

Thesis Submitted for the Degree of
Doctor of Philosophy

EFFICIENT PRE-SEGMENTATION
FILTERING IN MRCP

Author: Kevin Robinson, B.Eng., M.Sc.

Supervisor: Professor Paul F. Whelan

Dublin City University
School of Electronic Engineering
September 2005

I hereby certify that this material, which I now submit for assessment on the programme of study leading to the award of Doctor of Philosophy is entirely my own work and has not been taken from the work of others save and to the extent that such work has been cited and acknowledged within the text of my work.

Signed: _____
Candidate

ID No.: 51169452

Date: 21st September 2005

Acknowledgement

I wish to acknowledge the support and assistance of my Ph.D. supervisor Prof. Paul F. Whelan, and that of Dr Ovidiu Ghita and all the members of the Vision Systems Group at Dublin City University, with whom I have had the pleasure of working over the past number of years.

My thanks also to the Mater Misericordiae Hospital, Dublin, which provided funding for this project, and in particular to Dr John Stack, Director of Radiology, for offering his valuable clinical perspective on the development and assessment of this work.

Contents

Acknowledgement	i
Abstract	iii
List of Figures	iv
List of Tables	vi
Glossary of Acronyms	vii
1 Introduction	1
1.1 Background and Motivation	3
1.2 Contributions	19
1.3 Thesis Outline	22
2 Intensity Non-uniformity Correction	25
2.1 Types of Intensity Non-uniformity	28
2.2 Data Characterisation	33
2.3 Histogram Matching	36
2.4 Non-uniformity Correction Results	51
3 Adaptive Gaussian Smoothing	56
3.1 Gradient-Weighted Gaussian Filter	57
3.2 Elliptic Filter Model	68
3.3 Filter Characterisation and Performance	74
4 Greyscale Reconstruction	80
4.1 Reconstruction by Dilation	80
4.2 Hybrid Reconstruction	86
4.3 Downhill Filter	92
5 Implementations	106
5.1 Processing Framework	107
6 Conclusion	120
6.1 Summary	120
6.2 Discussion and Further Work	122
6.3 Publications Arising	123
Appendices	
A Body Fat Analysis	A-1
B NeatMRI	B-1
Bibliography	

Efficient Pre-segmentation Filtering in MRCP

Kevin Robinson

Abstract

Magnetic Resonance Cholangiopancreatography (MRCP) is an evolving MRI technique designed for the imaging of the biliary tree, a system of narrow ducts that collect bile, produced within the liver, store it in the gall bladder, and deliver it into the small intestine as needed. Current MRCP protocols, used to diagnose problems in this ductal system, generate cluttered and noisy, low resolution, non-isometric volume data, often with significant intensity non-uniformities. This combination of undesirable characteristics presents particular challenges for the application of automated image analysis techniques.

This thesis examines the development, characterisation, and testing of novel and efficient pre-segmentation filtering procedures designed to achieve increased robustness and precision in the subsequent segmentation and analysis of the biliary tree from MRCP data. A focused set of image preprocessing algorithms has been developed so as to facilitate the operation of non-complex segmentation and computer assisted diagnosis (CAD) procedures. Most notable in this regard are a number of novel techniques designed to address the key areas of this image processing task. These techniques consist of:

- a new histogram preserving approach to inter-image and inter-volume intensity non-uniformity correction,
- a highly versatile adaptive smoothing filter, implemented as an oriented, scaled and shaped ellipsoid filter mask,
- the *downhill filter*, an efficient new algorithm for morphological reconstruction by dilation, and
- a novel approach to the reconstruction of fine branching structures in noisy volume data.

Through this combination of flexible and efficient preprocessing algorithms, an effective route towards robust MRCP segmentation and analysis, and routine CAD in the assessment of the biliary tree from MRCP is presented.

List of Figures

1.1	The pancreato-biliary system	1
1.2	Stones in the gall bladder and common bile duct	2
1.3	A typical RARE image	4
1.4	Variable visualisation in RARE images	5
1.5	Two consecutive images from an axial HASTE sequence	7
1.6	Three consecutive images from a coronal HASTE sequence	7
1.7	Two images from a TRUFI sequence	8
1.8	An ERCP examination	9
2.1	Histogram with spikes and voids	27
2.2	Intra-image intensity non-uniformity	29
2.3	Intra-image non-uniformity in MRCP	30
2.4	Inter-image intensity non-uniformity	31
2.5	Inter-volume intensity non-uniformity	32
2.6	Typical data histogram	33
2.7	Histogram with third peak	34
2.8	Three coronal slices and their histograms	35
2.9	Piecewise linear histogram scaling	47
2.10	Histogram before and after scaling	47
2.11	Three coronal slices before and after matching	52
2.12	Whole body data before and after matching	54
3.1	Graph of the function: $y = e^{-x^2}$	57
3.2	Gaussians of varying widths	58
3.3	3-D distance maps	59
3.4	Non-isometric data grid	61
3.5	A voxel's 26-neighbourhood	62
3.6	Uniform gradients in three orientations	62
3.7	Non-isometric gradient filter masks	63
3.8	3-D x mask	65
3.9	Distance from a point to a plane	66
3.10	Five points along the mask shape continuum	69
3.11	The form of an ellipse	69
3.12	A family of ellipses	70
3.13	Anisotropic filter parameter space	71
3.14	Parameter space in λ and μ	73
3.15	Filtering results	74
3.16	Closeup of two ducts under filtering	76
3.17	Region smoothing versus edge retention	77

3.18 Filtering results	78
4.1 Narrow branch preservation in hybrid reconstruction	81
4.2 Illustration of the ductal tree	82
4.3 Segmenting a brain image	84
4.4 A comparison of dilation techniques in 1-D	85
4.5 Dilations iterated to stability	86
4.6 Reconstruction results in two datasets	89
4.7 Reconstruction difference images	90
4.8 Reconstruction results	91
4.9 Filtering results on three test images	99
4.10 Exhaustive and non-exhaustive neighbourhoods	104
4.11 The grassfire distance transform	105
5.1 Cubic interpolation model	108
5.2 Tri-cubic interpolation	109
5.3 Intensity non-uniformity correction in WB-MRI	110
5.4 Unsmoothed and smoothed axial HASTE data	111
5.5 Adaptive smoothing in MRI data	111
5.6 Hybrid reconstruction results	112
5.7 Reconstruction in whole body MRI	113
5.8 Volume and surface rendered biliary trees in good data	115
5.9 Volume and surface rendered biliary trees in poor data	116
5.10 Portion of a triangulated surface	117
5.11 Stones in the common bile duct	118
5.12 MIP of stones in the common bile duct	118
A.1 Five coronal slices	A-3
A.2 Unnormalised and normalised images	A-5
A.3 Unnormalised and normalised histograms	A-5
A.4 Volume reconstruction in WB-MRI	A-6
A.5 Data smoothing in WB-MRI	A-7
A.6 Segmentation results in a coronal section	A-9
A.7 Thresholding versus adaptive classifier	A-10
A.8 System results display	A-11
A.9 Orthogonal section viewer	A-13
A.10 Volume rendering tool	A-13
A.11 Medial cutaway view	A-14
A.12 Graph of BMI against percentage body fat	A-15
A.13 Graph of calculated against measured BMI	A-17
A.14 Thigh section rendering	A-18
A.15 Two coronal and two sagittal cross sections	A-19

List of Tables

2.1	Voxel ordering discriminant function effectiveness	55
3.1	Smoothing results from five approaches	79
4.1	Symbol definitions.	84
4.2	Execution timings for five test images	100
4.3	Execution timings for five test volumes	102
4.4	Standard distance metrics	105
5.1	DICOM header fields	108
A.1	Body fat results from 42 WB-MRI datasets	A-12
A.2	Standardised body mass index (BMI) categories	A-15
B.1	Parameter differences between C and Java.	B-3
B.2	Table of documented routines	B-10

Glossary of Acronyms

Acronym	–	Explanation
1-D	–	One Dimensional
2-D	–	Two Dimensional
3-D	–	Three Dimensional
AV	–	Ampulla of Vater
BMI	–	Body Mass Index
CAD	–	Computer Assisted Diagnosis
CBD	–	Common Bile Duct
CD	–	Cystic Duct
CHD	–	Common Hepatic Duct
CSF	–	Cortico-Spinal Fluid
CT	–	Computed Tomography
CTA	–	Computed Tomography Angiography
CTC	–	Computed Tomography Colonography
DICOM	–	Digital Imaging and COmmunications in Medicine
ERCP	–	Endoscopic Retrograde CholangioPancreatography
FIFO	–	First In First Out
GB	–	Gall Bladder
GI	–	GastroIntestinal
GUI	–	Graphical User Interface
HASTE	–	Half-fourier Acquisition Single-shot Turbo spin-Echo
HD	–	Hepatic Duct
IV	–	Intra-Venous
LHD	–	Left Hepatic Duct
LIFO	–	Last In First Out
LUT	–	Look-Up Table
MC	–	Marching Cubes
MIP	–	Maximum Intensity Projection
MNP	–	MiNimum intensity Projection
MR	–	Magnetic Resonance
MRA	–	Magnetic Resonance Angiography
MRC	–	Magnetic Resonance Cholangiography
MRCP	–	Magnetic Resonance CholangioPancreatography
MRI	–	Magnetic Resonance Imaging
NMR	–	Nuclear Magnetic Resonance

PD	–	Pancreatic Duct
RARE	–	Rapid Acquisition by Relaxation Enhancement
RHD	–	Right Hepatic Duct
ROI	–	Region Of Interest
SENSE	–	SENSitivity Encoding
SMASH	–	SiMultaneous Acquisition of Spatial Harmonics
SNR	–	Signal to Noise Ratio
SSD	–	Shaded Surface Display
TrueFISP	–	True Fast Imaging with Steady-state Precession
TRUFI	–	see TrueFISP
VC	–	Virtual Colonoscopy
VE	–	Virtual Endoscopy
WB-MRI	–	Whole Body Magnetic Resonance Imaging

Chapter 1

Introduction

The pancreato-biliary system (consisting of the pancreatic duct and biliary tree, see Fig. 1.1) is routinely examined by radiologists using a set of MRI acquisition protocols collectively referred to as Magnetic Resonance Cholangiopancreatography or MRCP. The data generated by this class of MR imaging protocol typically exhibits a number of undesirable qualities (poor signal to noise ratio, low spatial resolution, non-isometric voxels, greylevel inhomogeneity, limited coverage, and variable visualisation of the ductal system) all of which mean that MRCP data is poorly suited to the direct application of standard computer assisted diagnosis (CAD) procedures.

The aim of this work is to facilitate the effective utilisation of MRCP for the automated and semiautomated screening and assessment of the pancreato-biliary system, through the application of novel and well-focused image preprocessing techniques. The primary goal is to present a unified pre-segmentation data filter-

ing pipeline designed to allow the subsequent robust operation of segmentation and CAD techniques to the analysis of MRCP data in the visualisation, identification, and flagging of features of potential interest to the examining radiologist. The most immediate and important aspect of that task in the




Fig. 1.1: The pancreato-biliary system

context of this thesis lies in the rapid and consistent assessment of the gall bladder and common bile duct, and in the reliable recognition and localisation of stones located at these two sites. Also of interest is the identification of stenoses or narrowing of the ducts within the tree, which can be indicative of other pathologies¹. Bringing to the attention of the examining radiologist potential locations of such features through data flagging and flexible, high quality visualisations is the ultimate goal of CAD in MRCP.

Fig. 1.2 shows slices from two coronally acquired, volumetric MRCP examinations. In Fig. 1.2a a number of stones are visible within the enlarged gall bladder while in Fig. 1.2b a common bile duct stone can be seen. In both cases the information that can be built up from the preceding and succeeding slices clarifies the situation further, enabling the radiologist to form a detailed view. The optimal utilisation of this information through a full 3-D reconstruction of the tree is a key goal of this work.

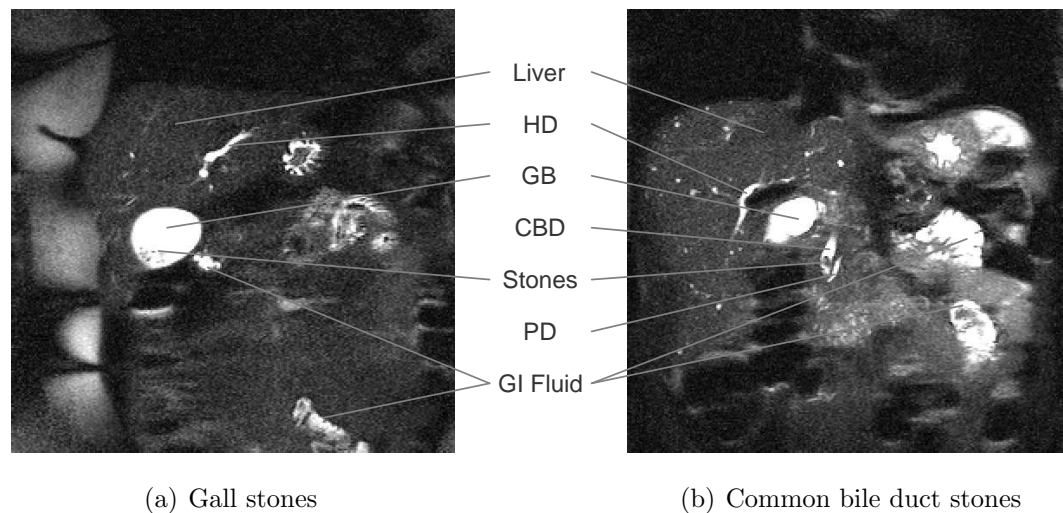


Fig. 1.2: Depiction of stones in the gall bladder and common bile duct

As the volume of data generated by MRCP and related imaging procedures continues to grow, it is essential that reliable automated screening techniques be developed in order to assist the radiologist in the thorough and timely assessment of these image series. In consort with the effect of advancing scanner technology, evolving protocol enhancements such as SMASH (Griswold et al., 1999) and SENSE (Pruessmann et al., 1999) make possible ever more detailed imaging of this anatomical region, but in so doing also generate greater vol-

¹Pathology — A departure or deviation from a normal condition.

umes of data to be reviewed and assessed by the radiologist. As this trend continues, CAD becomes ever more important in this area of medical imaging, as it has already become in such areas as CT Colonography (CTC) (Johnson and Dachman, 2000) and Whole Body MRI (WB-MRI) (Barkhausen et al., 2001), where the number and size of images in a typical examination tends to be exceptionally large. While these advances in MRCP acquisition inevitable improve the levels of detail resolvable, image noise remains an issue and the clarity achieved in MRCP is set to remain significantly below that observed in other areas of MRI usage due to the underlying processes involved and the inherent nonrigid motions ever present in this region of the body, which together limit useful scanning times and introduce image noise and motion artifacts. Much emphasis has thus been placed on addressing the issues mentioned above, in order to develop a viable set of image processing techniques towards the goal of reliable automated CAD in MRCP.

1.1 Background and Motivation

In order to provide a context for the material that follows, a short introduction is presented to the basics of MRCP and the factors and considerations that led to the initiation of this research project. The following discussion represents a brief outline of the three main classes of MRCP protocol addressed and utilised in this work. The flexibility provided and the restrictions imposed by the MRI scanner (Webb, 1988), and the specifics of MRCP acquisition protocol design and utilisation (Sai and Ariyama, 2000) are beyond the scope of this investigation.

There can be many variations within each of the classes of protocol described, and new acquisition protocols for MRCP examinations continue to be investigated and tested. The development of new MRI protocols is a large and active field, which also falls outside the scope of this work. There is much published literature in this area, see for example Boraschi et al. (1999a), Hundt et al. (2002). Most MRCP protocols, however, continue to utilise the same basic approach, designed to highlight stationary fluids in the scanned volume.

