5: Results using fluid registration and shape information for Patient 4: Left – Target image, Right – Warped source image.




Figure 7.6: Results using fluid registration and shape information for Patient 7: Left – Target image, Right – Warped source image.

7.5 RevealMVS™

RevealMVS™ is an interactive registration tool, so incorporating the shape points is easily done, by manually plotting the points that were found by the anatomical feature detector. A landmark registration (both rigid and non-rigid) can then be run, followed by a further registration based on this ‘constrained’ starting point.

7.5.1 Method

The specific sequence used was:

1. Landmark Rigid - performs a rigid registration based on the rigid landmarks that are the bone structures.
2. Automatic Rigid - performs a further rigid registration using the in-built algorithm, that will refine the result on the landmark based registration
3. Landmark Deformable - performs a non-rigid registration based on all the anatomical landmarks that are known.
4. Automatic Deformable - performs a further non-rigid registration using the previous result as the starting point, which thus constrains the registration, and as such allows the overall result to include large scale deformations.

7.5.2 Results

This was run on the same ten pairs of datasets. The measured node misalignments are shown in the Table 7.4. Overall, the mean misalignment was 4.59mm with a standard deviation of 1.60mm.

A typical result is shown in Fig. 7.7, for patient 2. It shows the target and source image and then the aligned source image. It can be seen that a good registration is achieved, especially for the central region.

However, the algorithm is still trying to align the outer regions of the image, which are often mis-registered to give very large incorrect deformations, despite the




Figure 7.7: Results using RevealMVS™ and shape information for patient 2: Top left – Target image, Top right – Source image, Bottom left – Warped Source image with overlaid warp field.

Table 7.4: Node Misalignment using RevealMVSTM and Shape Information.

Pa- tient	Affine misalignment (mm)	Without Shape Knowledge (mm)	With Shape Knowledge (mm)
1	26.85 ± 7.00	3.34 ± 1.17	4.50 ± 1.73
2	21.65 ± 9.00	6.84 ± 0.54	3.75 ± 2.06
3	15.00 ± 3.17	7.15 ± 2.24	3.33 ± 1.15
4	16.31 ± 4.41	7.63 ± 1.40	2.17 ± 0.29
5	17.50 ± 4.92	FAIL	6.17 ± 0.29
6	17.66 ± 4.03	8.33 ± 0.76	6.66 ± 1.15
7	36.21 ± 6.14	FAIL	5.25 ± 0.50
8	30.00 ± 3.65	6.67 ± 2.08	4.50 ± 0.50
9	17.34 ± 4.55	FAIL	4.33 ± 0.76
10	15.63 ± 6.00	5.67 ± 2.08	5.33 ± 1.15

fact that the central region is well aligned. This caused large errors in the overall registration in three out of the ten datasets. A typical image for patient 6 is shown in Fig. 7.8 showing the central region now registered much more accurately than the outer region.

We performed a quantitative analysis of the spatial distribution of the registration within the images. Fig. 7.9 shows a best fit curve of the distribution of errors measured from the centre of the region of interest (the centre of the mesorectum). It takes into account 8 datasets that were compared. Two of the datasets failed to register correctly before shape information was added for this algorithm, so they were not included in the comparison. The region of interest usually extends to 50-75 pixels from the centre, so it can be seen that error is much improved for this central region.

Using shape knowledge to drive registration using RevealMVSTM improves the accuracy of the registration within the region of interest by approximately 2.5mm. The lymph nodes are better aligned when the shape knowledge is used, however there are errors when looking at the outer regions of the image. That is, the central region is aligned well at the expense of the rest of the image. Although it is only crucial that the central region is aligned, it is important to still make the rest of the image align satisfactorily at least in a qualitative way. Although a very accurate result quantitative result is not required as measurements are not being taken there.




Figure 7.8: Results using RevealMVS™ and shape information for patient 6: Top left – Target image with overlaid warp field, Top right – Source image, Bottom left – Warped Source image with overlaid warp field.




Figure 7.9: Spatial distribution of the errors measured from the centre of the region of interest, for 8 datasets, using RevealMVS™.

7.6 Conclusions

Overall, the effect of incorporating shape knowledge into the registration algorithms increases their robustness as it stops them from falling into local minima. The accuracy is improved from using just the algorithms on their own as seen in Table 7.5, but little improvement is gained from just registering using the Thin Plate Splines in most cases. A qualitative comparison can be seen in Fig. 8.16, showing all the results for a slice on data-set 7. Hence, further improvements are to be made in order to increase accuracy. Simultaneous landmarks detection and registration is a method that will increase the accuracy while retaining a high level of robustness. This is described in detail in the next chapter.




Figure 7.10: Qualitative comparison of registration of data-set 7 using each of the four methods. Top left: Target image, Top right: Source image, Middle left: Registration result using B-Splines, Middle right: Registration result using Adaptive registration, Bottom left: Registration result using Fluid registration, Bottom right: Registration result using RevealMVSTM.

Table 7.5: Node Misalignment using Non-rigid registration techniques with and without shape information.

Algorithm	Mean error without shape information (mm)	Mean error with shape information (mm)
Affine	21	21
B-splines	30	8
Adaptive	9	5
Fluid	7	6
RevealMVS TM	7	5

Chapter 8

Mobilising Physiological Knowledge

We saw in the previous chapter that using anatomical knowledge in the form of landmarks vastly improves the robustness of the registration. However, further work is required to improve the precision. Just using the landmark segmentation for the registration, though robust, is not inherently accurate, and using the ‘generic’ registration algorithms often do not improve the registration much more from this initialisation, as seen in Chapter 7. For accurate registration based on landmarks alone, we would require a perfect segmentation, or else the errors in the segmentation may cause even greater errors in the registration. However, it is impossible to achieve a perfect segmentation for the wide range of images that we are observing due to the vast differences in the anatomies and tumours of the patients.

Hence, we introduce here a regularisation that we can use in the registration and which reduces the effects of errors due to the landmark placements. This regularisation will use the physiological knowledge that we have not previously incorporated. That is, we know how we expect the patients to change after treatment, and we can introduce this knowledge into the algorithm.

We incorporate this knowledge into a simultaneous framework whereby the identification of accurate landmark positions, and the registration, are performed at each iteration of the algorithm. This means that we are only incorporating the ‘interesting’ and likely to change features into the registration. Then, the landmark positions and

consequent registration result are improved simultaneously at each step.

We commence this chapter by looking at the nature of the deformation field and how this physiological knowledge can be incorporated as a regularisation. We then go on to describe the simultaneous framework, and finally the registration results.

8.1 Physiological Knowledge

8.1.1 Principal Warps of the Deformation Field

A complete and accurate deformation field can be found using a manual registration whereby corresponding points are plotted on the pre- and post-treatment images, for example, using RevealMVSTM. Although it is not useful in practice as a method of aligning images, since it is time and labour intensive, it is useful for analysing some of the properties of the deformation fields.

Seventy-six landmarks were plotted in two corresponding slices, already rigidly aligned, and the registration results of these two images can be seen in Fig. 8.1. The registration here was performed using the RevealMVSTM software, discussed previously. The deformation field computed can be analysed within the program to give a visualisation of where the changes in the images are greatest, and this can be plotted over the image in colour, Fig. 8.2. The greatest changes are indicated in red, and very little change is blue. Therefore, it can be clearly seen that the greatest changes are occurring in the central region of the image, where the tumour is located and in the mesorectum surrounding the tumour. This is what we would expect to happen from our knowledge of the effects of treatment on this region of the body described in Chapter 1.

A more quantitative approach can be taken by looking at the same deformation performed using Thin Plate Splines. The deformation is a vector that can be viewed as a scalar quantity in the x and y directions respectively. These two deformations f_x and f_y , described in Chapter 7 are visualised in Fig. 8.3. These images on their own are not that useful in gaining information about the nature of the deformation field, and so instead we focus on looking at the principal warps of the deformation field as




Figure 8.1: The 76 manually placed landmarks on the pre- and post-treatment images, and the corresponding registered image.

a whole.

This method of analysing the deformation field was introduced by Bookstein [11]. It is based on the fact that the bending energy matrix used in the calculation of the Thin Plate Splines, $L_n^{-1}KL_n^{-1}$, and mentioned in Chapter 7, is usually highly patterned. Therefore, in order to dissect the ‘features’ of the deformation, we exploit the eigenstructure of the bending energy matrix.