The above topics represent major areas of ongoing research in their own rights. The focus of this thesis, however, is with the most effective usage of the scanned

data once it has been generated. From this initial data the task is to apply image processing techniques in order to assist the radiologist in extracting the maximum amount of useful information from the acquired studies.

1.1.1 MRCP Protocols

Magnetic Resonance Cholangiopancreatography (MRCP) refers to the use of MR imaging techniques in order to image the biliary tree and the pancreatic duct in the area in and around the liver and pancreas. Quite a number of different protocols have been utilised in this regard, each with its own particular characteristics, merits and drawbacks, (Boraschi et al., 1999a, Sai and Ariyama, 2000, Tang et al., 2001, Hundt et al., 2002). Most existing techniques are based on acquisition protocols that operate by highlighting stationary fluid in the scanned volume. The data that has been considered in this project has been acquired utilising three classes of protocol, referred to as RARE, HASTE, and TRUFI. Of these three, HASTE has been primarily used in the work that has been conducted to date, as it provides the most direct route to a 3-D reconstruction of the pancreato-biliary system. Brief descriptions of the kinds of data yielded by each of these three types of acquisition protocol follow.

RARE

Rapid Acquisition by Relaxation Enhancement. This technique is used to acquire single slice, thick slab images of the biliary tree, as illustrated in Fig. 1.3. The biliary tree is clearly visible in the upper left quadrant of this image. The common bile duct is easily identified descending from the tree towards the centre of the image. The gall bladder can be seen as a large high intensity region located underneath the tree, extending to the left of the common bile duct.



Fig. 1.3: A typical RARE image

The pancreatic duct is not clearly visualised in this case. The bright signal

regions to the right of the image are due mainly to gastrointestinal fluids. Unwanted signals of this type can often overlap and interfere with the signals of interest coming from the pancreato-biliary system.

Typically an area of the anatomy surrounding the liver, with a volume in the region of $400\text{mm} \times 400\text{mm} \times 80\text{mm}$ is acquired as a single image, effectively yielding a raysum² projectional type view of an 80mm thick slab around the subject's liver. This type of acquisition tends to give a good overall view of the region of interest and is often used as a guide for more detailed examination of subsequent HASTE and TRUFI datasets.

RARE images provide a similar type of view of the subject area to that achieved through the use of the ERCP technique, which will be described later. The quality of the results achieved can, however, vary greatly from one study to the next (see Fig. 1.4) and depends strongly on there being significant amounts of bile and pancreatic juices present in the system at the time of the examination. This is a requirement for good results with all types of MRCP acquisition as it is the stationary fluid in the system that generates the signal. Subjects are usually asked to fast for several hours prior to examination in order to allow bile and pancreatic juice to collect.




Fig. 1.4: The degree of visualisation can vary considerably from one RARE examination to another. In (a) a faint common bile duct (CBD) is all but lost in high intensity gastrointestinal (GI) signal, and the rest of the tree is absent. In (b) the tree is visible to the first level of the hepatic duct (HD), while in (c) the gall bladder (GB) and pancreatic duct (PD) are also present in the image, but again little else of the tree can be seen.

²A raysum projection is formed as a set of parallel line integrals through the 3-D region of interest, each line (or ray) corresponding to a point in the final image. This amounts to a parallel projection of the 3-D region onto a 2-D plane.

These RARE images are not utilised directly in the work described in this thesis, as it is the 3-D reconstruction of the biliary tree which is being addressed, and RARE images are single slice projections of the volume of interest onto a plane. Some 3-D reconstruction can be performed from this type of data if a number of views are acquired, each taken from a different direction. In this case the 3-D layout of the biliary tree can be interpolated using a back projection type of approach (Ko et al., 1995, Lin et al., 1995). The shape of the ducts is then estimated using an elliptical cross-section model for the duct geometry, and from this a 3-D reconstruction of the biliary tree is achieved. Views of the tree structure, both external and virtual endoscopic can then be generated and from these views some assessment of the tree can be made. This technique is inherently of limited utility due to the nature of the estimations which have to be made in the reconstruction process, and the uncertainties which these estimations introduce.

HASTE

Half-Fourier Acquisition Single-Shot Turbo Spin-Echo. Fig. 1.5 shows two consecutive slices from midway through an axial HASTE dataset. The liver boundary is clearly visible in the left half of these images, with the numerous small high intensity regions inside representing the multitude of branches of the biliary tree spreading throughout the body of the liver. The larger high intensity regions in this area represent the common bile duct as it exits the liver. The cortico-spinal fluid surrounding the spinal cord is also clearly visible at the bottom centre of the images, and the intestines and spleen can be seen to the right.

This type of acquisition represents the main source of data in the current work. It yields a stack of slices acquired contiguously giving a volumetric dataset ideally suited to the task of 3-D reconstruction, which is the primary goal of this work. The images are usually acquired in an approximately axial (Fig. 1.5) or coronal (Fig. 1.6) orientation. That is to say slices may be acquired through the body with successive slices going from the feet towards the head (axial), or with successive slices running from the chest towards the back (coronal). Sagittal acquisitions (with slices being acquired running from the right side of the body to the left) are also possible but are rarely used.




Fig. 1.5: Two consecutive images from an axial HASTE sequence

In fact most acquisitions are made slightly off one of these orthogonal planes, oriented so as to achieve the best possible coverage of the region of interest, ensuring that all of the major elements of the pancreato-biliary system are captured. This is a particular issue due to the constraints that exist as to the amount of data that can be acquired in a single series. Tradeoffs with resolution and signal to noise ratio (SNR) are required, so it is important to maximise the amount of useful data acquired.




Fig. 1.6: Three consecutive images from a coronal HASTE sequence

HASTE datasets typically comprise thirteen to fifteen slices. Pixels are square in-slice and can typically range from about 1.2 to 1.6 millimetres in each direction. Slice thickness is usually around three to four millimetres. Due to the limitations on the coverage achievable in one acquisition, multiple volumes are often acquired in order to cover the totality of the region of interest. Volume merging is, however, difficult due to nonrigid organ motions in this region.

TRUFI (TrueFISP)

True Fast Imaging with Steady-State Precession. This protocol is not used as routinely as the previous two. It does provide excellent delineation of many boundaries of interest. However, it has one major drawback compared with the previous methods when addressing the task of automated or semiautomated analysis of the biliary tree using MRCP. This protocol highlights the flowing blood in equal measure with the stationary bile and pancreatic juices, and as such it is often difficult to reliably identify the path and condition of the ducts in the biliary tree because they run very close to the blood vessels, especially where they enter the liver.




Fig. 1.7: Two images from a TRUFI sequence

An example of TRUFI data can be seen in Fig. 1.7. As can be observed when compared to the axial HASTE images of Fig. 1.5, soft tissue boundaries in particular are far better delineated than in HASTE data. However, the high intensity signal within the liver is not now due solely to the bile present, but also to the substantial blood supply that the liver receives. This makes the reconstruction and analysis tasks far more difficult in this class of data, and for this reason the main focus in this work has been on HASTE MRCP series.

1.1.2 What MRCP is Used For

MRCP was primarily developed as a replacement for the far more invasive examination technique called ERCP or Endoscopic Retrograde Cholangiopancreatography. In ERCP an endoscope is passed down the oesophagus, through

the stomach, and into the small intestine. The endoscope is directed to the ampulla of Vater (see Fig. 1.1) where the pancreato-biliary system feeds into the intestine. A contrast agent is then injected into the tree and the subject undergoes an x-ray examination, which highlights the contrast agent now dispersed throughout the biliary system. In this way the biliary tree is imaged, but only to the extent to which it was successfully penetrated by the contrast agent. Therefore if the common bile duct is obstructed, for instance by stones that have migrated out of the gall bladder, then the tree may not be visualised above these obstructions. ERCP has a number of other undesirable features associated with it. These include the invasive nature of the procedure and the need for the use of ionising radiation. Insertion of the endoscope is uncomfortable for the subject and can result in tearing or perforation of the regions through which the endoscope must pass. This can be a very serious complication, which can in extreme cases result in patient mortality. The use of an x-ray examination and the associated exposure to ionising radiation is an additional undesirable necessity of this type of procedure and as such also counts against its use.

A typical ERCP examination is shown in Fig. 1.8. The main sections of the biliary tree are well delineated. The clear visibility of the vertebra of the spinal column and of the ribs is indicative of the nature of this type of examination, which utilises x-rays. By comparison, the RARE image in Fig. 1.3, which visualises a similar region shows no trace of the bones present in the field of view (although the cortico-spinal



Fig. 1.8: An ERCP examination

fluid is faintly visible descending in the lower middle section of the image). This characteristic along with the superior sharpness and SNR achieved in ERCP examinations easily differentiates between the two types of image. ERCP does have the advantage that once the endoscope is in place it can sometimes be used to remove stones that have been identified in the examination. It is sometimes the case that an MRCP exam will be followed by the conduction of an ERCP procedure to this end. Thus it is within this context that the ongoing developments in the quality and reliability of MRCP for the diagnosis of

problems in the pancreato-biliary system proceed. These advances mean that MRCP continues to increase its challenge to ERCP as the examination of first resort where such conditions are indicated.

1.1.3 How MRCP is Currently Utilised

MRCP is being increasingly used in examinations of the biliary tree and pancreatic duct. The protocols described in Section 1.1.1, along with others, are used to acquire a set of image series, which collectively form a study. Single slice, thick slab RARE images give a good overview of the tree while HASTE and TRUFI examinations provide a more detailed 3-D view. In current practice, little or no preprocessing of the data is performed. The radiologist, working at a review station, examines the collected series, browsing through the slices in order to come to an overall picture of the state of the subject's pancreato-biliary system.

Typically, regions and features of interest are identified in one series and the corresponding locations are pinpointed and examined in other series covering the same area in order to build up a more comprehensive picture of what is demonstrated in the scans, with a greater level of confidence in the conclusions drawn. In this way an assessment is made as to the state of the subject's pancreato-biliary system, and a diagnosis and course of action determined.

1.1.4 MRCP with CAD

The underlying goals of the research presented in this thesis involve the application of adaptive image processing techniques to the task of image enhancement and noise reduction. This aims to facilitate robust segmentation and to provide improved visualisation tools and visual cues in the review and assessment of MRCP data, and to render the data more suitable for the subsequent application of automated computer assisted diagnosis (CAD) techniques. The use of MRCP as a diagnostic tool is growing rapidly and interest in the area is expanding. However, while the technique has demonstrated the potential to provide a viable alternative to the more invasive diagnostic procedures of ERCP, its utility will ultimately be governed by the resolving power it can be shown to exhibit.

The data obtained in MRCP is generally noisy and of a relatively low resolution, especially when considering the relatively large inter-slice distance achieved for multi-slice datasets. Angio-style vessel tracking approaches quickly fail when the duct diameters approach the limits of the image resolution achieved, as is the case in the smaller ducts visualised in the majority of MRCP series. These properties render the reliable evaluation and interpretation of MRCP data a difficult task.

By suppressing extraneous signal from gastrointestinal and other stationary fluids in the scanning region, and by enhancing and highlighting signal due to the bile and pancreatic juices, the aim is to present the radiologist with images that are more easily, accurately, and consistently interpreted and assessed. In addition, by facilitating simple segmentation of the biliary tree and pancreatic duct in the 3-D data, more informative and more intuitively interpreted 3-D rendered views of the available data can be achieved. This will allow the radiologist to build up a more accurate and detailed picture of the condition of the pancreato-biliary system under examination. These improvements in the presented data will also facilitate the application of CAD based procedures, further assisting the radiologist by flagging regions and features of potential interest in the large volumes of data acquired across multiple series, which are typically generated in an MRCP study.

1.1.5 Literature Review

Much published literature exists addressing topics in MRCP and related areas, providing a large body of background reference and research material covering both the medical and image processing aspects of this work. The current section highlights some of the main publications relevant to the subject area addressed in this project. These publications are listed under four subheadings covering respectively, the clinical, and image processing aspects of MRCP, general medical imaging, and a broader collection of significant image processing material. The main contributions of these publications, and their primary significance within the field, are highlighted, providing a broader context within which the work presented in this thesis can be viewed.

Clinical MRCP literature

The various aspects of MRCP from the radiologist's perspective are covered by the material presented in a number of reference volumes that have been published on the subject (Pavone and Passariello, 1997, Hoe et al., 1998, Sai and Ariyama, 2000). These reference books provide an excellent overview of how MRCP examinations are utilised, and what kind and degree of clinical information they can yield. Review of the material provided in these texts also highlights in particular the levels of skill and training required by the examining radiologist in order to accurately interpret MRCP images, and as such illustrates the high degree of difficulty involved in attempting to automate the analysis of such data.

In addition to the above reference texts, more focused examinations of the evolving role of MRCP appear in numerous published research papers such as Guibaud et al. (1995), Reinhold and Bret (1996), Larena et al. (1998), Takehara (1999). These papers provide a detailed review and assessment of the performance of MRCP in the accurate and consistent visualisation and differentiation of various structures and pathologies of interest within the pancreaticobiliary system. They provide critical comparisons between MRCP and other competing types of examination such as ERCP, highlighting the strengths and weaknesses of existing MRCP protocols. They assess the suitability of MRCP to various diagnostic tasks, and the potential roles which these evolving imaging techniques might play within a broader clinical context.

There has also been a great deal of published work addressing the development and assessment of new MRCP protocols (Boraschi et al., 1999*a*, Papanikolaou et al., 1999, Tang et al., 2001, Hundt et al., 2002). More efficient data acquisition techniques for MR imaging have been proposed (Griswold et al., 1999, Pruessmann et al., 1999), along with examinations of the effectiveness and utility of various MRCP contrast agents (Mariani, 2001, Dalal et al., 2004), and reports on the conduct of clinical trials into the utility and performance of MRCP (Boraschi et al., 1999*b*, Williams et al., 2001, Kondo et al., 2005). Taken together these publications provide the clinical context within which the current work has been conducted, and as such have assisted in identifying the potential for the application of automated image analysis and CAD techniques in the assessment of MRCP data.

Image processing in MRCP

Less has been published on the application of image processing and analysis techniques to the presentation and assessment of MRCP. Some attention has been focused on the tasks of biliary tree reconstruction and visualisation. In Ko et al. (1995) and Lin et al. (1995), the authors propose a 3-D reconstruction technique for the biliary tree based on point correspondences and a branch skeletonisation procedure in two mutually orthogonal views, and they present useful 3-D renderings of the reconstructed trees. The views lack structural detail due to the estimations of the reconstruction process but provide a good 3-D overview of the biliary tree.

Chen and Wang (2004) illustrate a technique for segmenting the biliary tree from volumetric MRCP data based on a region growing and centreline tracking approach. The low resolution of the data results in a poor representation of the tree, especially in the inter-slice direction, and the approach fails to retain finer, less distinct portions of the tree, which are obscured due to noise and a lack of resolution in the data. The results do, however, provide an informative 3-D overview of the layout and general condition of that portion of the tree which is segmented.

A number of studies have reported on the utility of volume rendered review of MRCP data as an adjunct to planar review. Cesari et al. (2000) suggest a ray-sum reconstruction algorithm as being superior to the more familiar maximum intensity projection (MIP) rendering approach. The raysum algorithm better represents the presence of stones within a duct, which standard MIP renderings tend to obscure. In Neri et al. (2000), a study using shaded surface display (SSD) volume rendered MRCP is presented, concluding that the technique, while cumbersome to use, offers the potential for informative visualisations to be achieved.

Kondo et al. (2001) present a study comparing the results achieved for diagnoses performed on images generated from SSD and MIP renderings, and planar review data. The authors highlight the limitations of the widely used MIP rendering technique in adequately visualising the biliary tree and conclude that superior results can be achieved using more advanced volume rendering approaches such as SSD.

The use of virtual endoscopic renderings in MRCP has been investigated by a number of authors. In Dubno et al. (1998) the authors provide a brief review of the technique, and an assessment of the potential for virtual MR cholangiography, illustrating the intraluminal depiction of the common bile duct, demonstrating the ability to visualise stones and cavities in the duct.

In Neri et al. (1999*a*) and Neri et al. (1999*b*) the authors use a surface rendered approach to the task of generating virtual endoscopic views of the pancreato-biliary tract. They assessed the performance of the technique on data from 120 subjects and found it useful in depicting the internal anatomy of the biliary tract and in identifying changes due to pathological conditions. Prassopoulos et al. (2002) further demonstrate the application of virtual endoscopic assessment of the common bile duct, based on alternative MRCP protocols, and again conclude that the technique has significant potential.

Related studies examining the use of virtual endoscopic techniques for the assessment of CT cholangiography data (Prassopoulos et al., 1998, Koito et al., 2001) show similar levels of visualisation of the anatomy and pathologies of the pancreato-biliary tract present in that data. Virtual endoscopy in a more general context is a widely examined subject (Summers, 2000, Deschamps and Cohen, 2001, Oto, 2002, Fetita et al., 2004, Haigron et al., 2004). However, these studies in general address the application of virtual endoscopic techniques to data with significantly higher resolution in the spaces and cavities under examination than that which is achieved within the ducts of the biliary tree in MRCP studies. As such virtual endoscopic in MRCP presents particular challenges and requires the best possible quality and the highest possible resolution of input data in order to achieve useful results.

Medical image processing

Medical image analysis as a whole is a vast research area covering the entire spectrum of acquisition modalities and anatomical regions, and as such it is to be expected that much work which can be usefully applied to the processing of MRCP data is to be found within this wider area of investigation. The present work draws on a broad body of published material, addressing in particular a number of topics relevant to the specific problems encountered in the processing and analysis of MRCP data.

Greyscale correction approaches for MRI data, such as those techniques referred to as *bias field correction* and *coil correction* address many of the issues relating to image intensity non-uniformities frequently observed in MRI data in general. These techniques most often focus on *intra*-image non-uniformities and are not on the whole directly applicable to the kind of *inter*-image non-uniformities typically observed in the MRCP data that is addressed in this work. They do, however, provide a useful starting point in working towards an approach to address these issues.

Vokurka et al. (1999) provide a comprehensive treatment of the topic of addressing both intra-image and inter-image intensity non-uniformity correction. A versatile data model is developed to describe the non-uniformity effects observed in MRI data and a pair of iterative correction schemes is proposed. The main focus is on the intra-slice case and there is no examination of the effect the proposed inter-slice correction procedure (which applies a set of slice-wise correction factors at each iteration) has on the individual image histograms. This consideration is an important element of the correction scheme presented in Chapter 2, where considerable emphasis is placed on preservation of the image histograms, so as to facilitate later histogram-based processing. The application and effectiveness of the intra-slice correction procedure of Vokurka et al. (1999) is further examined in Vokurka et al. (2001), with a case study that looks at the correction of non-uniformities in MRI examinations of the eye.

Alternative approaches to address the problems surrounding non-uniformity correction can be found in papers such as Newman et al. (2002), where an adaptive histogramming technique is used to address slice-to-slice intensity variations, and Lai and Fang (2003), in which an acquisition-time solution is proposed, where a second lower resolution image, simultaneously acquired using an additional body coil, is utilised in order to guide the correction process.

Approaches to vessel tracking and segmentation are of particular interest in informing the direction of our work, and although most existing techniques address higher resolution data where the modelling of the vessels can be approached more straightforwardly, the general techniques described have helped to illuminate some of the problems that must be considered. These considerations in particular helped in formulating the development of the hybrid reconstruction procedure described in Chapter 4.

Various model-based vessel tracking strategies are presented by numerous authors (see Frangi et al., 1999, Wang and Bhalerao, 2002, for instance), which use information regarding expected vessel shape, layout and connectedness in order to segment the vasculature³ in various parts of the anatomy including the head and the heart (Flasque et al., 2001, Lorigo et al., 2001), the retina (Farid and Murtagh, 2001, Mohamed and Auda, 2002), the liver (Selle et al., 2002), and the limbs (Kanitsar et al., 2001). Chen and Molloy (2002) present a general purpose 2-D method for segmenting treelike structures by tracking valley courses in the image, and Canero and Radeva (2003) illustrate a vessel enhancement technique for 2-D images designed to preserve tubular structures in the data.

In addition to the topics covered in the paragraphs above, a number of subjects of more general interest should be mentioned as they have influenced the formulation of the overall approach developed, highlighting various other considerations that must be borne in mind in the processing and analysis tasks that are to be addressed.

Various anatomical segmentation techniques, and classification (Ashburner and Friston, 2000) and visualisation (Parker et al., 2000, Preim et al., 2000) procedures in medical imaging impinge on the processing tasks that are to be addressed. They are relevant either directly in terms of the image processing approaches that must be developed, or indirectly as elements of subsequent computer assisted diagnosis (CAD) procedures, consideration of which informs the more immediate goals of the preprocessing approaches that are developed and presented in this work. In addition to the various vessel tracking methods mentioned above, which are specific to branching tubular structures, many more general approaches are encountered, which address the segmentation of more compact structures such as the brain (Sijbers et al., 1997, Thacker and Jackson, 2001), heart (Frangi et al., 2001), or liver (Agrafiotis et al., 2001).

Data interpolation is of particular importance when working with volumes that are both of low resolution and non-isometric in their voxel dimensions, as is the case with the MRCP data considered here. Various approaches specific to the interpolation (Grevera and Udupa, 1998, Thacker et al., 1999) and registration (Hajnal et al., 1995) of MRI data are investigated and assessed

³Vasculature — Arrangement of blood vessels in the body or in an organ or a body part.

in the literature, including a zero-filled k -space⁴ interpolation approach by Du et al. (1994) that specifically addresses the enhancement of contrast and continuity in vessels after interpolation, and an iterative approach to k -space resampling (Pirsiavash et al., 2005) based on an alternating series of data refinement steps performed in k -space and image space.

Non-medical image processing

In addition to the medically-oriented material addressed above, a whole range of literature in the general field of image and signal processing provides the necessary foundations for much of the pre-segmentation filtering work which is addressed in the body of this thesis. These more general topics include various familiar techniques for gradient calculation and edge detection such as those presented in Frei and Chen (1977), Canny (1986), Sobel (1990). The robust identification of weak boundaries in noisy data is a key concern when considering potential approaches to the segmentation of the biliary tree in MRCP. Methods for edge enhancement (Greenspan et al., 2000) and line extraction (van der Heijden, 1995, Bigand et al., 2001) in image data can provide a useful starting point for the development of effective 3-D surface or boundary detection procedures.

Another image processing task that is of particular importance in this work is that of data smoothing and noise reduction. Many approaches to this topic appear in the literature, ranging from the simplest averaging and median filters (Gonzalez and Woods, 1992) through mathematical morphology (Serra, 1982, Soille, 1999), mean shift (Dominguez et al., 2003), and more involved spatial and frequency domain filtering schemes (Greenspan et al., 2000, Whelan and Molloy, 2000). Numerous adaptive approaches based on wavelets (Jung and Scharcanski, 2004), tangential smoothing (Bromiley et al., 2002), and variational methods (Schnorr, 1999) have all received attention.

One approach dominant in recent literature is that of diffusion filtering. Based on the mathematics of diffusion (Crank, 1975), it seeks to reduce noise within regions while preserving semantically important features such as regional boundaries by modelling the smoothing applied as a nonlinear diffusion process, with

⁴MR images are acquired as data in k -space, which is a frequency domain related to normal image space through the familiar Fourier transform pair.

the diffusivity being a function of local structure observed in the image data. Starting with Perona and Malik (1990) who described the original nonlinear diffusion filter for data smoothing, the technique has evolved with notable contributions to be found particularly in ter Haar Romeny (1994), Weickert et al. (1998), and Weickert (1999). The paper by Gerig et al. (1992) is notable in that it examines the use of diffusion filtering in the smoothing of MRI data in particular. Many other applications and variations have also been reported (Acton, 1998, Black et al., 1998, Sijbers et al., 1999, Krissian, 2002, Suri et al., 2002) addressing a variety of approaches and data smoothing tasks. All of this work forms an important backdrop for the adaptive filtering techniques described in Chapter 3.

Mathematical morphology in particular provides a number of useful tools for the purposes of this work. The theory and application of its techniques are widely examined, starting with the original work of Serra (1982) and including many important contributions from other authors addressing numerous topics. These include everything from the fundamental erosion and dilation operations (van Vliet and Verwer, 1988, Ji et al., 1989, Sivakumar et al., 2000), and the manipulation and decomposition of structuring elements (van Droogenbroeck and Talbot, 1996, Park and Yoo, 2001), to the implementation and application of much higher level morphological techniques.

Two methods in particular should be mentioned. Reconstruction by dilation (Vincent, 1993, Salembier and Serra, 1995, Soille, 2004), which is useful in the suppression of non-relevant structures in the data, and the widely investigated watershed segmentation procedure (Vincent and Soille, 1991, Beucher and Meyer, 1993, de Smet and de Vleeschauwer, 1997, Felkel et al., 2001, Lapeer et al., 2002, Nguyen et al., 2003). This latter technique offers an effective and robust approach to segmentation of the biliary tree once the issues of image noise and regional homogeneity have been successfully addressed.

An excellent introduction to the whole area of mathematical morphology is provided by Soille (1999), and the numerous investigated areas of application (Salembier et al., 1996, Araujo et al., 2001, Bueno et al., 2001, Angulo and Serra, 2003) provide further examples of the usefulness and versatility of the tools provided in the field.

Once a region of interest has been successfully defined, there are two basic routes to the 3-D visualisation of the structures in question. Either direct volume rendering techniques can be applied (Lacroute and Levoy, 1994, Gobbi and Peters, 2003), or a surface extraction procedure can be performed followed by the application of a surface rendering approach (Foley et al., 1993). Volume rendering techniques utilise all the information present in the original data but tend to be computationally expensive and thus can be slow and cumbersome to use. Surface renderings are generally fast, but discard much of the original data and can thus lack the detail of volume rendered views. In either case parallel or perspective projections can be applied in order to generate external or virtual endoscopic views respectively.

In the case of surface rendering approaches, surface extraction procedures enable a concise representation of the structures of interest to be constructed. Techniques have been proposed for the generation of a polyhedral mesh representation of the surface of an object from various input data including arbitrary point clouds (Boissonnat, 1984, Faugeras et al., 1984), stacked cross sectional contours (Boissonnat, 1988), and segmented voxel data. Approaches to this last case include the spider-web algorithm (Karron, 1992), the now ubiquitous marching cubes algorithm, proposed by Lorensen and Cline (1987) and since modified and enhanced by various authors (Delibasis et al., 2001, Rajon and Bolch, 2003), and the more recent growing cube algorithm of Lee and Lin (2001).

All of these topics were considered to a greater or lesser extent during the course of this research, and have influenced the form of the solutions that have been developed to address the particular problems encountered in the pre-segmentation processing of MRCP data towards effective computer assisted diagnosis (CAD) in the pancreato-biliary system.

1.2 Contributions

In assessing the research conducted over the course of this project, the most important aspects of this work have been identified, in the context of MRCP image processing for biliary tree pre-segmentation filtering. The body of work highlighted below in Section 1.2.1 represents the core of the research effort

presented in this thesis. Related work that was undertaken during the period, but that has a less direct bearing on the primary focus of this report is presented as subsidiary contributions in Section 1.2.2 and is expanded upon as appropriate in the appendices to this thesis.

The full scope of the work outlined in both of these sections can also be observed in the collection of publications that have stemmed from this project. Full references for these publications are given in Chapter 6, and all are available as pdf documents, along with presentations, posters, and other supporting materials, on the publications pages accessible at www.eeng.dcu.ie/~robinsok. All publications cited in Sections 1.2.1 and 1.2.2 below are taken from this list and represent the substantive contributions stemming from these aspects of the presented work.

1.2.1 Primary Contributions

In focusing on achieving an effective and consistent route to segmentation of the biliary tree from MRCP, the main aim has been to arrive at a data preprocessing scheme designed to address the particular problems relating to noise, resolution, and consistency as observed in MRCP data, so as to facilitate more robust and representative volumetric segmentation and computer assisted diagnosis (CAD) results in subsequent analysis of the data. In this context the major contributions documented in this thesis form the various steps in a multi-phase image preprocessing strategy for narrow, branching structures in noisy, low resolution volume data. Each of the topics below is addressed in the following chapters, forming the main body of this thesis.

1. Due to the characteristics of the MRI acquisition protocols utilised in the collection of the MRCP data, intensity non-uniformities often appear, resulting in the first several coronal slices in the data volume being significantly brighter than subsequent slices. Thus a data preparation procedure was developed in the form of a histogram-based inter-image intensity non-uniformity correction scheme (Robinson et al., 2004, 2005a) in order to minimise the effect of these greyscale inhomogeneities through a nonlinear histogram matching process that operates by aligning key features across the sequence of histograms corresponding to each slice in the dataset.

2. In the next phase attention was focused on the goal of achieving a significant reduction in the considerable noise present in the data, and to this end an investigation and comparison of many adaptive smoothing approaches was conducted (Robinson et al., 2002*a*, Lynch et al., 2004, Ghita et al., 2005*a*), and a novel 3-D adaptive filtering scheme was developed, based on the Gaussian smoothing model (Robinson, 2004). This approach has proven effective in attenuating signal noise while at the same time preserving well the semantically important discontinuities that are present in the volume data.
3. Following this a morphological reconstruction procedure was developed in order to suppress the signal originating from neighbouring structures in the scanned volume, while preserving the signal due to the tree structure that is to be segmented. This goal of retaining the narrow branch features during the morphological processing is addressed by the hybrid reconstruction approach detailed in Robinson and Whelan (2004*b*) where a generalisation of the traditional reconstruction by dilation procedure familiar from the greyscale mathematical morphology is described. This hybrid reconstruction approach allows the degree of greylevel connectivity required to be specified during the reconstruction process.
4. Through this work on morphological approaches to reconstruction, an optimal algorithm for reconstruction by dilation was developed. The generalisation of this *directed filtering* algorithmic pattern can be applied to a class of related image processing procedures including the grassfire distance transform and the watershed segmentation algorithm. The specific application of this new and efficient algorithm to morphological reconstruction by dilation, called the downhill filter, has been published in Robinson and Whelan (2004*a*).

1.2.2 Subsidiary Contributions

During the course of the research programme outlined above, a number of important topics had to be addressed, subsidiary to the main thrust of this effort, but nonetheless important in themselves and in the broader framework of the programme of research undertaken. Two of these subsidiary topics merit particular mention in this thesis despite not fitting well within the main body

of work being presented. Appendices addressing these two topics, as described below, appear at the end of this report.

1. A prospective study was conducted into the use of whole body MRI in the assessment of body fat level and distribution (Brennan et al., 2005). These investigations used many of the same techniques described in this thesis, achieving superior segmentations as a result, and thus demonstrating the wider applicability of the pre-segmentation approach described here. This work also led to some useful results in volumetric reconstruction (Robinson et al., 2004), and produced a number of more focused results and findings in its own right (Whelan et al., 2004, Robinson et al., 2005*b*, Ilea et al., 2005).
2. In order to encapsulate the tools and algorithms developed, the NeatMRI environment was constructed, a software library and a set of tools providing easy access to all the techniques and procedures investigated and implemented during this work. This software framework, as outlined in Appendix B, is fully documented in its own ‘Programmers Reference Manual’, which accompanies the toolkit. The NeatMRI environment represents a substantial and evolving image processing library for fast prototyping and robust development of powerful image analysis and visualisation systems.