There will always be three zero eigenvalues of the matrix which correspond to the affine displacement, but in order to characterise the non-rigid aspect of the de-




Figure 8.2: An example of registration with RevealMVSTM showing the extent of the deformations.

formation we look at the non-zero eigenvalues and their corresponding eigenvectors. Each non-zero eigenvector that is found can be interpreted as its own Thin Plate Spline, and together they describe the modes of the deformation, or principal warps. The largest eigenvalues correspond to the features with the highest bending energy, and therefore, as the eigenvalues decrease, they correspond to features at successively smaller physical scales.

The principal warps of the colorectal cancer deformation field corresponding to the first five eigenvalues are shown in Fig. 8.4. It can be seen that the first two warps have maximum deformations in the centre of the images, again consistent with our knowledge of the effects of chemotherapy and radiotherapy. The warps corresponding to the lower eigenvalues give more detail of the finer scale features of the warp field.

Cootes and Taylor [24] use the idea of eigenstructure analysis to describe shape and deformation fields when detecting structures in images. However, due to the large variation in patients and sizes and types of tumours, a very large training set would be required for this method to work on colorectal cancer images. Each case is very different so a much more robust method is required. However, the analysis of




Figure 8.3: The deformations in the x and y directions.

the deformation field in this way is a useful tool for driving further work into how the deformation field can be used to indicate the physical nature of changes that we see.

8.1.2 Vector Field Analysis

Vector field analysis is a different way of looking at the nature of the deformation, by describing the three dimensional vector field using a scalar field. The introduction of these methods, by Thirion and Calmon [84], was originally driven by the need to provide an accurate segmentation, in order to quantify changes, such as growth of tumours and plaques in the brain in Multiple Sclerosis patients. They observed that




Figure 8.4: The principal warps of the deformation field.

using segmentation on its own introduced problems when analysing volume changes between times t_1 and t_2 . The volume change is given by $\Delta V = V(t_1) - V(t_2)$, and gives an index of the tumour evolution. The problem was that the uncertainty σ_V of the segmentation was often such that $2\sigma_V > \Delta V$, and such a measurement is of little use. Therefore they proposed a four stage method in order to more accurately quantify ΔV :

1. Rigid Registration of the two datasets.
2. Non-rigid registration of the datasets, i.e. computation of the deformation field.
3. Detection of changing lesions, using the deformation field.
4. Quantification of the volume change, ΔV .

Having performed steps 1 and 2, we now have the deformation or flow field of the registration, $f : (u, v, w)$, which maps the points (x, y, z) onto (x', y', z') , such that:

$$\begin{aligned} x' &= x + u(x, y, z) = \phi_1(x, y, z) \\ y' &= y + v(x, y, z) = \phi_2(x, y, z) \\ z' &= z + w(x, y, z) = \phi_3(x, y, z). \end{aligned} \tag{8.1}$$

Large deformations, i.e. where the shape is changing the most and so $\|f\|$ is large, are of most interest. Also, for the case of the Multiple Sclerosis plaques, they are mostly spherical, and evolve outwards in an even way, so $\text{div}(f) = \frac{\partial u}{\partial x} + \frac{\partial v}{\partial y} + \frac{\partial w}{\partial z}$ will also be large.

Therefore the deformation field can be used to detect changing lesions by analysing the scalar operator of the deformation field:

$$\|f\|\text{div}(f). \tag{8.2}$$

The sign of $\text{div}(f)$ indicates whether the lesion is growing or shrinking, and, as such, the scalar information gained from the deformation field can then be used to give a more accurate segmentation of the changing lesion [84].

It has also been shown by Prima *et al.* [71] and Rey *et al.* [73] that the volume change over time is directly related to the Jacobian of the deformation:

$$\frac{V(t_2) - V(t_1)}{V(t_1)} = \frac{V(t_2)}{V(t_1)} - 1 \simeq Jac - 1. \quad (8.3)$$

Unlike the operator described by Thirion and Calmon, the Jacobian has a physical meaning. The Jacobian of the deformation, ϕ above, is given by:

$$Jacobian = \det(\nabla_P \phi) = \det \begin{pmatrix} \frac{\partial \phi_1}{\partial x} & \frac{\partial \phi_1}{\partial y} & \frac{\partial \phi_1}{\partial z} \\ \frac{\partial \phi_2}{\partial x} & \frac{\partial \phi_2}{\partial y} & \frac{\partial \phi_2}{\partial z} \\ \frac{\partial \phi_3}{\partial x} & \frac{\partial \phi_3}{\partial y} & \frac{\partial \phi_3}{\partial z} \end{pmatrix}. \quad (8.4)$$

It can also be written using the vector displacement field:

$$Jacobian = \det(\nabla_P \phi) = \det(I + \nabla f) = \det \begin{pmatrix} \frac{\partial u}{\partial x} + 1 & \frac{\partial u}{\partial y} & \frac{\partial u}{\partial z} \\ \frac{\partial v}{\partial x} & \frac{\partial v}{\partial y} + 1 & \frac{\partial v}{\partial z} \\ \frac{\partial w}{\partial x} & \frac{\partial w}{\partial y} & \frac{\partial w}{\partial z} + 1 \end{pmatrix}. \quad (8.5)$$

The physical interpretation can be shown in terms of the local variation of volume. If we take two points that are very close $(x, x + \delta x)$, and then compute the new distance between these two points after the deformation:

$$\begin{aligned} \delta x' &= \delta x + f(x + \delta x) - f(x) \\ &= \delta x + \nabla f \cdot \delta x + O(\|\delta x\|^2). \end{aligned} \quad (8.6)$$

We assume δx is small and therefore:

$$\delta x' \simeq (I + \nabla f) \cdot \delta x = (\nabla \phi) \cdot \delta x. \quad (8.7)$$

The local volume variation around that point can now be calculated:

$$\begin{aligned} \delta V' &= \det[\delta x', \delta y', \delta z'] \\ &\simeq \det[\nabla \phi \cdot \delta x, \nabla \phi \cdot \delta y, \nabla \phi \cdot \delta z] \\ &= \det[\nabla \phi] \cdot \det[\delta x, \delta y, \delta z] \end{aligned} \quad (8.8)$$

$$\delta V' = Jac(\phi) \cdot \delta V. \quad (8.9)$$

Therefore the evolution of a small volume is given by the Jacobian and when $Jac(\phi) > 1$ there is a local expansion, and when $Jac(\phi) < 1$ there is a local contraction.

Note that more complex models are available that model the tissue structure as well, but so far this has only been shown to be useful in brain imaging [21].

8.1.3 Jacobian of the Colorectal Cancer Deformation Field

A typical Jacobian of a two dimensional colorectal cancer deformation field was calculated and can be seen as a log plot in Fig. 8.5. Much of the Jacobian is negative which indicates that the area is shrinking. This is consistent with the information gained from analysing the principal warps, and, as such, is also consistent with the effects of chemotherapy and radiotherapy which will cause shrinking in the colorectal and mesorectal region. A cross section through the log of this Jacobian can be seen in Fig. 8.6. The large dip in the centre of the plot can be seen more clearly and corresponds to the colorectum and mesorectum area where the local shrinking is occurring.

Using the deformation field from the 76 landmarks described previously the Jacobian representation is shown as a two dimensional greyscale image in Fig. 8.7. Landmarks were randomly taken away and the Jacobian was recalculated. By the time there were 67 landmarks (i.e. 9 landmarks had been taken away), the Jacobian had lost some of the main central features, as seen in Fig. 8.8. It can also be seen that the smaller physical features of the deformation are more pronounced.

However, the calculation of the deformation field with just 17 selected landmark points (from the original 76) within the mesorectum (the area that we know is most likely to change) gives a good estimate of the Jacobian, Fig. 8.9. This is much more similar to the Jacobian seen in Fig. 8.7, especially for the larger scale deformations in the central region of the image.

Comparing this to using 26 randomly selected landmarks from the 76 original




Figure 8.5: The central section of $\log(\text{Jacobian})$ of a typical deformation of colorectal cancer images.



Figure 8.6: A cross section of $\log(\text{Jacobian})$ of a typical deformation of colorectal cancer images.

gives quite a different result, Fig. 8.10 indicating that some of the landmarks are much more important to the registration than the others.

A similar experiment was performed using the automatically placed control points in the image, placed using the algorithm from Chapter 6. The Jacobian was compared to that calculated from many manually plotted, known, corresponding points. The two Jacobians can be seen in Fig. 8.11. The correlation, discussed in detail in the next section, between the two Jacobians was found to be at least 0.59 indicating that the Jacobian estimate from the segmentation is reasonable. As corresponding control points are found in the non-rigid registration algorithm the Jacobian estimate can be updated at each iteration, but at the same time constrains the landmark positions




Figure 8.7: The Jacobian of the colorectal cancer images with the 76 manually placed landmarks.




Figure 8.8: The Jacobian of the colorectal cancer images with 67 of these manually placed landmarks.