1.3 Thesis Outline

The chapters that follow this introduction present the design, development, characterisation and testing of a series of data preprocessing steps that together represent an effective means for the preparation of MRCP data for subsequent segmentation, visualisation, and analysis of the biliary tree from an MRCP volume. Each chapter examines an important topic representing a step in the overall data preprocessing pipeline. Together they present an effective route towards the robust application of standard automated CAD procedures in MRCP.

Chapter 2 addresses a histogram-matching procedure developed to overcome intensity non-uniformities observed in the coronal HASTE data that is being

used. The greylevel distributions within a series of slices are matched to compensate for inter-slice greylevel shift. A novel technique that preserves the integrity of the data histogram is described, ensuring that spikes and voids are not introduced into the histogram during the scaling and matching process. This is important so that subsequent data processing and analysis steps can utilise the resulting volume histogram robustly.

Chapter 3 concentrates on noise reduction through the application of a flexible new gradient-weighted adaptive Gaussian smoothing technique. Noise is suppressed while semantically important boundaries are preserved through a process of directed filtering, where stronger gradients result in more highly directional smoothing, while weaker gradients allow more isotropic smoothing to occur. This leads to a nonlinear anisotropic form of data filtering where edges do not get blurred while noise in the body of regions is effectively eliminated.

Next the hybrid reconstruction procedure of Chapter 4 operates so as to isolate the biliary tree by attenuating the signal from neighbouring high intensity regions, while maintaining the signal level throughout the fine branching structure of interest. This is achieved using a morphological approach based on conditional and geodesic dilations. This step helps to better differentiate between signal from relevant and non-relevant structures in the volume. Through this work on morphological approaches to reconstruction, an optimal algorithm for a class of directed filtering problems was also developed. The specific application of this to morphological reconstruction by dilation, called the downhill filter, has been documented in Robinson and Whelan (2004a). The generalisation of this algorithmic pattern, which was called *directed filtering*, is also presented in this chapter.

In Chapter 5 an implementation overview is presented illustrating the application of the techniques described in the previous three chapters. Examples from both MRCP and more general MR imaging sources are given, demonstrating the wider utility of the techniques. Versatile, high resolution reslicing and rendering procedures are applied to the processed volumes, demonstrating the application of advanced visualisation techniques in the computer-assisted assessment and diagnosis of the processed data.

The main body of the thesis is closed in Chapter 6, where a summary of the research work conducted and a review of the results achieved is presented, and

what further work remains to be done in order to carry forward the goals of the project is discussed. An assessment of the results achieved through the application of these procedures is presented and progress towards the goal of robust and consistent isolation of the biliary tree using non-complex segmentation procedures is examined. A full list of the publications stemming from this work is also provided at this point.

Finally, a number of appendices covering subsidiary topics appear at the end of the thesis. During the course of the work, a prospective study was conducted into the use of whole body MRI in the assessment of body fat level and distribution (Brennan et al., 2005). These investigations used many of the same techniques described in this thesis, they led to some useful results in volumetric reconstruction (Robinson et al., 2004), and they also produced a number of more focused results and findings in their own right (Robinson et al., 2005*b*, Ilea et al., 2005, Whelan et al., 2004). The major aspects of this study are outlined in Appendix A.

An introduction to the programming library and environment NeatMRI is given in Appendix B, where the full functionality of the library is outlined, and its usage and dual interfaces through C and Java are illustrated. In order to encapsulate the tools and algorithms that have been developed, the NeatMRI environment was constructed. This is a software library and a set of tools providing easy access to all the techniques and procedures investigated and implemented during this work. This software framework is outlined in Appendix B and fully documented in its own ‘Programmers Reference Manual’, which accompanies the toolkit.

Chapter 2

Intensity Non-uniformity Correction

Histogram-based techniques are widely used in two distinct contexts. Firstly, the data histogram can be altered, globally or locally, in order to change the data's greylevel distribution in some way. This is often done so as to improve the visual appearance of a displayed image by increasing its contrast or focusing in on a particular grey range in the data. These tasks are typically performed using such techniques as histogram stretching, windowing, and global and local area equalisation (Gonzalez and Woods, 1992, Whelan and Molloy, 2000), and generally result in the introduction of spurious new extrema into the histogram through the merging or separation of histogram bins. Similar histogram-based greyscale homogenisation techniques have also been applied to the correction of intensity non-uniformities in the processing of MR images (Vokurka et al., 1999, Dauguet et al., 2004), and these procedures exhibit the same tendency to corrupt the resultant data histogram with spurious spikes and voids.

The second area of application of histogram-based techniques is in the context of data segmentation and classification tasks, where they are useful for such operations as threshold selection (Otsu, 1979, Dulyakarn et al., 1999, Robinson et al., 2005*b*) and automated segmentation procedures (Mangin et al., 1998, Newman et al., 2002). In this second class of applications, the shape of the histogram is examined in order to guide the processing being applied. If techniques of the first type have been used, that have altered the characteristics of the histogram, introducing new minima and maxima, prior to the appli-

cation of techniques from this second class, then these subsequent histogram analysis based procedures are likely to encounter difficulties. In this chapter we describe a method of the first type designed to accommodate the later application of this second type of procedure, by preserving the overall shape of the data histogram during the remapping process, avoiding the introduction of spurious new minima and maxima into the histogram.

We present a novel histogram-matching procedure that was specifically developed to correct for the inter-slice intensity non-uniformities typically observed in the coronal HASTE data with which we are working. In order to achieve this intensity non-uniformity correction effect, the greylevel distributions *within* the individual slices in a HASTE volume must be modified so as to compensate for inter-slice greylevel shift. To this end a histogram matching approach was developed that aligns corresponding grey ranges across the individual slices, while at the same time preserving the integrity of the overall volume histogram, ensuring that spurious features are not introduced during the matching procedure. This property of preserving the integrity of the data histogram is particularly important in order to ensure that subsequent processing steps can utilise the histogram from the resulting data volume in order to perform robust automated histogram-based analysis and threshold level selection operations.

Histogram corruption by spurious extrema

The aforementioned unwanted new extrema are commonly observed with the more traditional histogram scaling schemes typically employed in applications such as greyscale windowing for visualisation purposes as mentioned above. Such spike and void features do not adversely affect the visual characteristics of the data as they merely cause certain sets of pixels to change their greyvalues by a single greylevel relative to their intensity neighbours (i.e. those pixels close to them in intensity rather than space). As such, the visual consequences of such histogram spikes and voids are minimal. However, in our work their presence would constitute a significant difficulty, as the data histogram is used in later processing steps. By addressing this issue here we ensure that subsequent histogram-based calculations, especially in the automatic selection of threshold bands through histogram analysis, can be performed simply and robustly in the succeeding phases of the processing pipeline.

In Fig. 2.1 the first histogram is that of an unaltered image taken from one of our coronal HASTE datasets. The second is from the same slice after all its voxels have been multiplied by a value of 0.96, while in the third the voxels were multiplied by a value of 1.04. Since the voxel values are stored as integers, rounding occurs and as a result spikes and voids are created in the histogram, where two adjacent values in the original data are either both mapped to the same value in the result, or are mapped to two more widely separated values. Where this issue is encountered within an image processing environment, it is most commonly addressed by smoothing the compromised data histogram through a process of bin averaging, before any histogram based calculations such as peak detection are performed.

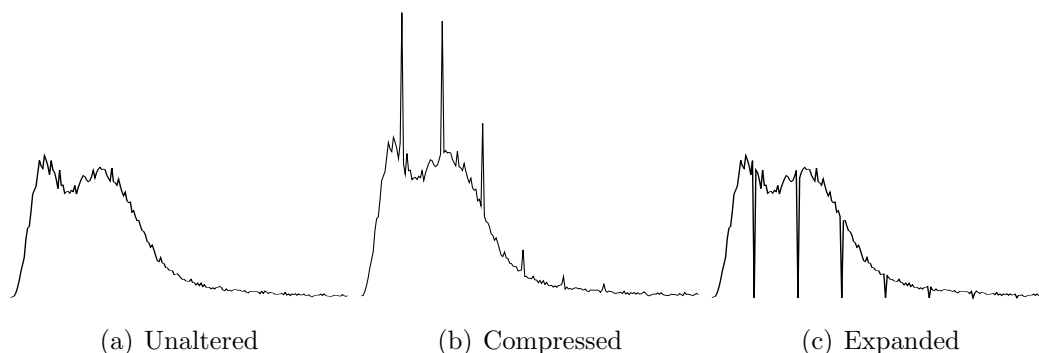


Fig. 2.1: A histogram compromised with spurious spikes and voids caused by remapping of the data into a new grey range. The original histogram is shown in (a), while (b) illustrates the effect of compressing the gray range thus introducing spikes, and (c) shows the effect of expanding the grey range, introducing voids.

While such histogram smoothing approaches seek to overcome the problem as best they can, it is difficult to completely eliminate the spurious features that have been introduced into the histograms in the scaling process. They can thus result in the misinterpretation of such spurious local minima or maxima found in the compromised data histogram. Our aim is to avoid these difficulties arising in the first place by performing a more elaborate redistribution of the voxels into the available bins in the rescaled histogram, thus preventing the introduction of artificial spikes and voids.

In the specific case of the coronal HASTE data that we address here, the histogram matching procedure that we have developed, operates in the following fashion. Firstly a separate intensity histogram is constructed for each slice in the volume. The characteristic peaks representing air and soft tissue are algo-

rhythmically identified in each case, along with the position of the intervening trough or local minimum. Thus fixed points on the individual histograms are first located and then aligned across the complete set of histograms, so as to achieve a matched greyscale distribution across all slices.

The question of identifying the appropriate histogram peaks, troughs, and characteristic fixed points in a robust and consistent fashion is addressed in the analysis that follows. This is in turn followed by an examination of how we scale the individual grey maps so as to preserve the integrity of the rescaled slice histograms, and thus that of the overall volume histogram after matching. This procedure avoids the introduction of artificial spikes and voids, which can be caused by histogram bins merging and separating respectively as illustrated in Fig. 2.1.

2.1 Types of Intensity Non-uniformity


Intensity non-uniformities in MRI data can be classified under three headings: intra-image, inter-image, and inter-volume. Inter-volume non-uniformities may be further subdivided into those that appear between adjacent sections in multi-section acquisitions, and those observed between the individual volumes forming a time series. All of the above kinds of non-uniformities occur due to the variability observed in the characteristics of the detected MR signal at the receiver coils in the MRI scanner, coming from different regions within the field of view, or originating from the same location at different times.

Although in this chapter we concentrate primarily on the inter-slice non-uniformities observed in coronal HASTE data, the approach that we present can easily be adapted to operate on other inter-slice and inter-volume non-uniformities. It has also been successfully applied to data in the multi-section inter-volume category mentioned above. This application is demonstrated in the whole body MRI study outlined in Appendix A, where non-uniformities observed between the individual sections of the scan must be corrected for. An equivalent treatment could also be applied to multi-image or multi-volume time series non-uniformities. Of the three types, only intra-image non-uniformities are not suited to correction by this technique. Here regional inhomogeneities exist within a single slice and are hence encoded within a single image his-

togram, requiring a different approach to their correction. The three categories of intensity non-uniformities are further outlined in the sections that follow.

Intra-Image Non-Uniformities

Intra-image intensity non-uniformities typically manifest as an intensity drop off observed as one moves away from the area close to where the receiver coils have been placed during acquisition. In Fig. 2.2, which depicts two views of the same slice taken from a volumetric shoulder scan, we can see how the intra-image non-uniformities make it difficult to clearly visualise all regions of the imaged field simultaneously.



(a) High intensity windowing

(b) Low intensity windowing

Fig. 2.2: Two views of the same image from an MRI scan of a shoulder, demonstrating intra-image intensity non-uniformity. Intensities fall off from top left to bottom right. In (a) intensity windowing focuses on the top left portion of the image leaving the bottom right obscured. In (b) the windowing focuses on the lower intensities in the bottom right, washing out the detail in the top left.⁵

The two different windowing levels that have been applied in (a) and (b) demonstrate how some regions are obscured (either excessively darkened or washed out) so that others can be brought into a better degree of visualisation by expanding their displayed grey range. In this case the acquired data has

⁵The MRI images shown here and in the rest of this thesis have had their greylevels inverted to improve their printed clarity, such that darker regions in the reproduced images correspond to higher voxel values in the original data, and vice versa.

intensities that lie in the range 1–3185, while typically a maximum of only 256 different greys can be displayed simultaneously, and indeed the human visual system can only differentiate between far fewer again (only about 30-50 greylevels at once). The large intensity range observed in the image in Fig. 2.2 corresponds as much to a global intensity gradient across the image as it does to a finer degree of local intensity resolution.

Axial HASTE MRCP series manifest the effect of these nonlinear receiver coil characteristics in a similar fashion. The hyper-intense signal regions visible in the upper left and right hand portions of the abdominal wall in the three axial images in Fig. 2.3 exemplify this. As already mentioned, the correction of this type of intensity non-uniformity requires a different approach to that presented here, in order to address the issue of continuous greyscale offset compensation within a single image. Procedures most commonly referred to as bias field correction or coil correction (Vokurka et al., 1999, Lai and Fang, 2003) have been developed in order to address this kind of in-slice intensity non-uniformity.




Fig. 2.3: Three consecutive slices from an axially acquired HASTE dataset demonstrating intra-image non-uniformities, most prominently visible in the region of the anterior⁶abdominal wall (arrow) as compared to the posterior abdominal wall (arrowhead).

Inter-image Non-uniformities

The coronal HASTE series that we are working with are far less dramatically affected by this phenomenon. Coronal acquisitions are better suited to our purposes because they provide a more complete and more fully connected representation of the pancreato-biliary system than do their axial counterparts. This is due to the differing ranges of spacial coverage achievable by the respective orientations of these two types of acquisition. This same difference in spacial coverage is also the major reason why the greyscale non-uniformity ef-

⁶Anterior — Located on or towards the front of the body.

fect is far less pronounced and less troublesome in coronal examinations, which do not generally sample in such close proximity to the receiver coils.

The series that we utilise in this work do, however, often exhibit a marked greylevel shift when moving from one slice to the next within the data volume (Fig. 2.4). The effect, where present, is particularly pronounced in the first several slices of these coronal volumes, closest to the abdominal wall, and is due to the nonlinear characteristics of the detected MR signal originating from regions proximal to the scanner’s receiver coils. We wish to compensate for this intensity non-uniformity in order to maximise greyscale homogeneity across the entire volume for each distinct anatomical region imaged, such that the same tissue appears at the same grey intensity range throughout the volume. This is to provide the most consistent data possible to the subsequent filtering and classification procedures. We achieve this greyscale correction effect using a piecewise linear histogram scaling technique based on the alignment of robustly located fixed points across the histogram set (Robinson et al., 2005a).




Fig. 2.4: The first four slices in a fifteen slice MRCP dataset demonstrating inter-image intensity non-uniformity. Intensities in slices one and two in particular are significantly higher than in subsequent slices in the volume.

Inter-volume Non-uniformities

Inter-volume intensity non-uniformities may be observed in any situation where an MRI examination consists of multiple data volumes, either of the same region acquired at different times, or of sequentially acquired, spatially contiguous regions. The time gap between acquisitions may be brief, as is the case with multi-section whole body MRI and cardiac motion imaging, or it may be far more extended, as in the case of disease progression tracking in brain MRI examinations for instance. Regardless of the delay between acquisitions, the phenomenon of intensity non-uniformities may manifest itself, and where

it does, it is necessary to compensate for its effect if automated image analysis procedures are to be successfully applied to the data.

In Fig. 2.5 we can see five slices taken from various positions through a 32 slice coronally acquired multi-section whole body MRI examination, which has been reconstructed from seven individually acquired coronal sections. Greylevel discontinuities can be observed at the interfaces between the individual sections, resulting in the appearance of artificial regional boundaries that would present a significant difficulty to automated segmentation procedures if left unaddressed. In this whole body data, *intra-slice* and *inter-slice* correction proved unnecessary, however, the need for multiple types of correction within a single dataset can arise. In such a case, corrections would be applied, as necessary, in the following order: *intra-slice*, *inter-slice*, *inter-volume*. As previously mentioned, the particular application of these techniques, to whole body MRI data is further explored in Appendix A.

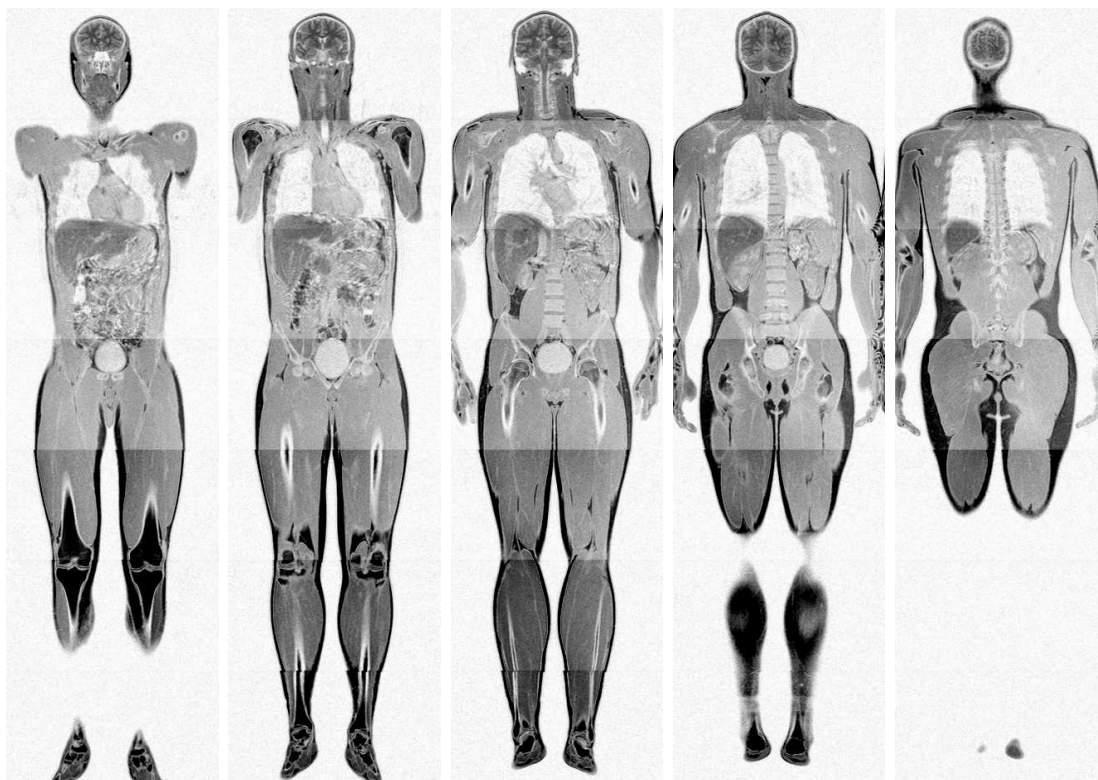


Fig. 2.5: Inter-volume intensity non-uniformity. The seven sections in this WB-MRI dataset were acquired sequentially. Five of the thirty two slices acquired in each section are shown. Intensity steps are clearly visible at the sectional interfaces, most notable between the neck and shoulders, and hips and thighs.

2.2 Data Characterisation

Initially we examined a large number of coronal HASTE datasets in order to characterise the nature of the greylevel intensity distributions typically found in this class of MRCP data. We observed two characteristic peaks, as illustrated in Fig. 2.6, representing signal from airspaces (originating both inside and outside the body), and soft tissues respectively. The first peak is usually the larger of the two. Its size is governed by how much airspace was encompassed in the acquired volume, which can vary significantly. The high intensity signal in which we are most interested, due to the stationary fluids in the scanned volume, lies mostly in the long falling tail of the data histogram. The underlying technique which is described in this chapter can be applied to many types of data. However, a specific application will always start with a re-examination of the histogram characteristics, so as to fine tune the procedure to the data being processed. MRCP and WB-MRI data, for instance, demonstrate slightly different characteristics, and as a result, WB-MRI data requires a slightly modified correction scheme.

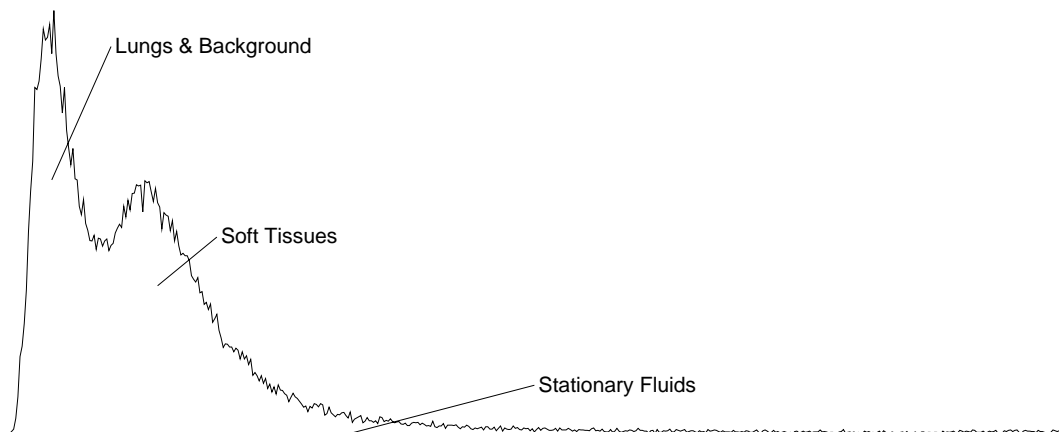


Fig. 2.6: Typical data histogram for a single slice taken from a fifteen slice, coronally acquired HASTE dataset, illustrating the characteristic features found in the greylevel distributions observed in such data.

Occasionally the signal due to stationary fluids also demonstrates a small third peak rising up at the start of the tail region (Fig. 2.7). However, this feature is small enough and infrequent enough to be ignored in any formalised modelling or characterisation of the data. From our observations, it is also clear that there exists significant overlap between these three signal regions representing air, tissue, and stationary fluid. As such, no simple level-based data partition-

ing will prove accurate or robust in delineating the true extents of the various regions imaged within the data, when it comes to the task of segmenting the structures of interest within the MRCP volume.




Fig. 2.7: In this histogram we see a small third peak at the start of the high intensity tail, primarily due to a massively distended gall bladder introducing an unusually high number of bright voxels into the dataset.

The precise greyscale intensity level at which each of the peaks appears in the series of data histograms corresponding to the individual slices in any given volumetric dataset can vary considerable from slice to slice within the dataset, as illustrated in Fig. 2.8. It also becomes clear upon examination of these slice-wise histograms that while the general pattern of two peaks and a long tail holds well, the trough separating the two characteristic peaks can become less distinct, such that the trough and second peak appear more as a simple plateau on the way down from the first peak. This is most clearly demonstrated in the third slice shown in Fig. 2.8.

We note also that the shift observed in the locations of the two maxima from one histogram to another does not necessarily affect the two peaks uniformly. In general the low intensity peak representing the air signal remains very stable in its location and spread while the soft tissue peak changes significantly in both these respects. For the first image shown in Fig. 2.8 the soft tissue peak is broad and shifted to the right as compared to that of the second and third images.

These observations are clearly reflected in the visual appearance of the three slices in Fig. 2.8. The air space in the lungs at the top of each image looks much the same in each case, but the first image looks considerably darker (in its inverted form) in the region of the liver and other organs. In the case of the third image, the soft tissue peak is similarly placed to that of the second image, but is broader, causing the inter-peak trough to be virtually obscured as




Fig. 2.8: Three slices and their corresponding histograms (unsmoothed and smoothed), from a fifteen slice coronal HASTE dataset. All histogram plots have been truncated at 500 bins, the highest greylevel intensities present in the images in fact reach slightly over 1,000.

mentioned previously. This is because of the presence of an increasing number of lower intensity voxels appearing at the periphery of the soft tissue areas, caused by partial volume effects due to the encroaching background signal as the scan nears the subjects back.

The histogram tail becomes increasingly sparse as it extends farther into the higher intensity range. The histograms illustrated in this chapter have all had their tails uniformly truncated for easier review and comparison. In fact, most span more than twice the depicted range and include many empty bins in this upper region. Before processing commences we compress the histograms, squeezing out all the empty bins. This does not adversely affect the data's characteristics and makes the histogram processing easier to manage.

Finally, we note also the level of noise present in the histograms, which represents an additional obscuring factor when it comes to the question of reliably identifying the characteristic points in these data histograms. This noise renders the trough in the histogram corresponding to the third image of Fig. 2.8 particularly indistinct and difficult to localise with any degree of precision. It also introduces a greater level of uncertainty when specifying the location of the characteristic fixed points that we wish to use in matching between the histograms.

As explained in the procedure that is described in the next section, the fixed points just referred to are placed at locations where the histogram slope is large, so as to minimise the effect of the noise, and thus achieve stable positioning and reliable and consistent matching between slices. However, even at these locations the noise present can introduce a destabilising influence, adversely affecting the matching results. Thus in our calculations we use smoothed versions of the histograms, as illustrated on the right in Fig. 2.8 so as to achieve the most robust and stable localisation possible for use in our histogram matching. Smoothing is achieved by averaging over a number of neighbouring bins. We have used a three bin neighbourhood in this process.

2.3 Histogram Matching

The histogram matching procedure commences with the locating of the histogram's global maximum. This will be positioned in either the air peak or


the soft tissue peak. Next the intervening trough and second peak are found in tandem using a robust maximum differential test, searching out from the global maximum in both directions. Once the characteristic peaks representing air and soft tissue have been identified, two fixed points on the outer slopes of the two peaks are recorded. These are the fixed points used in the matching process that follows. The fixed point location procedure is repeated for each slice in the dataset. The histogram scaling necessary in each case is then determined based on the relative locations of these fixed points, and thus each slice has its intensities remapped in order to bring the fixed points into alignment across all the slices, in such a fashion as to preserve the original histogram shape, avoiding the introduction of spurious spikes and voids.

2.3.1 Fixed Point Detection

Once the data histogram has been constructed for a slice the largest bin is found and its location in the histogram is recorded (A). This maximum is usually located in the air peak (as illustrated here) but occasionally, when there is very little airspace in the scanned region, it will lie within the soft tissue peak. Because the intensity distribution in air voxels tends to be much tighter than it is in soft tissue voxels, the air peak is generally tall and narrow as compared to that of the soft tissue signal. As a result, a relatively small region of airspace in the image can result in a peak that is taller than the soft tissue peak, even though the area under the peak (i.e the number of voxels in that intensity range) may be significantly less than the corresponding area under the soft tissue peak.




Next we need to identify locations for the trough (B) and second peak (C). Since we are operating in the presence of noise we must be careful that we are locating the correct features in the histogram. A secondary noise maximum in the first peak might for instance be mistakenly identified. After examining the potential difficulties and testing the performance of a number of detection strategies we arrived at a scheme that performs in a robust fashion even when presented with a weakly defined trough and second peak in the presence of significant noise, as is the case with the third image in Fig. 2.8.



Most approaches to histogram mode detection use arbitrarily chosen threshold levels or tuning parameters to overcome the problems that the presence of noise introduces. It may be decided that, on moving away from the global maximum, the level must drop to 50% of the maximum value before the next peak is sought, or that any peak must be wider than 100 bins, to some definition of width, or must demonstrate a consistent slope direction when averaging over a predetermined number of bins for a given distance, or some other such arbitrarily or experimentally determined mode detection scheme.

Our method utilises a robust and meaningfully defined metric that does not require any tuning, initialisation, or parameterisation. By performing the directed search described below we identify the pair of bins that demonstrate the largest trough to peak range in the data, in which the identified local minimum position lies between the histogram’s global maximum and the identified local maximum. This represents the most reliable and also the most predictable method of identifying the two most significant modes in a multi-modal histogram in the presence of noise.

Our scheme searches for the maximum differentiating measure in the histogram data. Using the fact that we are looking for two peaks separated by an intervening trough, we move out from the global maximum A, (which we assert must lie within one of the two histogram modes that we seek to identify), and we test each pair of candidate positions for the trough and second peak, so as to maximise the vertical distance (δ) from trough to second peak.



The pair of histogram positions that yield the largest value for δ gives us the pair of bins we seek. We note that while the three identified points can reliably be taken each to lie within its corresponding histogram feature, it is only the levels, and not the precise locations of A, B, and C that can be assumed to be stable and meaningful.

In the case illustrated, C in particular could very easily jump to any one of a number of different local maxima located along the top of the broad, flat soft tissue peak with just the smallest change in the underlying histogram. Thus C reliably identified the feature of interest, but should not be taken to represent some special location within this feature (the location of the ‘*top*’ of the peak in this case). We must proceed with a further step in order to ensure the stability of our final set of fixed points with such noisy histogram characteristics.

Algorithm 2.1⁷: *min/max pair finding*


1. *Size is the size of the histogram*
2. *H_n is the n^{th} bin of the histogram (from 1 to Size)*
3. *A is the position of the global max*
4. $B = C = 1$
5. $\forall b \in A + 1 \dots \text{Size} - 1$
6. $\forall c \in b + 1 \dots \text{Size}$
7. $\left| \begin{array}{l} \text{If } (H_c - H_b > H_C - H_B) \end{array} \right.$
8. $\left| \left| \begin{array}{l} B = b \end{array} \right. \right.$
9. $\left| \left| \left| \begin{array}{l} C = c \end{array} \right. \right. \right.$

Algorithm 2.1 above illustrates the process for identifying B and C in the forward direction. The same procedure is applied looking backwards from the global maximum towards the histogram origin. The pair of positions yielding the overall largest value for δ represent the minimum/maximum set to be used in the subsequent calculations performed to locate the fixed points needed by the histogram rescaling step that follows.

Line five iterates through every bin from the global maximum up to the end of the histogram. Then in line six we search out from each of these bins looking for the largest δ we can find. An explicit check is not required to test that b is a minimum and c is a maximum as the test in line seven naturally identifies the minimum and maximum that we intend for the algorithm to find. It will of course locate many non-minima non-maxima pairs along the way but they will always be superseded by a pair consisting of a local minimum at b and a local maximum at c before the search terminates. And so, by the end of this process we have identified integer positions for A , B , and C in the histogram.


⁷The notation used in these algorithms is based on set theory and the predicate calculus, and follows the formal specification style described in Woodcock and Loomes (1988).

However, as already mentioned, with the histogram characteristics that we have observed, while the levels of A, B, and C represent robust metrics for a given




histogram under perturbation, their positions are not necessarily so stable, due to the significant noise found in the histograms, especially in the presence of broad flat topped peaks and broad flat bottomed troughs. Greater stability can be seen in the regions of the histogram where the values are changing more rapidly, and as such we can provide a more stable and consistent set of alignment points by selecting positions on the outer slopes of the two peaks (P_1 and P_2), which we place at a level one third of the way down from A and C, moving to the left and right respectively, rather than using the more unstable positions of the peaks themselves. In other words we seek to align the bodies of the histogram modes, not simply their local maxima positions. The choice of fixed point location, one third of the way down, is a trade-off between getting off a potentially broad, noisy peak, and not encroaching on the long, falling tail of the histogram. Locating the fixed points in either of these two regions would result in increased instability and uncertainty in their placement.

It can for instance be seen that P_2 would be placed almost identically regardless of which of the three widely separated potential highest peaks C_1 , C_2 , and




C_3 was selected. In this case there is only a difference of a single voxel between the lowest and highest of these three local maxima, C_1 corresponds to a value of 410 voxels at intensity level 82, while C_2 at intensity 98, and C_3 at intensity 106 both take a value of 409 voxels. Thus an extremely small change in the voxel intensity distribution in the data could easily result in the highest of these maxima being overtaken by one of its close rivals, still identifying the correct feature in the histogram but at a significantly different position. This leads us to define the fixed points P_1 and P_2 in the manner described, ensuring that such instabilities do not adversely affect the matching process.