Figure 8.9: The original colorectal cancer images with a selected group of 17 manually placed landmarks, and the corresponding Jacobian of the deformation just using these landmarks.




Figure 8.10: The Jacobian of the colorectal cancer images using a random selection of 26 of the manually placed landmarks




Figure 8.11: A visual comparison of Jacobians using manually selected landmark points and automatically generated points.

according to the physiology.

Therefore, we can use this Jacobian representation to help us to constrain the deformation in the colorectal cancer alignment, and this is simultaneously incorporated into the non-rigid registration algorithm, described in detail in the next section.

8.2 Simultaneous Methods

8.2.1 Simultaneous Segmentation and Registration

It has recently been suggested that simultaneous segmentation and registration is a superior method of registering two images using the known segmented features within those images.

Chen *et al.* [18] showed that by interleaving the segmentation and registration steps within a Bayesian framework, both the segmentation and registration results were improved. They were identifying tissue types in Contrast Enhanced Breast MRI images, whilst aligning the images that were taken over the course of 6 minutes. The images needed registering due to motion of the patient over this period of time. A maximum a-posteriori (MAP) estimation was used for the segmentation and registration stages, and a Markov Random Field was also incorporated as a noise reduction measure. The results showed a reduction in misclassification of pixels within the images.

The problem of aligning and segmenting contrast-enhanced breast images is however, quite a different problem to that of aligning the colorectal MR images, since the colorectal images are taken six months apart, and after extensive chemotherapy and radiotherapy. As such, the deformations are much greater and unpredictable, causing many registration algorithms to fail. Although the problem of aligning the breast MR images is critical for diagnosis, the scale of the changes is much smaller, and the robustness problem is therefore less of an issue.

We have shown that incorporating shape knowledge increases the robustness of the registration, and as such intend to incorporate the shape knowledge, or relational representation into a simultaneous registration and landmark detection algorithm. Since we are not interested in segmenting every pixel in the image into a different class, but are only interested in key features in terms of the segmentation, we will not use a Bayesian MAP approach, but instead simply interleave the landmark detection and registration steps within an adaptive registration framework.

8.2.2 Simultaneous Landmark Detection and Registration Initialisation

An algorithm has already been implemented that finds landmark positions based on a relational representation of the anatomy visible within the images. This was outlined in Chapter 6. It gives a robust anatomical representation due to the relational restrictions, but lacks accuracy due to the segmentation method used and the large variations in the images. Accuracy will be improved by interleaving the landmark detection with the registration and the anticipated shape changes, but an initialisation is required for this process. Therefore, the landmark detection from the relational representation will be used as the initialisation step.

Similarity Measure

The initial landmark estimates, although robust, are still liable to have an error of alignment, and are not guaranteed to be placed accurately. As such, we need to update them. The crucial aspect is not so much their positions on the image, i.e. where exactly on the bladder the landmark is placed, it is more that the positions of the landmarks in the pre-treatment image correspond to the positions of the landmarks in the post-treatment image. We can look at the correspondence of the landmarks using Mutual Information, or Normalised Mutual Information. In this case, since the images are all T2 weighted MRI, we use Mutual Information, given by

$$C_{similarity}(A, B) = H(A) + H(B) - H(A, B) \quad (8.10)$$

to provide a similarity measure between corresponding landmarks points. $H(A)$ and $H(B)$ are the entropies of image segments surrounding the landmark points on the two images and $H(A, B)$ is the joint entropy. We need to apply this similarity measure over a series of scales to give a robust and accurate alignment. Hence the size of the image sub-block, surrounding the landmark points, used to calculate the Mutual Information will be varied.

Registration

The method of Thin Plate Splines [11] are used to register the images throughout this algorithm. This provides a straightforward method by which to align corresponding landmark points. It is fast to calculate the warp field and, as such, fast to calculate information about the warp field that can be used in the regularisation stage (Section 8.2.2).

The Thin Plate Splines are included in an adaptive registration framework similar to that developed by Park *et al.* [66], and described in detail in Chapter 5. This algorithm works iteratively over a series scales, searching for the region of maximum misalignment. Park *et al.* defined a mismatch measure based on Mutual Information $C_{similarity}$.

$$M = 1 - \frac{C_{similarity}(A, B)}{\min(H(A), H(B))} \quad (8.11)$$

This identifies areas of high entropy and low local misinformation which makes for good placement of landmarks. The maximum mismatch can then be found and a landmark is placed and aligned. However, the initial problem with this algorithm was that these landmarks are not within the region of interest. Since we have already defined our landmarks and our goal is to align them, we instead use the minimum value of Mutual Information to decide which are our maximum misaligned landmarks, and we can consequently align these.

The alignment is based on a cost function that is due to the similarity between the landmarks, and the anatomy and physiology relationships on the two images. The cost function is minimised using simple gradient descent. We can use this simple method because local minima are unlikely within the small search region around each landmark point. We have already accounted for the large scale deformations by including these particular landmarks in the first place.

The registration operates across a series of scales from coarse to fine. As the scale becomes finer, the sub-block size used to calculate the mutual information becomes smaller. The scale is refined until a threshold is met or when no further improvement

can be made. That is, the algorithm is stopped when the change in landmark positions is negligible.

Regularisation

Shape knowledge was used in identifying the initial landmarks and this knowledge needs to be incorporated throughout the registration as well, to ensure a fully simultaneous algorithm, making the most of all the information available.

There are two types of information that can be used. Firstly, the anatomical knowledge, which involves information about the relative positions of objects in the images. That is, the colorectum is always lower in the image than the bladder. The bone structures will be level assuming the images are rigidly aligned based on the hip positions, the coccyx is along the line perpendicular to that connecting the hips, from the midpoint of the hips and so on. These are included as a series of weak constraints so that if they are not met then the cost function will increase sharply.

The other type of information that can be used is the physiology. This is the knowledge of how the patient is likely to have changed from one scan to the next. That is, the tumour and surrounding area may have contracted considerably, or alternatively there may be little change. Also, the bladder may be full in one image, and empty in another and so on. We can represent both the anatomical and the physiological knowledge using a scalar representation of the warp field, the Jacobian. The Jacobian $J(A, B)$, described above, gives an indication as to whether the tissue is expanding or contracting in a particular region, and by how much. A first estimate of this physiology change can be found by looking at the warp field of the initial landmarks, estimating the position of the mesorectal fascia and the colorectum, which will give an impression as to which regions are acting in which way.

A final cost measure of the anatomy and physiology is introduced based on the linear correlation between the expected and updated Jacobian representations: the two Jacobians are taken to be the two dimensions of a random variable:

$$(J_{initial}(A, B), J_{update}(A, B)) = (X, Y), \quad (8.12)$$

where $J_{update}(A, B)$ is the Jacobian calculated from the current updated landmarks. The correlation coefficient, which is represented in the cost function is then:

$$C_{physiology}(A, B) = \frac{\sigma_{XY}}{\sigma_X \sigma_Y} \quad (8.13)$$

where $\mu_X = E(X)$, $\mu_Y = E(Y)$, $\sigma_X^2 = E([X - \mu_X]^2)$ and $\sigma_Y^2 = E([Y - \mu_Y]^2)$ are the means and variances of the marginal distributions of X and Y , and σ_{XY} is the covariance of X and Y , given by:

$$\sigma_{XY} = E([X - \mu_X][Y - \mu_Y]) = E(XY) - E(X)E(Y). \quad (8.14)$$

Final Algorithm

Combining landmark detection, adaptive registration and regularisation, the new algorithm is illustrated in Fig. 8.12. The algorithm looks very similar to the adaptive registration of Park *et al.* except for the initialisation using the anatomical knowledge and the cost function which also incorporates that knowledge. The cost function to be minimised is

$$C = -C_{similarity} - \lambda C_{physiology}. \quad (8.15)$$

The choice of the parameter λ will be discussed in the next section.

8.2.3 Implementation

Parameter choices

Various parameters need to be selected for the implementation of this algorithm. Firstly the sub-block size, which is the size of the block of the image over which the similarity measure is calculated. The size of this block indicates the feature size that is being searched for, so in the multi-scale approach a number of sub-block sizes will be used.

```

Do Find rigid features – hips and coccyx
Do Affine registration based on rigid features
Do Find initial estimates of anatomical relational landmarks, and initial Jacobian
    estimate.
For I = large scale to small scale
    Do
        Identify Maximally mismatched landmark using MI
        Optimise Landmark points to minimise cost function
        While Cost function decrease > threshold
End For

```

Figure 8.12: Pseudo-code of the Simultaneous Landmark Detection and Registration algorithm

Similarly, the step size for the gradient descent method needs to be defined. We tested the results for 12 different combinations of sub-block and step sizes and the results can be seen in Fig. 8.13. We assumed three levels of multi-scale refinement. It can be seen in the graph that a smaller step size for a smaller sub-block size is the best combination, which is not surprising since the registration is already initialised well. Consequently, we kept the step size to 1 pixel for all scales and looked across scales from 10 down to 2 pixels. Larger choices of sub-blocks can produce larger errors as they find features that are on a larger scale and therefore less accurate, due to the large scale changes that occur.