We can see how, particularly at the location of P_2 above, noise can still introduce some jitter where the slope is more gentle, and so we smooth the histogram



data by averaging over a neighbourhood of three bins, before performing the peak and fixed point location procedure, so as to further improve the stability of the results achieved.

The points P_1 and P_2 subdivide the histogram into three regions, each of which is scaled independently, leading to the piecewise linear scaling mentioned at the outset. In addition to these two fixed points, we use the origin (P_0) and the maximum intensity value (P_3) in each histogram in order to complete the set of fixed points used in the matching process. In fact the tail in the displayed histogram is severely truncated such that P_3 actually lies much further to the right than is shown here. We use these two additional constraints of the histogram's origin and size in the matching process, as voxel intensities should not be allowed to stray into the negative range, nor should they be allowed to take on arbitrarily high intensity levels.



2.3.2 Modelling Histograms as Continuous Functions

In order to perform the desired rescaling, we model each histogram as a continuous function so that we can resample that function at a new set of intervals in order to arrive at the required relative levels for the new set of histogram bins. So as to present a consistent treatment throughout and so as to maintain the maximum levels of precision in the process, we perform all calculations in continuous space, treating the histogram bin values as a series of sampled data points along this notional continuous function. Thus in the final implementation we use a quadratic model to locate the sub-integer coordinates of the two maxima A and C in each case. See the text box *'Locating Local Maxima to Sub-integer Precision'* for the mathematical details.

We then use linear interpolation (see the text box *'Linear Interpolation'*) in order to calculate the exact locations of the fixed points P_1 and P_2 , followed by a combination of linear and cubic interpolation (see the text box *'Cubic Interpolation'*) for the calculation of the entire new set of histogram bin levels, which are positioned at locations intermediate to the existing bins.

In order to formalise our presentation we declare a set of useful functions that are needed to perform the calculations necessary to implement the procedure outlined above. The underlying mathematics required in order to implement all of these functions is provided in the three previously mentioned highlighted text boxes that appear on the following three pages.

Linear Interpolation

Linear interpolation provides a quick and easy method to calculate intermediate values between pairs of sample points in a data series. Given two non-coincident points (x_1, y_1) and (x_2, y_2) we can fit a line $y = mx + c$, and so define m and c :

$$m = \frac{y_1 - y_2}{x_1 - x_2}$$


$$c = \frac{x_1 y_2 - x_2 y_1}{x_1 - x_2}$$

Thus for any given x we can calculate a value for y :

$$\begin{aligned} y &= mx + c \\ &= \frac{y_1 - y_2}{x_1 - x_2}x + \frac{x_1 y_2 - x_2 y_1}{x_1 - x_2} \end{aligned}$$

And for any given y we can calculate a value for x :

$$\begin{aligned} x &= \frac{y - c}{m} \\ &= \frac{y - \frac{x_1 y_2 - x_2 y_1}{x_1 - x_2}}{\frac{y_1 - y_2}{x_1 - x_2}} \\ &= \frac{y(x_1 - x_2) - (x_1 y_2 - x_2 y_1)}{y_1 - y_2} \end{aligned}$$



Cubic Interpolation

Cubic interpolation, which is particularly widely used for interpolation tasks in medical imaging applications, is used in a number of places throughout this work in order to resample data at intermediate locations between existing sample points. In general it provides superior performance to linear interpolation, but requires four sample points, in comparison to the linear scheme's two. The following treatment shows how to estimate y for any given x on the basis of four uniformly spaced samples, typically two either side of the desired location, based on a cubic model, $y = ax^3 + bx^2 + cx + d$.

Consider four sample points $(-1, y_1)$, $(0, y_2)$, $(1, y_3)$, and $(2, y_4)$. Inserting these into the cubic leads to four simultaneous equations in the four unknowns a , b , c , and d :

$$y_1 = -a + b - c + d \tag{2.1}$$

$$y_2 = d \tag{2.2}$$

$$y_3 = a + b + c + d \tag{2.3}$$

$$y_4 = 8a + 4b + 2c + d \tag{2.4}$$


We can combine and rearrange these equations in order to solve for the four unknowns in the standard fashion:

$$d = y_2$$

$$b = \frac{y_1 - 2y_2 + y_3}{2}$$

$$a = \frac{-y_1 + 3y_2 - 3y_3 + y_4}{6}$$

$$c = \frac{-2y_1 - 3y_2 + 6y_3 - y_4}{6}$$



We substitute these values into the cubic to calculate a value for y at any given intermediate position of x . Typically x will be in the range zero to one, spanning the gap between the inner two samples. Moving away from the sampled points reduces the likelihood of a close fit to the real data, although boundary cases can be covered by allowing x to range from minus one to two in order to facilitate interpolation right up to the edges of the sampled dataspace. We perform interpolation in two or more dimensions by composing a series of 1-D interpolations in orthogonal directions.

Locating Local Maxima to Sub-integer Precision

Consider the three successive sample values v_1 , v_2 , and v_3 , where v_2 represents a local signal maximum ($v_1 < v_2 > v_3$). If we fit a quadratic ($y = ax^2 + bx + c$) through these three points then the location of the quadratic's maximum ($\frac{dy}{dx} = 0$) represents the best estimate of the true local maximum's location.

In order to calculate the position as an offset from v_2 we consider the three points as having coordinates $(-1, v_1)$, $(0, v_2)$, and $(1, v_3)$. This yields the following three simultaneous equations:

$$v_1 = a - b + c \quad (2.5)$$

$$v_2 = c \quad (2.6)$$


$$v_3 = a + b + c \quad (2.7)$$

We can thus solve for a , b , and c in the sought quadratic:

$$c = v_2$$

$$b = \frac{v_3 - v_1}{2}$$

$$a = \frac{v_3 + v_1 - 2v_2}{2}$$



Now differentiating and setting equal to zero we get:

$$\frac{d}{dx}ax^2 + bx + c = 2ax + b = 0$$

$$\Rightarrow x = -\frac{b}{2a} = \frac{v_3 - v_1}{4v_2 - 2(v_3 + v_1)}$$

$$y = -\frac{b^2}{4a} + c = \frac{(v_3 - v_1)^2}{16v_2 - 8(v_3 + v_1)} + v_2$$

Hence we know x , the offset, and y , the level. Note that if v_1 equals v_3 then the offset is zero, the symmetrical case. Further, since v_1 and v_3 are always smaller than v_2 we can see that the denominator must be positive and hence the offset is negative if v_1 is larger than v_3 and positive if v_1 is smaller than v_3 .

If $v_1 = v_2$ then $x = -\frac{1}{2}$ and if $v_2 = v_3$ then $x = \frac{1}{2}$. The situation where $v_1 = v_2 = v_3$ is a degenerate case (the quadratic becomes a straight line). These observations conform to the expected range and motion of the calculated maximum.

The first pair of functions use linear interpolation to calculate either the x or y coordinate of a point on a straight line ($y = mx + c$), given the other coordinate plus two additional points on the line. Linear interpolation is used in two places, first to calculate the x positions of P_1 and P_2 , and second in the piecewise linear mapping of each bin in the rescaled histogram back to a corresponding position in the original as illustrated in Fig. 2.9.

$$x = \text{LinearX}(y, x_1, y_1, x_2, y_2)$$

$$y = \text{LinearY}(x, x_1, y_1, x_2, y_2)$$

The next pair of functions is used to calculate the x and y coordinates respectively of the maximum/minimum point on the quadratic ($y = ax^2 + bx + c$) that passes through the three points $(-1, y_1)$, $(0, y_2)$, and $(1, y_3)$. We use these functions to calculate the best estimate for the true location and height of the maxima corresponding to the positions **A** and **C** identified previously. As discussed, we only actually need the height, which is used in the definition of points P_1 and P_2 , but derivations for both are provided in the text box.

$$x = \text{QuadraticMaxMinX}(y_1, y_2, y_3)$$

$$y = \text{QuadraticMaxMinY}(y_1, y_2, y_3)$$

The final function uses a cubic model ($y = ax^3 + bx^2 + cx + d$) to calculate the y value at any given x , for a point that lies on a cubic curve passing through the four points $(-1, y_1)$, $(0, y_2)$, $(1, y_3)$, and $(2, y_4)$. This is used in interpolating the inter-bin levels from the original histogram for the construction of the final rescaled version.

$$y = \text{CubicY}(x, y_1, y_2, y_3, y_4)$$

By the end of the *min/max pair finding* algorithm fragment presented above we had arrived at values for **A**, **B**, and **C**, the locations of the discrete global and local maxima pair, and the intervening local minimum. We can now formulate the necessary algorithm to calculate the x positions of the four fixed points P_0 , P_1 , P_2 , and P_3 . This in turn allows us to define the piecewise linear mapping function illustrated in Fig. 2.9 and so to proceed with the histogram re-binning phase of the procedure.

Algorithm 2.2: *calculating fixed points*

1. *Size* is the size of the histogram
2. H_n is the n^{th} bin of the histogram
3. A and C are the histogram maxima previously found
4. $P_0 = 1$
5. $val = \frac{2}{3} \text{QuadraticMaxMinY}(H_{A-1}, H_A, H_{A+1})$
6. $p = A - 1$
7. *while*($H_p > val$)
8. | *dec*(p)
9. $P_1 = \text{LinearX}(val, p, H_p, p + 1, H_{p+1})$
10. $val = \frac{2}{3} \text{QuadraticMaxMinY}(H_{C-1}, H_C, H_{C+1})$
11. $p = C + 1$
12. *while*($H_p > val$)
13. | *inc*(p)
14. $P_2 = \text{LinearX}(val, p - 1, H_{p-1}, p, H_p)$
15. $P_3 = \text{Size}$

In lines five and ten we work out the height of the modelled maxima in our notional continuous function at locations close to the discrete histogram maxima previously identified at points A and C , using the quadratic calculation introduced earlier. The *while* loops starting at lines seven and twelve then find the histogram bins flanking points P_1 and P_2 respectively, and then linear interpolation is used in lines nine and fourteen to calculate the precise locations of these two fixed points. Along with the trivial assignment for P_0 in line four and P_3 in line fifteen, this gives us values for the locations of the four fixed points required for the resampling process.

2.3.3 Histogram Resampling

Once these four ‘P’ values have been determined for each slice in the volume, the set of values with the largest P_2 is selected and used to define the target histogram configuration for all the slices. The rescaled histograms can now be calculated using the identified P_{1Max} , P_{2Max} pair, along with each image histogram’s own set of calculated fixed points (P_0 to P_3), in order to define the required histogram scaling function, mapping the original histograms into their final rescaled versions (Fig. 2.9). Since initially the number of bins in the histogram is to be held constant (P_3 maps to P_3) the next task is to determine how many voxels should be allocated to each bin in the rescaled version.




Fig. 2.9: Piecewise linear scaling

Having first calculated values for P_1 and P_2 , and applying the two additional constraints, that the origin stays fixed (P_0 maps to P_0) and the number of bins in the histogram also initially remains fixed (P_3 maps to P_3), we can now use the piecewise linear model described in order to map each bin in the new histogram back to a real valued position in the original

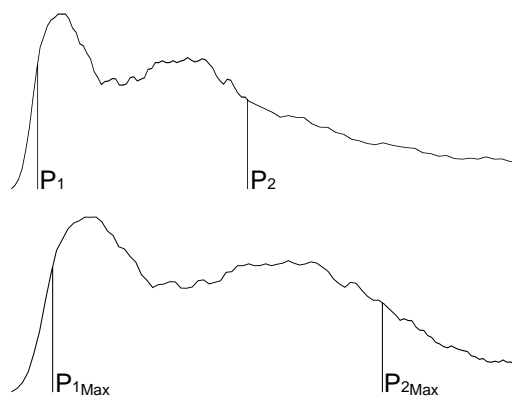


Fig. 2.10: Before and after scaling

and then, using cubic interpolation, we can calculate the size of each new bin in the final histogram (Fig. 2.10). These relative sizes must next be scaled uniformly so that the sum of all the histogram bins equals the number of voxels in the slice. Some additional manipulation is then required in order to arrive at integer values for the size of each bin, (since only a whole number of voxels can be assigned to a bin). Step one in this process is to perform the piecewise linear histogram re-binning procedure in order to arrive at the real-valued resampling of the notional continuous function as illustrated in Fig. 2.10.

Algorithm 2.3: *histogram re-binning*

1. H_n is the n^{th} bin of the histogram
2. P_0, P_1, P_2 , and P_3 are the fixed points previously found
3. $P_{2_{\text{Max}}}, P_{1_{\text{Max}}}$ are the largest P_2 and corresponding P_1
4. $\forall i : \mathbb{N} \in P_0 \dots P_{1_{\text{Max}}}$
5. $\left| \begin{array}{l} val = \text{LinearY}(i, P_0, P_0, P_{1_{\text{Max}}}, P_1) \\ ip = \text{floor}(val) \\ H'_i = \text{CubicY}(val - ip, H_{ip-1}, H_{ip}, H_{ip+1}, H_{ip+2}) \end{array} \right.$
6. $\left| \begin{array}{l} ip = \text{floor}(val) \\ H'_i = \text{CubicY}(val - ip, H_{ip-1}, H_{ip}, H_{ip+1}, H_{ip+2}) \end{array} \right.$
7. $\left| \begin{array}{l} H'_i = \text{CubicY}(val - ip, H_{ip-1}, H_{ip}, H_{ip+1}, H_{ip+2}) \end{array} \right.$
8. $\forall i : \mathbb{N} \in P_{1_{\text{Max}}} \dots P_{2_{\text{Max}}}$
9. $\left| \begin{array}{l} val = \text{LinearY}(i, P_{1_{\text{Max}}}, P_1, P_{2_{\text{Max}}}, P_2) \\ ip = \text{floor}(val) \\ H'_i = \text{CubicY}(val - ip, H_{ip-1}, H_{ip}, H_{ip+1}, H_{ip+2}) \end{array} \right.$
10. $\left| \begin{array}{l} ip = \text{floor}(val) \\ H'_i = \text{CubicY}(val - ip, H_{ip-1}, H_{ip}, H_{ip+1}, H_{ip+2}) \end{array} \right.$
11. $\left| \begin{array}{l} H'_i = \text{CubicY}(val - ip, H_{ip-1}, H_{ip}, H_{ip+1}, H_{ip+2}) \end{array} \right.$
12. $\forall i : \mathbb{N} \in P_{2_{\text{Max}}} \dots P_3$
13. $\left| \begin{array}{l} val = \text{LinearY}(i, P_{2_{\text{Max}}}, P_2, P_3, P_3) \\ ip = \text{floor}(val) \\ H'_i = \text{CubicY}(val - ip, H_{ip-1}, H_{ip}, H_{ip+1}, H_{ip+2}) \end{array} \right.$
14. $\left| \begin{array}{l} ip = \text{floor}(val) \\ H'_i = \text{CubicY}(val - ip, H_{ip-1}, H_{ip}, H_{ip+1}, H_{ip+2}) \end{array} \right.$
15. $\left| \begin{array}{l} H'_i = \text{CubicY}(val - ip, H_{ip-1}, H_{ip}, H_{ip+1}, H_{ip+2}) \end{array} \right.$

The three *for* loops starting at lines four, eight, and twelve correspond to the three sections in the piecewise linear mapping function illustrated in Fig. 2.9. Each loop iterates through all the integers in the given range, and consists of three steps. First we use linear interpolation to determine the real valued position along our notional function modelling the original histogram that corresponds to the current discrete bin location in the new histogram. We round this number down in order to identify the nearest actual histogram bin below it in the original histogram.