The choice of λ is also important for the implementation as this indicates how much the prior information is used over the similarity and vice versa. For $\lambda = 0$ the anatomical and physiological knowledge is not incorporated at all, whereas for much larger values of λ the similarity measure has no effect. We ran the algorithm for a large range of λ values and the results can be seen in Fig. 8.14. The optimum value of λ that we chose was $100\sqrt{10}$. This is large due to the relative magnitudes of $C_{similarity}$ and $C_{physiology}$




Figure 8.13: Effects of changing the sub-block and step size on the results of the algorithm.

Additional landmarks in the Region of Interest

To improve the accuracy of the algorithm further in the mesorectal region of interest, seven extra, equally spaced, landmarks were added within the fat layer automatically. Although they do not necessarily correspond to features or landmarks within the image, they will have a unique pattern due to any lymph nodes, blood vessels or artefacts such as streaking within that region. As such, they are useful to align to provide an even more accurate representation of the registration. The additional landmarks are added using the current points that are known such that we know these landmarks will lie within the region of interest. An example of all the landmarks points used can be seen in Fig. 8.15

8.2.4 Results and Discussion

This algorithm gives a robust estimate of the registration for all 10 pairs of datasets tested. That is, a clinically useful, sensible result is calculated for all cases. The accuracy of the algorithm was measured again using corresponding lymph nodes, and




Figure 8.14: Effects of changing the values of λ on the results of the algorithm, with the dashed line indicating the optimum value of λ




Figure 8.15: Typical distribution of landmarks on an image.

the results are summarised in Table 8.1. The algorithm was run using three sub-block sizes for 3 different scales, of 10, 5 and 2 pixels, and a step size of 1 pixel for the gradient descent. λ was chosen to be $100\sqrt{10}$. The overall mean node misalignment for all 10 datasets is 3.42mm with a standard deviation of 1.61mm. Most enlarged lymph nodes are at least 5mm in diameter and hence an error of 3mm is a satisfactory registration result.

To compare, the adaptive registration algorithm with no shape knowledge input failed on 3 out of the 10 datasets. That is, corresponding lymph nodes could not be identified. This is not clinically useful. The precision of this algorithm was such that the mean distance between corresponding nodes was 8.6mm, when it was run for many iterations. Using the segmentation as an initial alignment, but not interleaving it simultaneously, the algorithm was much more robust and gave a clinically useful result in all cases. However, the precision was such that the misalignment of the nodes was 5mm. These comparisons can be seen in Fig. 8.16 and Table 8.2.

Although the algorithm is more computationally costly when shape knowledge is incorporated, as there is the extra segmentation stage, the results show that when using this shape knowledge the registration is 100% reliable on all the datasets tested to date. The total runtime is approximately 20 minutes in Matlab. Incorporating the shape knowledge into the simultaneous algorithm increases the computation further as the Jacobian is checked at each iteration. However, looking at the increased precision and robustness of the results by incorporating shape knowledge in this way, it can be seen that this compensates for the increased computation, [9], [10].

8.2.5 Conclusions

We have shown that by introducing a simultaneous registration and landmark detection technique we can provide an extremely accurate and robust registration of colorectal cancer images pre- and post-treatment. This is due to the incorporation of the anatomical and physiological knowledge at each iteration of the registration, which prevents the algorithm falling into incorrect local minima.

Table 8.1: Accuracy of Simultaneous Landmark Detection and Registration Algorithm.

Patient	Mean misalignment (mm)	Standard deviation of misalignment (mm)
1	3.15	1.64
2	4.45	3.56
3	1.92	0.49
4	2.00	0.77
5	3.48	2.20
6	3.48	1.16
7	3.65	2.46
8	3.29	0.54
9	5.00	2.54
10	3.48	0.76

Table 8.2: Comparison of the node misalignments (mm) for each of the 10 datasets, running the Park-Meyer algorithm on its own, the Park-Meyer algorithm with shape initialisation, and finally using the simultaneous algorithm.

Patient	Park-Meyer	With initialisation	Simultaneous
1	8.76	3.22	3.15
2	8.77	5.76	4.45
3	12.74	7.08	1.92
4	8.33	3.58	2.00
5	FAIL	7.01	3.48
6	8.76	3.88	3.48
7	5.06	4.06	3.65
8	FAIL	5.36	3.29
9	FAIL	6.36	5.00
10	7.59	3.77	3.48




Figure 8.16: Reliability and Precision of non-rigid registration algorithms with and without the use of shape knowledge.

Chapter 9

Conclusions

In this chapter, we summarise the work that has been described in this thesis and the main contributions. We analyse how far we have got with the original patient management problem, and the limitations of the work to date. We then go on to describe future work that we propose in this area.

9.1 Summary

The aim of this thesis has been to apply computer vision techniques to aid the diagnosis and patient management decisions for Colorectal Cancer. We have focussed on the local MR imaging of the primary tumours and surrounding lymph nodes for this thesis, although we realise this is only a small part of the total challenge.

In Chapter 2, we outlined the underlying principles of MRI and how the images are acquired. We looked at the artefacts that appear on the images, and discussed where we can use them to our advantage in some cases, such as chemical shift, but also where we have to correct for them in the images in other cases, most importantly for the bias field. We reviewed methods of how this bias field can be removed, and finally implemented the method of Styner *et al.* [80]. We devised an initialisation for this method that is suitable for the colorectal images, and then removed the bias across a series of slices, using the knowledge that the field varies only a small amount between adjacent slices.

Having removed the bias field, we then looked at how to start segmenting the

interesting anatomical features from the images, such as the colorectum and the mesorectum, in Chapter 3. A snake was used to segment the colorectum in the first image of a series of slices, then a Kalman filter was used to propagate this snake through the set of slices, updating it at each stage. The snakes that make up the colorectum were plotted in three dimensions to give the clinicians a new way to view the colorectum and the surrounding nodes. Using well defined features in the images, the mesorectum could also be found, either using a complex model based on spherical harmonics, or a simpler model, finding the inscribed circle of the triangle that makes up the anatomical landmarks of the hips and coccyx.

Lymph nodes are very important features that we also wish to detect in the images. Their characteristics are considered crucial for decision making and in determining the circumferential resection margin for surgery. Chapter 4 outlined a method for segmenting the lymph nodes, and then discussed a series of criteria by which to estimate their involvement, i.e. whether or not they contain tumour. The estimate is given as a probability that can be plotted on the images using colour to highlight for the clinicians which areas they should be focussing their attention on. This was validated using a comparison between our algorithm and two expert radiologists.

Patient management decisions for colorectal cancer are made at different stages in the course of treatment. That is, often a patient is re-assessed for surgery having been given a course of downstaging chemotherapy and/or radiotherapy. This assessment involves comparing the patients' scans before and after treatment and analysing what the changes have been. In Chapter 5 we described the methods and results of running a number of 'off-the-shelf', generic registration algorithms to align these images that had been taken, usually six months apart. The results were disappointing with many failures due to the large non-rigid displacements between the corresponding images. The algorithms also focussed on the regions of the image with highest entropy over a range of scales, that were actually not of interest for the alignment of the tumour and lymph nodes.

We looked at how we might possibly constrain these algorithms by incorporating

knowledge of the features within the images. This was done in Chapter 6, where we developed a relational representation of the key features within the images, that it was important to align. Similarity measures and distance functions were used to detect features such that the features found were in the correct positions with respect to each other. In Chapter 7 we incorporated this anatomical knowledge into the generic non-rigid registration algorithms, in order to drive them. This increased the robustness dramatically, eliminating failures.

Finally, we looked at how to incorporate physiological knowledge in Chapter 8. That is, how to include knowledge of the predicted effects of the treatment that the patient has had. We did this by first analysing the deformation field and seeing what information we can get from this. We noted that the Jacobian representation of the deformation field gives quantitative information about the local volume change in the images. This Jacobian can then be used to regularise the registration. We incorporated this regularisation into a simultaneous registration and landmark detection algorithm, which is equally robust, and also improves the accuracy of the registration. A robust and accurate registration provides the clinician with a tool that they can use to observe changes in the tumour and lymph nodes.