Then using cubic interpolation and the two discrete histogram bins either side of the real valued location in the original histogram we model the expected histogram level at the intermediate position. This becomes the level at the current discrete bin location in the new rescaled histogram. By the end of this process we have a series of *Size* real-valued numbers corresponding to the required relative levels of the final histogram bins. That is to say, these numbers are directly proportional to the number of voxels that should finally be assigned to each bin or intensity level.

However, we can only assign an integer number of voxels to each bin, and the total number of voxels assigned must also be equal to the number of voxels in the slice (this is similar to the bit assignment problem found in compression algorithms (Jayant and Noll, 1984)). Consider an image with \mathbf{N} voxels, and a series of real valued histogram bin levels stemming from the above process, H_1 to H_{Size} , which sum to the real value R . In order to determine the number of voxels that we must assign to each intensity level we proceed as in Algorithm 2.4 below.

Algorithm 2.4: *bin size recalculation*

```

1.  Sum = 0
2.   $\forall i \in 1 \dots Size$ 
3.  |  $H_i = round(\frac{H_i \times \mathbf{N}}{R})$ 
4.  |  $Sum = Sum + H_i$ 
5.  Remainder =  $\mathbf{N} - Sum$ 
6.  If Remainder < 0
7.  |  $Remainder = -Remainder$ 
8.  |  $i = Size$ 
9.  | While Remainder >  $H_i$ 
10. | |  $Remainder = Remainder - H_i$ 
11. | |  $H_i = 0$ 
12. | |  $dec(i)$ 
13. |  $H_i = H_i - Remainder$ 
14. |  $Size = i$ 
15. ElseIf Remainder > 0
16. |  $i = Size + 1$ 
17. | While Remainder >  $H_{Size}$ 
18. | |  $Remainder = Remainder - H_{Size}$ 
19. | |  $H_i = H_{Size}$ 
20. | |  $inc(i)$ 
21. |  $H_i = Remainder$ 
22. |  $Size = i$ 

```

The *for* loop starting at line two first normalises the bin values so that they sum to the required total, and then rounds the values to the nearest integer and accumulates the actual sum achieved. Due to this rounding the actual number of binned voxels will not immediately achieve the required value. The maximum discrepancy possible is equal to $\frac{Size}{2}$ voxels but generally the mismatch is much smaller than this. The actual remainder is calculated in line five, and depending on whether it is positive or negative one of the two *if* clauses starting at lines six and fifteen are executed in order to arrive at the correct number of voxels in the final histogram.

This final correction phase of the algorithm is a natural consequence of the necessary shift from a continuous to a discrete mode in the analysis. Due to the relatively small number of voxels involved and the natural variability of the tail length in the original data, this procedure does not compromise the underlying greyscale characteristics of the data after matching, and thus allows the histogram preserving non-uniformity correction algorithm to operate effectively.

2.3.4 Sample Value Reassignment

The final step involves assigning each voxel to its new bin, and therefore by definition assigning it its new intensity value. In order to achieve this we need to define a strict ordering on all the voxels in the slice. We already have a partial ordering given by their initial intensities, so that what we now need is to be able to define an ordering within each bin of the original histogram, thus providing a total ordering for all the voxels, allowing us to assign the required number of voxels to each bin in the new histogram.

We use an ordering function based on the offset-weighted intensity mean measured in a local neighbourhood around each voxel in order to specify the strict ordering required:

$$\mu_{wt} = \frac{\Sigma(val/d)}{\Sigma(1/d)} \tag{2.8}$$

where *val* is the neighbour's value and *d* is the offset from the current voxel to the neighbour. This scheme preserves the underlying greylevel distribution

and the relative intensities for individual voxels within each slice. No voxel's intensity will leapfrog any other's. We are only providing a means of splitting the voxels in a given bin so that we can distribute them in any required proportion between two or more neighbouring bins. Alternatively, an error diffusion (Ulichney, 1987) based reassignment of the voxels, similar to the process employed in dithering, would be a valid approach to this reassignment phase. The underlying goal is to avoid the introduction of stripe artifacts. Either of these two approaches will achieve this goal, while a simple scan-order based reassignment typically will not.

The ordering function works on the principle that the brighter a voxel's neighbours are, the more likely it is that that voxel should take on a brighter value itself. Thus the intra-bin ordering is performed such that voxels within each bin are given a weighting that specifies a relative tendency towards brighter or darker, based on the local neighbourhood. Essentially each voxel is assigned a fractional part to their intensity values, allowing for the required strict ordering to be specified.

2.4 Non-uniformity Correction Results

Once the intensity non-uniformity correction procedure has been applied and the set of volume slices have been homogenised such that the characteristic points on their histograms are matched and their greylevel distributions have been brought into alignment, corresponding tissue regions in each slice are updated to occupy the same grey range from slice to slice. Fig. 2.11 shows the same three slices as were first shown in Fig. 2.8. Regional intensity variations previously visible from slice to slice have been greatly reduced, resulting in a far more homogenous set of data regions within the volume as a whole. The two sets of histograms shown, illustrate the slice-wise greyscale distributions before and after intensity non-uniformity correction, and demonstrate the improved levels of histogram feature alignment achieved after the homogenisation procedure has been applied. We can see that the soft tissue peaks in the second and third images in Fig. 2.11 have been broadened into the higher intensity region, so as to more closely match that of the first image (which yielded the largest P_2 in this case).




Fig. 2.11: Three slices before and after the histogram matching procedure has been applied to the data. Also shown are their corresponding histograms before and after matching, demonstrating the improved alignment of features achieved in the matching process.

The correction effect achieved is even more strikingly obvious in Fig. 2.12, which shows the result of applying the same intensity non-uniformity correction scheme to a whole body MRI dataset, where significant greyscale mismatch is visible at the interfaces of the individual coronal sections that are combined to arrive at the final reconstructed dataset. Again the histograms from before and after matching clearly illustrate how the soft tissue peaks, which start out at varying intensity positions, are brought into good alignment after the application of this technique. Even the extremely low soft tissue peak in the lowermost of the seven coronal sections has been successfully identified and aligned. This last peak is as relatively small as it is because the only non-air voxels in this section are in the feet and ankles, which occupy only a small proportion of the total volume covered by the section. Even with such a large mismatch between the first and second histogram modes, our histogram mode detection scheme has successfully identified the correct features and thus our intensity non-uniformity correction algorithm arrived at an appropriate set of greylevel intensity correction mappings across all sections.

In a similar fashion to the characteristic MRCP histograms, a number of the WB-MRI histograms in Fig. 2.12 also demonstrate a third peak, in this case due to high intensity signal from fat tissue rather than the high intensity stationary fluid signal typical of the MRCP class of MRI protocols. Again this peak is not always present and where it is, it is consistently the smallest of the three and as such does not present any difficulty for the operation of the correction procedure.

Implementations for key sections, corresponding to the individual steps detailed in this chapter, are presented in Appendix B where we describe the software environment that was developed in the course of this project. Typical time to process one MRCP volume ($255 \times 255 \times 15$ voxels) using this implementation is approximately three to four seconds.

The voxel ordering function used in this code examines a five by five neighbourhood so as to determine a voxel's weighting, and does not guarantee an absolute ordering as it is possible for two voxels of the same base intensity to also manifest the same weighting value. We applied the intensity non-uniformity correction procedure to a database of 24 MRCP volumes in order to assess its performance. As illustrated in Table 2.1 the situation where an absolute ordering of the voxels is not achieved arises most infrequently in the




Fig. 2.12: Slices from the seven sections of a whole body MRI study along with histograms for each section, before and after matching. After matching the seven lean tissue peaks are well aligned and discontinuities at the sectional boundaries are minimised.

real world MRCP data examined (on average about once in sixty five thousand voxels), and as such it does not present any difficulty in this context. In all but seriously pathological cases (e.g. a completely uniform image where no discrimination is possible) this consideration does not compromise the effectiveness of the approach discussed in this section. The very few indistinguishable voxels are ordered by their scan line positions within the slice and allocated to the new histogram bins as such.

Volumes in database	24
Total slices	360
Voxels per slice	65025
Indistinguishable voxel pairs per slice	
Minimum	0
Maximum	7
Mean	1.094444
Variance	1.818857

Table 2.1: Voxel ordering discriminant function effectiveness

In developing this greyscale non-uniformity compensation procedure we have provided a method of correcting for potentially troublesome inter-slice and inter-volume intensity inhomogeneities, enabling a more robust and more consistent analysis of the processed data to be carried out by maximising regional homogeneity in the data, while avoiding the introduction of additional difficulties at later stages due to the creation of compromised data histograms. In so doing we maximise the levels of confidence with which we can extract information from the data, and provide the best possible starting point for the processing to follow.

Chapter 3

Adaptive Gaussian Smoothing

In this chapter we describe a novel and highly versatile boundary preserving approach to data smoothing and noise suppression that operates on non-isometric volume data (Robinson, 2004). The method evolves naturally from the familiar Gaussian smoothing model and utilises a dynamically modified Gaussian mask, scaled and shaped by the magnitude and direction of the local gradient vector. The filter is parameterised for strength and anisotropy, and operates on 3-D data with non-cubic voxels to selectively smooth in a manner that retains prominent boundary features while attenuating signal noise.

Many approaches to data smoothing and noise reduction appear in the literature ranging from the simplest averaging and median filters (Gonzalez and Woods, 1992) through mathematical morphology (Serra, 1982, Soille, 1999) and more involved spatial and frequency domain filtering schemes (Greenspan et al., 2000, Whelan and Molloy, 2000). Among these one of the simplest and most widely used techniques has long been that of Gaussian smoothing (Davies, 1990, Haralick and Shapiro, 1992).

In recent years, however, such approaches have been to a large extent superseded by more advanced locally adaptive data smoothing methods that aim to preserve semantically important details while at the same time attenuating noise in the data. Approaches based on wavelets (Jung and Scharcanski, 2004), tangential smoothing (Bromiley et al., 2002), variational methods (Schnorr, 1999) and nonlinear and anisotropic diffusion models (Perona and Malik, 1990, Gerig et al., 1992, Weickert et al., 1998) have all received considerable attention.

Notwithstanding the fact of more advanced methods, the underlying principles of Gaussian smoothing remain useful and much used, due in no small part to its simplicity and elegance (Smith, 1999). Indeed some of the more advanced techniques mentioned evolved directly from this same root, usually adding significant complexity on their way to arrive at more effective approaches to smoothing the data.

In our approach, we attempt to retain the fundamental simplicity of the original Gaussian smoothing methodology while at the same time offering levels of performance and adaptability on a par with those achieved by the more complex and unwieldy alternatives (Lynch et al., 2004, Ghita et al., 2005*a*). This leads to a family of filters that are simple and intuitive to use and to tune to the specific characteristics of the data being processed.

3.1 Gradient-Weighted Gaussian Filter

At the root of the basic Gaussian smoothing approach is the idea of convolution with a Gaussian mask, replacing each data sample by the weighted mean of samples within a given neighbourhood. The weights in the weighted mean are calculated so as to fall off with distance from the centre point, based on a simple exponential form, $y = e^{-x^2}$, which reaches a maximum of one at x equal to zero and approaches zero as x tends to $\pm\infty$ (Fig. 3.1).





Fig. 3.1: Graph of the function: $y = e^{-x^2}$

A parameter modelling the standard deviation of the exponential function allows for control over the degree of smoothing achieved, by varying the width of the central peak in the convolution kernel used, as illustrated in Fig. 3.2. Thus in its simplest form the mask weights used to arrive at the weighted mean can be calculated as shown in Eq. 3.1 (the final values are divided by

the sum of these weights in order to normalise the filter and maintain the data within its original grey range). In this equation, d is the distance from the filter mask's origin to each neighbour in the convolution kernel, and σ is the chosen standard deviation or smoothing strength parameter.

$$f(d) = e^{-\left(\frac{d}{\sigma}\right)^2} \quad (3.1)$$

A larger value for σ results in a wider distribution, larger contributions from neighbouring sample points, and thus more aggressive smoothing. The degree of smoothing is, however, constant across the filtered space for any given value of the parameter σ . No account is taken of the local structure of the data being smoothed. To improve on this scheme the next step is to allow the smoothing applied to be modified locally, based on a neighbourhood of samples, an adaptive filtering approach.



(a) $\sigma = 0.5$ (b) $\sigma = 1.0$ (c) $\sigma = 2.0$

Fig. 3.2: Gaussian masks of varying widths, shown in one, two, and three dimensions.

3.1.1 Adaptive Filtering

In order to extend the familiar Gaussian approach we wish to modify the mask weights locally, based on the magnitude and direction of the gradient vector calculated at the current point. In general, a strong local gradient vector suggests the presence of a significant boundary in the locality, which we would not wish to blur or dislocate. In order to achieve this edge preserving effect, the gradient vector's *magnitude* can be used to govern the strength of the applied smoothing by modifying the width of the Gaussian kernel used in the smoothing mask at that location.

Similarly, the *direction* of the gradient vector can be used to enable the modelling of an isosurface through the current point as the plane perpendicular

to the gradient. Thus using this model we can reduce the effective contribution to the weighted mean of those sample points that fall farther from the isosurface, while allowing those neighbours close to the isosurface to have a significant effect on the final result. In this way we restrict the smoothing to occur primarily along the isosurface rather than across it thus avoiding the blurring or dislocation of boundaries whilst still allowing noise in the vicinity of strong edges to be efficiently smoothed.

For completeness sake we also mention the less useful (in the context of our work) third alternative, that of giving favour to smoothing across the isosurface, along the direction of the local gradient vector, thus actively encouraging the suppression of strong boundaries. This can be seen as the effective dual of the approach we take. Such a technique might be found useful in a feature suppression application, perhaps in the removal of power lines or other strong edge features in photographic images for instance. It does not, however, lend itself to any useful application in our work and is not investigated further here.

The nature of the 3-D Gaussian masks required in order to implement each of these three nonlinear filter types, isotropic, boundary suppressing, and boundary preserving, are shown in Fig. 3.3, where 3.3a illustrates the isotropic case, simple drop-off with increasing distance from the current point uniform in all directions, 3.3b encourages smoothing across boundaries (the cylinder axis being oriented parallel to the local gradient vector), and 3.3c shows the boundary preserving case, with the uniform intensity plane direction being perpendicular to the gradient, to model the inferred isosurface.

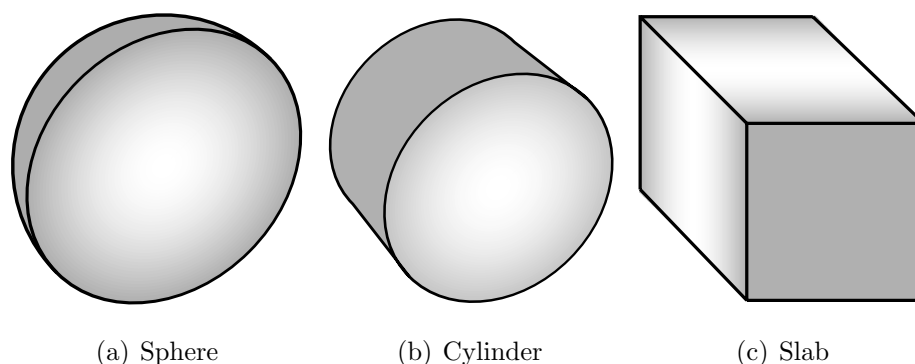


Fig. 3.3: 3-D distance maps demonstrating Gaussian drop-off with distance from a point, a line, and a plane respectively. Intensities in these figures are directly proportional to the actual weight values in the corresponding filter masks.

This last case is the one in which we are most interested here. Furthermore, in all three cases the gradient magnitude can be used to modify the “size” of the mask and thus the degree of smoothing applied at each point.