9.2 Future Work

Although much progress has been made, this is still only a start with respect to how computer vision can be used to aid clinicians dealing with colorectal cancer. There are many directions that could be taken to build on this work. Three such directions are described in the sections below. They include how contrast-enhanced MRI (CE-MRI) can be used to further aid the diagnosis of lymph nodes, as well as a method of complete validation for such analyses, using histological data. Finally, we outline how imaging techniques can be used not only in diagnosis, but also in treatment, for the case of local excision surgery.

9.2.1 Dynamic Contrast-Enhanced MRI

To date, the image analysis methods demonstrated in this thesis have been based entirely on T2-weighted MR images because this is currently the basis of clinical decision making at the Oxford Radcliffe Hospitals. However, it has been suggested that contrast-enhanced MRI may be useful in aiding the diagnosis, particularly with respect to lymph nodes, for colorectal cancer. Contrast-enhanced imaging is a technique that has been known for a long time, but has only recently come into clinical use with MRI. The idea was originally published by Kety [50], who proposed using inert gases to observe the gas exchange processes taking place in the lungs. This was built upon by Zierler [94], who proposed the idea of using an indicator to trace blood flow, as well as the first mathematical model for doing this. The process of tracing blood flow is often referred to as perfusion imaging. This type of imaging is particularly useful when looking at cancers, since the circulation of blood around tumours has characteristic features. Cancerous tissue usually shows a high and fast uptake due to a vast number of ‘leaky’ microvessels that supply the tumour with blood (angiogenesis). In contrast, normal fatty tissue shows a small and slow uptake of contrast agent. There are a number of different techniques to visually enhance blood flow in MRI, and many more are being developed. For example, Arterial Spin Labelling is a new technique used in brain imaging whereby no extrinsic contrast agent is needed, rather the contrast is provided intrinsically by the MR tagging certain molecules in the aorta that then flow into the brain [32].

The most common clinical method of using CE-MRI, however, is to inject a bolus of paramagnetic contrast agent (such as gadopentetate dimeglumine, Gd-DTPA) and then observe the temporal flow of this agent. Many experiments have been published in the medical literature where contrast agents have been used to provide static images. However, the information that is of particular use is the flow of the contrast agent in the blood over time. The signal-time curves of the contrast agent allow the identification of different types of tissue, and a typical example can be seen in Fig. 9.1. The voxels in the image will need to be classified, and this will be done




Figure 9.1: Idealised signal time curves for Gd-DTPA uptake.

by segmenting the image into areas of fat, tumour, etc., and then comparing the enhancement changes in these regions. Complications may arise due to the partial volume effect often experienced on MR images, and a way to work around this will need to be found.

Known pathological data can also be incorporated as prior knowledge when building these models. The pathology is taken as the ‘ground truth’, and a high-level comparison with the MR data will be useful. The characteristics seen in the pathology will show where radiotherapy has been successful as well as which lymph nodes were cancerous. This information cannot always be found from the normal MR scans. Hence any erroneous results from analysing the MR scans can also be investigated by using the pathology as a reference. A method of automatically analysing this data is needed as the images will be such that the radiologist will not immediately be able to tell which areas have the greatest perfusion. This means the data needs to be analysed, models built and predictions made as to which areas are the most highly perfused. This information will then be represented in a visual format in order for it to be useful to the clinicians.

One possible way of visualising the data is to overlay it on the normal MR image. Depending on whether or not the patient moves between the normal MRI and the

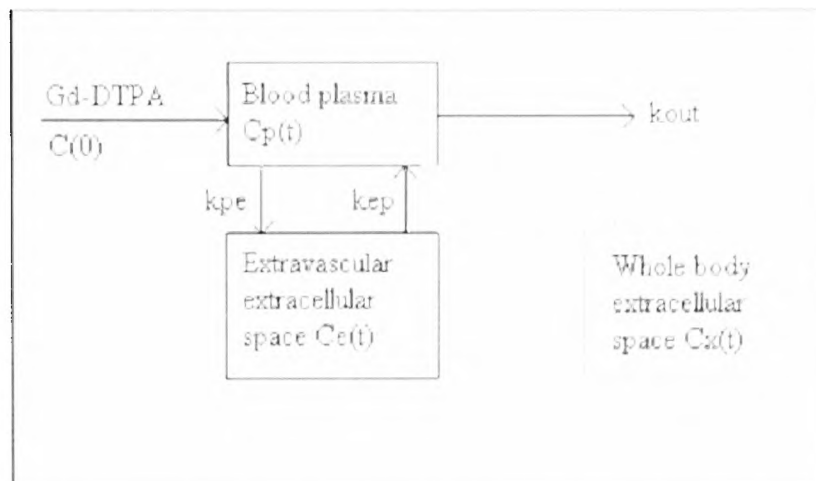


Figure 9.2: Pharmacokinetic model setup.

CE-MRI, registration techniques and motion-correction algorithms may be required to align the two sets of data. A registration algorithm such as the one developed in this thesis could be used.

It is hoped that much information can be gained from using the CE-MRI data and the change in volume of lesions will be able to be quantified. Investigations can also be performed into the perfusion effects of different types of tumours, and also different benign tissues. It is very important that the benign and malignant tissues can be accurately differentiated between.

Pharmacokinetic modelling is the method used to interpret this flow and a number of models have been proposed. A typical example model is shown in Fig. 9.2. Such methods have been shown to work well on breast imaging [2], but have yet to be tested on colorectal data.

As mentioned above, one of the main purposes of performing contrast enhanced imaging is to try and predict which of the lymph nodes are involved. Having gathered the CE-MRI data, this node prediction would be done in two ways, using both explicit data and implicit data. The explicit data is that which has been used for the lymph node detection in this thesis. It involves analysing the data that can be immediately found from the images, such as the size, shape, heterogeneity and connectivity of the regions that could be nodes. However, Kim *et al.* [51] have observed that implicit

data can also play a key role in detection of different regions. That is the probability density functions (pdf) of the different areas on the scans will be analysed to gain extra information about the intensity distributions of involved nodes. The information acquired from the uptake curves will also add to the implicit data associated with the individual nodes. This implicit representation can then be combined with the explicit representation to give a more robust lymph node detector that picks up the very small nodes, hardly visible on the scans, as well as the enlarged ones.

The image analysis methods described here and throughout this thesis need to be in a form that can be easily accessed and used by the clinicians. A program being developed by Kelly *et al.* [49], in collaboration with GE Healthcare, called MDTSuite will be able to view images and include such algorithms. The results will then be available for multidisciplinary meetings each week, where the clinicians make their crucial patient management decisions.

9.2.2 Validation

It has been seen in this thesis that the problem of validation is open and on-going. Currently the best test of our algorithms is by comparison with the opinion of the radiologists and other clinicians. However this does not provide ‘gold standard’ validation. In order to do this, it is currently accepted that histology data is required.

Although the histology data is analysed after the tumour and surrounding region has been excised, this analysis is done separately to that of the MRI analysis and the two results are rarely linked up and compared. In order to perform a full comparative assessment, the excised specimen would need to be scanned using MRI. The specimen is then sliced into 3-5mm macro slices and photographs are taken of these slices. Finally, micro slices that are a few microns thick are stained and analysed and the tumour and lymph nodes tested for their involvement, using standard pathological practice.

Current practice is such that macro and micro slices are taken, the tumour and nodes analysed, and the result that comes back is how many nodes out of those found,




Figure 9.3: Three consecutive macro slices.

were involved. There is no information about the spatial location of these nodes. If the locations of the involved nodes were known then this information could be related back to the MR images, as a gold standard validation.

This would involve a number of registration challenges. Firstly the macro slices that are photographed are not aligned to each other. Some work has already been started by Chen *et al.* [70] in the alignment of these slices. The outer edge of the colorectum is segmented and then aligned across the series of images. Three consecutive slices, with the colorectum outlined, can be seen in Fig. 9.3. The aligned images stacked together in a three dimensional representation are shown in Fig. 9.4.

Having aligned the macro slices to each other these slides need to be registered to the scans of the excised specimen. Also, the MRI scan of the specimen would need to be registered to the pre-surgery scans. These registrations would both involve large scale changes. However, unlike the case of the pre- and post-chemotherapy scans, there are no longer any large anatomical landmarks that can be used for registration. Instead, less obvious landmarks within the mesorectum and tumour would need to be used, and therefore these would have to be patient specific landmarks.

The macro slices that are cut are not necessarily along the same plane as the MRI slices. MRI slices are taken off-plane axial so as to be perpendicular to the

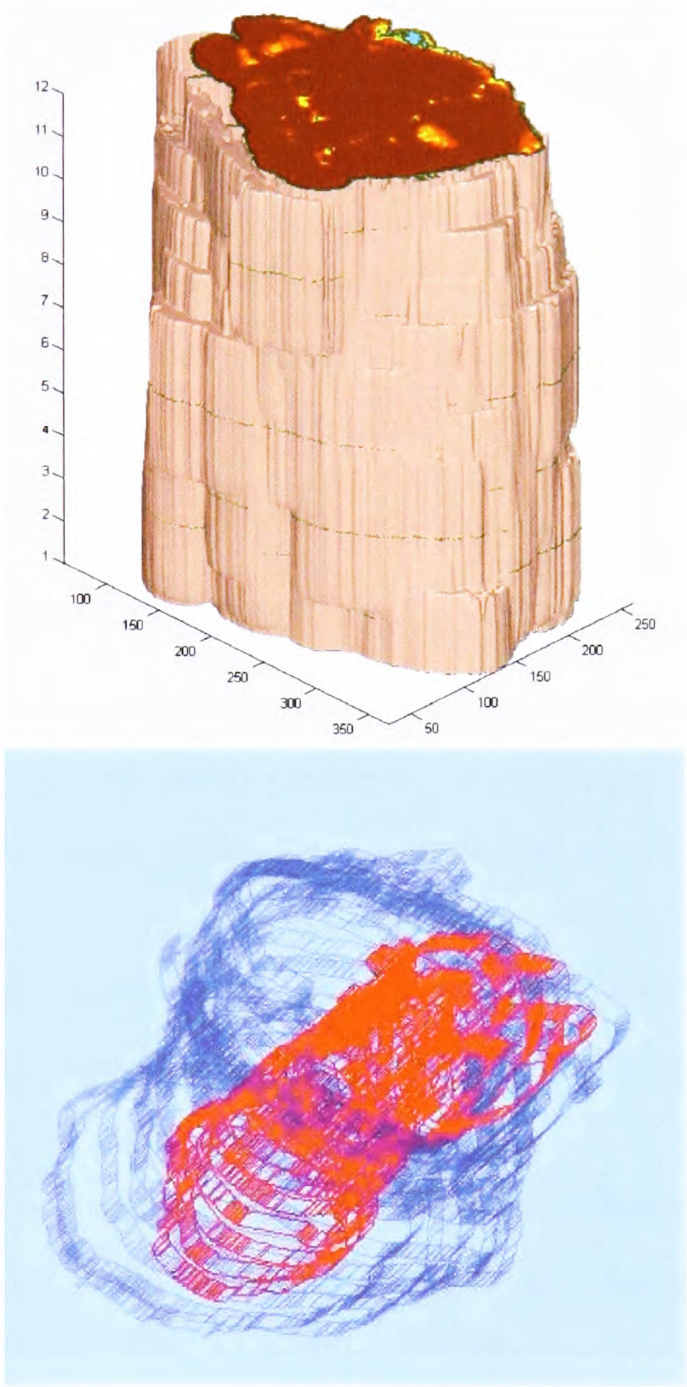


Figure 9.4: Aligned macro slices viewed as a 3-D stack, courtesy of Chen [70].

colorectum, and the histology slices are also perpendicular to the colorectum, but due to errors in both, these angles can vary by up to 20° . Petroudi *et al.* [69] have proposed a method of correcting for this effect, building on the work in this thesis, especially that of Chapter 3, where the colorectum is automatically tracked using a snake and a Kalman filter. This method is then used to affinely align the datasets such that the main axis of the colorectum is aligned.

Having performed these alignments, the lymph nodes present in each dataset can be found. The pathological result as to whether or not each node is involved can then be related back to the original MR scans. This is useful for comparisons for the radiologists who diagnose the nodes in the first place, but also for validation of lymph node detection algorithms, as we will have a gold standard of the status of each node after it has been excised.

9.2.3 Local Excision Surgery

One of the main treatment options for colorectal cancer is surgery. Surgery is performed directly in the case of small tumours, or following on from down-staging chemo/radio therapy for larger tumours. The surgical procedure involves excising a large volume that consists of the tumour, the colorectum and the surrounding mesorectum, and as such is a very invasive procedure. Complications can occur as a result of surgery, especially since many of the patients are old and, in many cases, very frail.

It has been proposed to introduce a much smaller operation to reduce the risks associated with surgery, particularly for those patients who just have a small, T1, or bordering on T2, tumour, with no nodal involvement or metastases. This smaller operation is called local excision surgery and just involves the tumour, and the very small area surrounding it, being taken out. Since only a small volume is being excised the operation is done in a very different way, using minimally invasive procedures.

This idea is yet to be seen in clinical practice due to concern about the risk of metastatic disease. That is, if there is any spread of the cancer further than the

primary tumour, then this surgery won't be curative. Therefore, in order for the operation to be introduced, the diagnosis needs to be certain. A similar procedure for breast cancer was introduced approximately 10 years ago, called lumpectomy. It is a breast conserving surgery, involving removing just the main tumour and the immediately surrounding tissue. Although it was met with some scepticism at first, it is now the most common form of breast surgery, and is usually performed in conjunction with either chemo or radiotherapy. It is hoped that the local excision surgery for colorectal cancer will have similar success.

However, despite being minimally invasive, the local excision surgery for colorectal cancers is not straightforward, and there are a number of challenges that need to be addressed. The first is in diagnosis and decision making. The tumour needs to be small and not extending into the mesorectal fat. Also, there can be no involved lymph nodes. Incorporation of lymph node detection algorithms, discussed previously, would add an extra layer of 'checking' as the tumour is staged.

The second challenge lies in surgery planning. As the procedure is so precise, a plan is made beforehand by the surgeons so they know exactly which area they are removing. This plan is done using the CT and MR images that were used in the diagnosis. They need the most recent scans in order to do these plans, and 3-D visualisation is also helpful when looking at where the tumour is and how the surgeon is going to get to it.

Finally, as part of the procedure, the patient will have a trans-anal ultrasound probe inserted into them in order for the surgeon to see exactly the position and level of invasion of the tumour. An example of such an ultrasound image is given in Fig. 9.5. The image shows the layers that make up the wall of the colorectum, the submucosa and the muscle layers, and the tumour is visible in these regions indicating how far it has extended through the layers. Segmentation algorithms will be useful in order to analyse these images to highlight the different regions and the tumour.

Registration will also be required in order to compare the ultrasound images to the MR and CT planning scans so that the surgeon knows which part of the body




Figure 9.5: A typical trans-anal ultrasound image.

they are looking at with respect to the tumour location. Multimodality registration will therefore be required. Mellor *et al.* [64] suggested using local phase as a method for registration of Ultrasound to other modalities, and this would certainly be a good starting point for this research. Fast segmentation and registration methods are also required, such that they can be performed in real time as the operation proceeds.

Appendix A

Imaging Orientations

The standard orientations of MR images are sagittal, coronal and axial, and these can be seen in Fig. A.1. When imaging the colorectum however, off-plane axial images are used that are perpendicular to the rectum. The direction of these images can be seen on the sagittal MR image in Fig. A.2.




Figure A.1: MR Imaging Directions.



Figure A.2: Sagittal MRI showing the direction of the off-plane axial images, perpendicular to the rectum.

Bibliography

- [1] M. N. AHMED, S. M. YAMANY, N. MOHAMED, A. A. FARAG, AND T. MORIARTY. A modified fuzzy C-means algorithm for bias field estimation and segmentation of MRI. *IEEE TMI*, **21**:193–199, 2002.
- [2] P. ARMITAGE, C. BEHRENBRUCH, M. BRADY, AND N. MOORE. Extracting and visualizing physiological parameters using dynamic contrast-enhanced magnetic resonance imaging of the breast. *MIA*, **9**:315–329, 2005.
- [3] C. P. BEHRENBRUCH, S. PETROUDI, S. L. BOND, J. D. DECLERCK, F. J. LEONG, AND J. M. BRADY. Image filtering techniques for medical image post-processing: an overview. *British Journal of Radiology*, **77**:126–132, 2004.
- [4] R. BEICHEL, H. BISCHOF, F. LEBERL, AND M. SONKA. Robust active appearance models and their application to medical image analysis. *IEEE TMI*, **24**:1151–1169, 2005.
- [5] A. BLAKE, R. CURWEN, AND A. ZISSERMAN. A framework for spatio-temporal control in the tracking of visual contours. *IJCV*, **11**:127–145, 1993.
- [6] S. BOND AND M. BRADY. Image analysis for patient management in colorectal cancer. *CARS*, 2005.
- [7] S. BOND AND M. BRADY. Non-rigid registration for colorectal cancer MR images. *MIUA*, 2005.
- [8] S. BOND AND M. BRADY. Non-rigid registration for colorectal cancer MR images. *CVBIA LNCS 3765*, 2005.
- [9] S. BOND AND M. BRADY. Simultaneous registration and landmark detection. *MIUA*, 2006.
- [10] S. BOND AND M. BRADY. Simultaneous registration and landmark detection. *EMBC*, 2006.
- [11] F. BOOKSTEIN. Principal warps: Thin-plate splines and the decomposition of deformations. *IEEE PAMI*, **11**:567–585, 1989.
- [12] Y. BOYKOV, O. VEKSLER, AND R. ZABIH. Fast approximate energy minimization via graph cuts. *IEEE TPAMI*, **23**:1222–1239, 2001.

- [13] C. BRECHBUHLER, G. GERIG, AND O. KUBLER. Surface parametrization and shape description. *Visualization in Biomedical Computing*, 1992.
- [14] G. BROWN, I. R. DANIELS, C. RICHARDSON, P. REVELL, D. PEPPERCORN, AND M. BOURNE. Techniques and troubleshooting in high spatial resolution thin slice MRI for rectal cancer. *BJR*, **78**:245–251, 2005.
- [15] G. BROWN, C. RICHARDS, M. BOURNE, R. NEWCOMBE, A. RADCLIFFE, N. DALLIMORE, AND G. WILLIAMS. Morphologic predictors of lymph node status in rectal cancer with use of high-spatial-resolution MR imaging with histopathologic comparison. *Radiology*, **227**:371–377, 2003.
- [16] G. BROWN, C. J. RICHARDS, R. G. NEWCOMBE, N. S. DALLIMORE, A. G. RADCLIFFE, D. P. CAREY, M. W. BOURNE, AND G. T. WILLIAMS. Rectal carcinoma: Thin section MRI for staging 28 patients. *Radiology*, **211**:215–222, 1999.
- [17] T. F. CHAN AND L. A. VESE. Active contours without edges. *IEEE TIP*, **10**:266–277, 2001.
- [18] X. CHEN, M. BRADY, AND D. RUECKERT. Simultaneous segmentation and registration for medical images. *MICCAI*, pages 663–670, 2004.
- [19] G.E. CHRISTENSEN, B. CARLSON, K.S.C. CHAO, AND ET AL. Image-based dose planning of intracavitary brachytherapy: Registration of serial-imaging studies using deformable anatomic templates. *Int. J. Radiation Oncology, Biol & Phys.*, **17**:227–243, 2001.
- [20] G.E. CHRISTENSEN, R. D. RABBITT, AND M. I. MILLER. Deformable templates using large deformation kinematics. *IEEE TIP*, **5**:1435–1447, 1996.
- [21] O. CLATZ, M. SERMESANT, P.-Y. BONDIAU, H. DELINGETTE, S. WARFIELD, G. MALANDAIN, AND N. AYACHE. Realistic simulation of the 3-d growth of brain tumors in MR images coupling diffusion with biomechanical deformation. *IEEE TMI*, **24**:1334–1346, 2005.
- [22] L. D. COHEN. On active contour models and balloons. *CVGIP: Image Understanding*, 1991.
- [23] L. D. COHEN AND I. COHEN. Finite-element methods for active contour models and balloons for 2D and 3D images. *IEEE TPAMI*, **15**:1131–1147, 1993.
- [24] T. F. COOTES, C. J. TAYLOR, D. H. COOPER, AND J. GRAHAM. Active shape models their training and application. *CVIU*, **61**:38–59, 1995.
- [25] CRUK. Bowel cancer information sheet. 2003.
- [26] W. R. CRUM, L. D. GRIFFIN, D. L. G. HILL, AND D. J. HAWKES. Zen and the art of medical image registration: Correspondence, homology, and quality. *NeuroImage*, **20**:1425–1437, 2003.

- [27] W.R. CRUM, D.L.G. HILL, AND D.J. HAWKES. Information theoretic similarity measures in non-rigid registration. *IPMI, LNCS 2732*, pages 378–387, 2003.
- [28] M. B. CUADRA, J. GOMEZ, P. HAGMANN, C. POLLO, J.-G. VILLEMURE, B. M. DAWANT, AND J.-P. THIRAN. Atlas-based segmentation of pathological brains using a model of tumor growth. *LNCS*, **2488**:380–387, 2002.
- [29] F. P. CURRA AND L. A. CRUM. Therapeutic ultrasound: Surgery and drug delivery. *Acoust. Sci. & Tech.*, **24**:343–348, 2003.
- [30] SHATTUCK. D., R. SANDOR-LEAHY, K. SCHAPER, D. ROTTENBERG, AND R. LEAHY. Magnetic resonance image tissue classification using a partial volume model. *NeuroImage*, **13**:856–876, 2001.
- [31] B. M. DAWANT, S. L. HARTMANN, AND S. GADAMSETTY. Brain atlas deformation in the presence of large space-occupying tumors. *LNCS*, **1679**:589–596, 1999.
- [32] J. A. DETRE, W. ZHANG, D. A. ROBERTS, A. C. SILVA, D.S. WILLIAMS, D. J. GRANDIS, A. P. KORETSKY, AND J. S. LEIGH. Tissue specific perfusion imaging using arterial spin labelling. *NMR Biomed*, **7**:75–82, 1994.
- [33] A. DZIK-JURASZ, C. DOMENIG, M. GEORGE, J. WOLBER, A. PADHANI, G. BROWN, AND S. DORAN. Diffusion MRI for prediction of response of rectal cancer to chemoradiation. *Lancet Research Letters*, **360**:307–308, 2002.
- [34] GLEESON. F. New techniques for imaging colorectal cancer: Use of MRI, PET and radioimmunoscinotography. *Oxford Radcliffe Hospitals Trust, Tech. Report*, 2003.
- [35] J. FELDMAR AND N. AYACHE. Rigid, affine and locally affine registration of free-form surfaces. *IJCV*, **18**:99–119, 1996.
- [36] P. F. FELZENSZWALB AND D. P. HUTTENLOCHER. Pictorial structures for object recognition. *CVPR*, 2000.
- [37] M. FISCHLER AND R. BOLLES. Random sample consensus: a paradigm for model fitting with applications to image analysis and automated cartography. *Communications of the ACM*, **24**:381–395, 1981.
- [38] S. FRANTZ, K. ROHR, AND H. S.. STIEHL. A new approach to the localization of 3D anatomical point landmarks in medical images based on deformable models. *DAGM Symposium, Springer-Verlag*, **22**:452–459, 2000.
- [39] X. GENG, D. KUMAR, AND G. E. CHRISTENSEN. Transitive inverse-consistent manifold registration. *LNCS*, **3565**:468–479, 2005.
- [40] G. GERIG, O. KUBLER, R. KIKINIS, AND F. A. JOLESZ. Nonlinear anisotropic diffusion of MRI data. *IEEE TMI*, **11**:221–232, 1992.

- [41] S. GILLES, M. BRADY, J. DECLERCK, J. P. THIRION, AND N. AYACHE. Bias field correction of breast MR images. *LNCS*, **1131**:153–158, 1996.
- [42] A. GUILLEMAUD AND M. BRADY. Estimating the bias field of MR images. *IEEE TMI*, **12**:238–251, 1997.
- [43] A. GUIMOND, R. ROCHE, N. AYACHE, AND J. MEUNIER. Three-dimensional multimodal brain warping using the demons algorithm and adaptive intensity corrections. *IEEE TMI*, **20**:58–69, 2001.
- [44] R. HASHEMI, W. BRADLEY, AND C. LISANTI. *MRI The Basics*. Lippincott Williams & Williams, 2003.
- [45] C. JARQUE AND A. BERA. A test for normality of observations and regressions residuals. *International Statistical Review*, **55**:163–172, 1987.
- [46] N. JOSHI AND M. BRADY. A non-parametric model for partial volume segmentation of MR images. *BMVC*, 2005.
- [47] T. KADIR AND M. BRADY. Saliency, scale and image description. *IJCV*, **45**:83–105, 2001.
- [48] M. KASS, A. WITKIN, AND D. TERZOPOULOS. Snakes: Active contour models. *Int. J. of Computer Vision*, **55**:321–331, 1988.
- [49] M. KELLY AND M. BRADY. To follow: mkelly@robots.ox.ac.uk.
- [50] S. S. KETY. The theory and application of the exchange of inert gas at lungs and tissues. *Pharmacol. Review*, **3**:1–41, 1951.
- [51] J.-G. KIM AND M. BRADY. To follow: jgkim@robots.ox.ac.uk.
- [52] D. KOH, G. BROWN, L. TEMPLE, A. RAJA, P. TOOMEY, N. BETT, A. NORMAN, AND J. HUSBAND. Rectal cancer: Mesorectal lymph nodes at MR imaging with USPIO versus histopathologic findingsinitial observations. *Radiology*, **231**:91–99, 2004.
- [53] E. KREYZIG. *Advanced Engineering Mathematics*. John Wiley & Sons Inc., 1997.
- [54] S. K. KYRIACOU, C. DAVATZIKOS, S. J. ZINREICH, AND R. N. BRYAN. Non-linear elastic registration of brain images with tumor pathologyusing a biomechanical model [MRI]. *IEEE TMI*, **18**:580–592, 1999.
- [55] J. LEE, J. HEIKEN, D. LING, H. GLAZER, D. BALFE, R. LEVITT, W. DIXON, AND W. MURPHY. Magnetic resonance imaging of abdominal and pelvic lymphadenopathy. *Radiology*, **153**:181–188, 1984.
- [56] L. LJUNG AND T. SODERSTROM. *Theory and Practice of Recursive Identification*. MIT Press, 1987.

- [57] S. MANAY. Phd thesis: Applications of anti-geometric diffusion to computer vision: Thresholding, segmentation and distance functions. 2003.
- [58] S. MANAY AND A. YEZZI. Anti-geometric diffusion for adaptive thresholding and fast segmentation. *IEEE TMI*, **12**:1310–1323, 2003.
- [59] G. E. MARAI, D. H. LAIDLAW, AND J. J. CRISCO. Super-resolution registration using tissue-classified distance fields. *IEEE TMI*, **25**:177–187, 2006.
- [60] A. MARTLING, T. HOLM, S. BREMMER, J. LINDHOLM, B. CEDERMARK, AND L. BLOMQVIST. Prognostic value of preoperative magnetic resonance magnetic imaging of the pelvis in rectal cancer. *BJS*, **90**:1422–1428, 2003.
- [61] C. MCARDLE, D. KERR, AND P. BOYLE. *Colorectal Cancer*. Isis Medical Media Ltd, 2000.
- [62] T. MCINERNEY AND D. TERZOPOULOS. Deformable models in medical image analysis: A survey. *MIA*, **1**:91–108, 1996.
- [63] R. MEEGAMA AND J. RAJAPAKSE. Nurbs snakes. *Image and Vision Computing*, **21**:551–562, 2003.
- [64] M. MELLOR AND M. BRADY. Non-rigid multimodal image registration using local phase. *MIA*, **4**:330–343, 2005.
- [65] S. OSHER AND J. A. SETHIAN. Fronts propagating with curvature-dependent speed- algorithms based on hamilton-jacobi formulations. *J. Computational Physics*, **79**:12–49, 1988.
- [66] H. PARK, P.H. BLAND, K.K. BROCK, AND C.R. MEYER. Adaptive registration using local information measures. *MIA*, **8**:465–473, 2004.
- [67] S. PERIASWAMY AND H. FARID. Medical image registration with partial data. *MIA*, **10**:452–464, 2005.
- [68] P. PERONA AND J. MALIK. Scale-space and edge detection using anisotropic diffusion. *IEEE TPAMI*, **12**:629–639, 1990.
- [69] S. PETROUDI, G. BROWN, S. BOND, AND M. BRADY. Circumferential resection margin assessment on MRI of rectal cancer. *To appear: EMBC*, 2006.
- [70] J. PITT-FRANCIS, D. CHEN, M. SLAYMAKER, A. SIMPSON, M. BRADY, I. VAN LEEUWEN, F. REDDINGTON, P. QUIRKE, G. BROWN, AND D. GAVAGHAN. Multimodal imaging techniques for the extraction of detailed geometrical and physiological information for use in multi-scale models of colorectal cancer and treatment of individual patients. *J. of Theoretic Medicine*, To appear.
- [71] S. PRIMA, J. P. THIRION, G. SUBSOL, AND N. ROBERTS. Automatic analysis of normal brain dissymmetry of males and females in MR images. *LNCS*, **1496**:770–779, 1998.

- [72] P. QUIRKE. Training and quality assurance for rectal cancer: 20 years of data is enough. *Lancet Oncol.*, 4:695–702, 2003.
- [73] D. REY, G. SUBSOL, H. DELINGETTE, AND N. AYACHE. Automatic detection and segmentation of evolving processes in 3D medical images: Application to multiple sclerosis. *IPMI LNCS 1613*, pages 154–167, 1999.
- [74] G.K. ROHDE, A. ALDROUBI, AND B.M. DAWANT. The adaptive bases algorithm for intensity-based nonrigid image registration. *IEEE TMI*, 22:1470–1479, 2003.
- [75] D. RUECKERT, L.I. SONODA, C. HAYES, D.L.G. HILL, M.O. LEACH, AND D.J. HAWKES. Nonrigid registration using free-form deformations: Application to breast MR images. *IEEE TMI*, 18:712–721, 1999.
- [76] J. G. SLED, A.P. ZIJDENBOS, AND A. C. EVANS. A non-parametric method for automatic correction of intensity nonuniformity in MRI data. *IEEE TMI*, 17:87–97, 1998.
- [77] R. A. J. SPENCE AND P. G. JOHNSON. *Oncology: Oxford Core texts*. OUP, 2001.
- [78] L. H. STAIB, , AND J. S. DUNCAN. Boundary finding with parametrically deformable models. *IEEE TPAMI*, 14:1061–1075, 1992.
- [79] L. H. STAIB, , AND J. S. DUNCAN. Model-based deformable surface finding for medical images. *IEEE TMI*, 15:720–731, 1996.
- [80] M. STYNER, C. BRECHBUHLER, G. SZCKELY, AND G. GERIG. Parametric estimate of intensity inhomogeneities applied to MRI. *IEEE TMI*, 19:153–165, 2000.
- [81] R. SUMMERS, A. JEREBKO, M. FRANASZEK, J. MALLEY, AND D. JOHNSON. Colonic polyps: Complementary role of computer-aided detection in CT colonography. *Radiology*, 225:391–399, 2002.
- [82] R. SUMMERS, J. YAO, AND C. D. JOHNSON. CT colonography with computer-aided detection: Automated recognition of ileocecal valve to reduce number of false-positive detections. *Radiology*, 233:266–272, 2004.
- [83] G. SZEKELY, A. KELEMEN, C. BRECHBUHLER, AND G. GERIG. Segmentation of 3D objects from MRI volume data using constrained elastic deformations of flexible fourier surface models. *MIA*, 1:19–34, 1996.
- [84] J.-P. THIRION AND G. CALMON. Deformation analysis to detect and quantify active lesions in three-dimensional medical image sequences. *IEEE TMI*, 18:429–441, 1999.

- [85] M. TINCHER, C.R. MEYER, R. GUPTA, AND D.M. WILLIAMS. Polynomial modeling and reduction of RF body coil spatial inhomogeneity in MRI. *IEEE TMI*, **12**:361–365, 1993.
- [86] A. TSAI, W. WELLS, C. TEMPANY, E. GRIMSON, AND A. WILLSKY. Coupled multi-shape model and mutual information for medical image segmentation. *LNCS*, **2732**:185–197, 2003.
- [87] A. TSAI, A. YEZZI, W. WELLS, C. TEMPANY, D. TUCKER, A. FAN, W. E. GRIMSON, AND A. WILLSKY. A shape-based approach to the segmentation of medical imagery using level sets. *IEEE TMI*, **22**:137–154, 2003.
- [88] P. VIOLA AND W. WELLS. Alignment by maximization of mutual information. *IJCV*, **24**:137–154, 1997.
- [89] G. WELCH AND G. BISHOP. An introduction to the kalman filter. *University of North Carolina: Tech. Report*, 2001.
- [90] W. M. WELLS, W. E. L. GRIMSON, R. KIKINIS, AND F. A. JOLESZ. Adaptive segmentation of MR data. *IEEE TMI*, **15**:429–442, 1996.
- [91] Z. WU. Compactly supported positive definite radial functions. *Adv. Comput. Math.*, **4**:283, 1995.
- [92] J. YANG, L. H. STAIB, AND J. S. DUNCAN. Neighbor-constrained segmentation with 3D deformable models. *LNCS*, **2732**:198–209, 2003.
- [93] L. YING, Z.-P. LIANG, D. C. MUNSON, R. KOETTER, AND B. J. FREY. Unwrapping of MR phase images using a markov random field model. *IEEE TMI*, **25**:128–136, 2006.
- [94] K. L. ZIERLER. Theory of use of indicators to measure blood flow and extracellular volume and calculation of transcapillary movement of tracers. *Circ. Res.*, **12**:464–471, 1963.