The form of the resulting algorithm used to apply the desired filtering is shown below as Algorithm 3.1. The neighbourhood is taken to include the centre point p itself. The precise nature and characteristics of the filtering achieved are controlled by the shape of the weighting function $f(\vec{pq}, \nabla u)$. Manipulation of this function is the key to our approach’s performance, and allows for exceptional flexibility within the simple framework outlined here. Also of importance, and addressed in the next section of this chapter, is the consistent calculation of the local greyscale gradient vector ∇u , in volumetric data with non-cubic voxels.

Algorithm 3.1: *gradient-weighted adaptive Gaussian filter*

For each sample point p

- | $\nabla u = \text{local gradient vector at } p$
- | $wt_{sum} = 0$
- | $val = 0$
- For each neighbour q
 - | $\vec{pq} = \text{vector from } p \text{ to } q$
 - | $wt_q = f(\vec{pq}, \nabla u)$
 - | $wt_{sum} += wt_q$
 - | $val += wt_q q$
- | $p = \frac{val}{wt_{sum}}$

3.1.2 Gradient Vector Calculation

Consider the case of a 3-D non-isometric data array with voxel dimensions dx , dy , and dz as illustrated to the right in Fig. 3.4. In order to perform the calculations for anisotropic filtering, a gradient vector $\nabla u = [\nabla u_x, \nabla u_y, \nabla u_z]^T$ must be calculated at each location in the data, to be used in governing the strength and directionality of the filtering.




Fig. 3.4: Non-isometric data grid with unequal orthogonal grid spacing.

Familiar gradient calculation methods such as the Sobel operator (Sobel, 1990, Duda and Hart, 1973) make an assumption of isometric data, so that a simple extension of any such operator into three dimensions is not suitable for use in the case of volumetric data with non-cubic voxels. In place of the familiar pair of static 3×3 filter masks used in most traditional gradient operators, three dataset specific $3 \times 3 \times 3$ masks are required, covering the current voxel's 26-neighbourhood.

Unlike the masks of the Sobel and other standard operators, in the case of non-isometric data, the mask weights differ between the different mask orientations. This is due to the fact that the same intensity difference corresponds to a different intensity gradient if the distance over which it occurs changes, as it does along the different axes in a non-isometric data grid. As such, unique weights are required for the masks in each of the three orthogonal orientations in order to arrive at consistent gradient vector components in each direction.

The eight 26-connected neighbours (see Fig. 3.5) are all equidistant from the centre voxel ($\sqrt{dx^2 + dy^2 + dz^2}$), but the true distance to each of the individual 6 and 18-connected neighbours varies, depending on the direction of travel. A 6-connected neighbour can be offset from the centre voxel by any one of dx , dy , or dz , while the offset to an 18-connected neighbour is one of $\sqrt{dx^2 + dy^2}$, $\sqrt{dx^2 + dz^2}$, or $\sqrt{dy^2 + dz^2}$, depending on position.

In order to simplify the presentation, in developing the gradient operator for non-isometric data, we initially consider the 2-D case, and extend our treat-




Fig. 3.5: A voxel’s 26-neighbourhood in a rectilinear data grid: 6-connected neighbours share a face (a), 18-connected neighbours share an edge (b), and 26-connected neighbours share a corner (c). Thus a voxel has six 6-connected, twelve 18-connected, and eight 26-connected neighbours.

ment to three dimensions subsequently. Thus we first consider a 2-D data array with pixel spacing dx and dy . With this pixel size the three pixel configurations illustrated in Fig. 3.6 all demonstrate a uniform gradient field. The gradient magnitude is the same across all three regions but the direction differs in each case. Taking into account the pixel dimensions, we can see that all demonstrate a unit length gradient magnitude in terms of greylevels per unit distance in the direction of maximum gradient.

$-dx$	0	dx	0	$\frac{dxdy}{\sqrt{dx^2+dy^2}}$	$\frac{2dxdy}{\sqrt{dx^2+dy^2}}$	dy	dy	dy
$-dx$	0	dx	$-\frac{dxdy}{\sqrt{dx^2+dy^2}}$	0	$\frac{dxdy}{\sqrt{dx^2+dy^2}}$	0	0	0
$-dx$	0	dx	$-\frac{2dxdy}{\sqrt{dx^2+dy^2}}$	$-\frac{dxdy}{\sqrt{dx^2+dy^2}}$	0	$-dy$	$-dy$	$-dy$

Fig. 3.6: Three pixel configurations demonstrating local greyscale gradients of equal magnitude but differing orientations. The pixel intensities are given as functions of the pixel dimensions dx and dy such that the intensities equal the perpendicular distance of each pixel from the zero line. Note that what is being depicted here is actual image sections, not filter masks.

If we let $dx = dy = 1$ and apply the straightforward Sobel operator for example, we get a gradient vector of magnitude eight in all cases, the three vectors being $[8, 0]^T$, $[4\sqrt{2}, 4\sqrt{2}]^T$, and $[0, 8]^T$ respectively. As stated in Sobel (1990), that operator yields an estimate that is 8 times as large as the average gradient. In general, the standard gradient operators do not normalise their results as it is usually the case that only the relative rather than the absolute gradient magnitudes are of significance. Thus it makes sense to use the simplest mask configurations possible, ignoring any normalisation factors in favour of concise filter masks.

$-\frac{dx}{3dx^2+dy^2}$	0	$\frac{dx}{3dx^2+dy^2}$	$\frac{dy}{3dy^2+dx^2}$	$\frac{dy^2+dx^2}{3dy^3+dydx^2}$	$\frac{dy}{3dy^2+dx^2}$
$-\frac{dx^2+dy^2}{3dx^3+dx^2dy^2}$	0	$\frac{dx^2+dy^2}{3dx^3+dx^2dy^2}$	0	0	0
$-\frac{dx}{3dx^2+dy^2}$	0	$\frac{dx}{3dx^2+dy^2}$	$-\frac{dy}{3dy^2+dx^2}$	$-\frac{dy^2+dx^2}{3dy^3+dydx^2}$	$-\frac{dy}{3dy^2+dx^2}$

Fig. 3.7: Gradient filter masks for use in non-isometric 2-D data with pixel dimensions dx and dy .

With arbitrary values for dx and dy a more general formulation is needed. The usual pair of static filter masks is replaced by dataset-specific masks where the weights are a function of dx and dy , as shown in Fig. 3.7. To illustrate the consistent operation of these masks in non-isometric data consider applying these gradient masks to each of the pixel configurations in Fig. 3.6. This yields the three gradient vectors $[2, 0]^T$, $[\frac{2dy}{\sqrt{dx^2+dy^2}}, \frac{2dx}{\sqrt{dx^2+dy^2}}]^T$, and $[0, 2]^T$, which all have a magnitude of two, and are each oriented correctly for the pixel configurations given. The expressions for the individual weights in Fig. 3.7 are derived as follows.

Consider a pixel and its 8-neighbourhood. An estimate of the greyscale gradient vector at that pixel's location can be calculated in terms of the eight sub-gradients⁸ observed between the pixel and its eight nearest neighbours. These sub-gradients are each of fixed orientation, along the line connecting each neighbour to the centre pixel, and have magnitudes equal to the intensity difference between the two pixels divided by the distance between these pixels. In order to form the local gradient estimate, $\nabla u = [\nabla u_x, \nabla u_y]^T$, each such sub-gradient is decomposed into x and y components, and the resultant components normalised to account for the non-isometry of the data grid. The normalisation is necessary in order to compensate for the fact that the relative distribution of the 8-connected sub-gradients is not uniform, since they tend to point more in the direction of the larger of the pixel dimensions, resulting in an unequal level of decomposition between the orthogonal directions. Thus the expressions at each position in the x and y masks are constructed as

$$\begin{aligned}
 wt_x &= \left(\frac{1}{\|\delta\|}\right)\left(\frac{\delta_x}{\|\delta\|}\right)(C_x) \\
 wt_y &= \left(\frac{1}{\|\delta\|}\right)\left(\frac{\delta_y}{\|\delta\|}\right)(C_y)
 \end{aligned} \tag{3.2}$$

⁸Calculated as $\frac{\text{Intensity difference}}{\text{Euclidean distance}}$, between the centre voxel and each of its 8-neighbours.

where δ is the offset vector to each neighbour, and C_x and C_y are the normalisation constants. Thus in each of the two lines in Eq. 3.2 the first term derives the sub-gradient magnitude, the second term decomposes it into its x and y components, and the third term corrects for the non-isometry in the data grid, yielding the following masks,

$-\frac{dx C_x}{dx^2+dy^2}$	0	$\frac{dx C_x}{dx^2+dy^2}$	$\frac{dy C_y}{dx^2+dy^2}$	$\frac{C_y}{dy}$	$\frac{dy C_y}{dx^2+dy^2}$
$-\frac{C_x}{dx}$	0	$\frac{C_x}{dx}$	0	0	0
$-\frac{dx C_x}{dx^2+dy^2}$	0	$\frac{dx C_x}{dx^2+dy^2}$	$-\frac{dy C_y}{dx^2+dy^2}$	$-\frac{C_y}{dy}$	$-\frac{dy C_y}{dx^2+dy^2}$

such that each sub-gradient is normalised based on its distance from the centre, and then decomposed into its x and y components. In order to arrive at mutually consistent masks, appropriate values for C_x and C_y are calculated, such that the two sets of mask weights exhibit a ratio equal to $\frac{dy}{dx}$ to compensate for the non-isometric distribution of the samples used in the calculations. This can be achieved by setting the sums of the positive weights in the two masks to $\frac{1}{dx}$ and $\frac{1}{dy}$ respectively, which leads to the expressions in Eq. 3.3.

$$\begin{aligned}
 C_x \left(\frac{2dx}{dx^2 + dy^2} + \frac{1}{dx} \right) &= \frac{1}{dx} \\
 C_y \left(\frac{2dy}{dx^2 + dy^2} + \frac{1}{dy} \right) &= \frac{1}{dy}
 \end{aligned} \tag{3.3}$$

These in turn can be rearranged to give values for the constants C_x and C_y .

$$\begin{aligned}
 C_x &= \frac{dx^2 + dy^2}{3dx^2 + dy^2} \\
 C_y &= \frac{dy^2 + dx^2}{3dy^2 + dx^2}
 \end{aligned} \tag{3.4}$$

By multiplying these normalisation constants into the two masks shown above we arrive at the final mask weights as given at the outset in Fig. 3.7.

This treatment extends naturally into three dimensions where three masks, each of size $3 \times 3 \times 3$, are required. The mask for use in calculating the x

component of the 3-D gradient vector in non-isometric volume data is shown below in Fig. 3.8, where the value for C_x is given in Eq. 3.5. The masks for y and z are formulated similarly.

$\frac{-dx C_x}{dx^2+dy^2+dz^2}$	0	$\frac{dx C_x}{dx^2+dy^2+dz^2}$	$\frac{-dx C_x}{dx^2+dy^2}$	0	$\frac{dx C_x}{dx^2+dy^2}$	$\frac{-dx C_x}{dx^2+dy^2+dz^2}$	0	$\frac{dx C_x}{dx^2+dy^2+dz^2}$
$\frac{-dx C_x}{dx^2+dz^2}$	0	$\frac{dx C_x}{dx^2+dz^2}$	$\frac{-C_x}{dx}$	0	$\frac{C_x}{dx}$	$\frac{-dx C_x}{dx^2+dz^2}$	0	$\frac{dx C_x}{dx^2+dz^2}$
$\frac{-dx C_x}{dx^2+dy^2+dz^2}$	0	$\frac{dx C_x}{dx^2+dy^2+dz^2}$	$\frac{-dx C_x}{dx^2+dy^2}$	0	$\frac{dx C_x}{dx^2+dy^2}$	$\frac{-dx C_x}{dx^2+dy^2+dz^2}$	0	$\frac{dx C_x}{dx^2+dy^2+dz^2}$

Fig. 3.8: Filter mask to calculate the x component of the local greyscale gradient vector in non-isometric volume data.

$$C_x = \frac{1}{\frac{4dx^2}{dx^2+dy^2+dz^2} + \frac{2dx^2}{dx^2+dy^2} + \frac{2dx^2}{dx^2+dz^2} + 1} \quad (3.5)$$

Thus the three orthogonal components of the local greyscale gradient vector, ∇u_x , ∇u_y , and ∇u_z can be calculated by employing the masks obtained in this fashion. We can then proceed to calculate $\|\nabla u\|$, the magnitude of the gradient vector as:

$$\|\nabla u\| = \sqrt{(\nabla u_x)^2 + (\nabla u_y)^2 + (\nabla u_z)^2} \quad (3.6)$$

We are now in a position to proceed with the development of the crucial weighting function $f(\vec{p}\vec{q}, \nabla u)$ from Algorithm 3.1, which is the final step necessary to complete the implementation of the algorithm. A number of filter formulations are possible using the magnitude and direction of the gradient vector. If only the magnitude is used, a nonlinear but isotropic filtering is achieved, so that the amount of smoothing applied can be limited in the vicinity of strong edges (mask a from Fig. 3.3). By introducing directional information from the gradient vector, we can arrive at the anisotropic solution we seek, where smoothing is applied along but not across the isosurface modelled at the current point (mask c from Fig. 3.3).

3.1.3 Distance Metrics

In order to implement these masks we need to be able to calculate a number of distance metrics as illustrated in Fig. 3.9. Consider the data point $p = (p_x, p_y, p_z)$ with associated gradient vector ∇u . Coordinates are taken to have been appropriately scaled by the voxel dimensions dx, dy, dz . We model an isosurface at p as being that plane perpendicular to ∇u and containing the point p . Take $q = (q_x, q_y, q_z)$ to be another data point in the neighbourhood of p , and $\vec{p\bar{q}} = [q_x - p_x, q_y - p_y, q_z - p_z]^T$ to be the vector from p to q .




Fig. 3.9: Distance from a point to a plane

Three related distance measures emerge, each associated with one of the three mask shapes illustrated in Fig. 3.3. They form the sides of the right angled triangle in Fig. 3.9. We call them d_h, d_a , and d_o respectively (after the sides of the triangle: hypotenuse, adjacent, and opposite). The first, d_h , is simply the magnitude of the vector $\vec{p\bar{q}}$.

$$d_h = \|\vec{p\bar{q}}\| = \sqrt{(q_x - p_x)^2 + (q_y - p_y)^2 + (q_z - p_z)^2} \quad (3.7)$$

Using d_h , in conjunction with $\|\nabla u\|$ leads to the nonlinear, isotropic filter, where smoothing is simply reduced in high gradient regions (d_h on its own leads to the linear case — standard Gaussian smoothing). Note that the directionality of the vectors does not come into the calculation, only their magnitudes.

For the anisotropic filter that we seek to implement, we need d_o , the perpendicular distance from q to the plane modelling the isosurface through the current point p . This distance, the shortest distance from the point to the plane, can be calculated as the projection of \vec{pq} onto ∇u , (that is to say the component of \vec{pq} in the direction of ∇u). This is calculated as the dot product of \vec{pq} and ∇u divided by the magnitude of ∇u :

$$d_o = \frac{\vec{pq} \cdot \nabla u}{\|\nabla u\|} \tag{3.8}$$

Recall that the dot product of two 3-D vectors a and b is a scalar quantity defined as: $a \cdot b = a_x b_x + a_y b_y + a_z b_z$. It thus also follows that the boundary suppressing form corresponding to the cylindrical mask of Fig. 3.3b can be arrived at similarly, using the distance metric $d_a = \sqrt{d_h^2 - d_o^2}$. We do not utilise this form in the current work.

3.1.4 Calculation of Mask Weights

We are now in a position to construct filter masks appropriate to any of the various forms that we have described so far. Weights for the simple linear form are calculated as shown in the first line of Eq. 3.9. The nonlinear isotropic filter derives its weights using the formula in line two, while a boundary preserving filter mask can be formed using either of the two expressions shown on line three. The first uses only the directionality while the second also incorporates the gradient magnitude into the calculation.

$$\begin{aligned} f(\vec{pq}) &= e^{-\left(\frac{d_h}{\lambda}\right)^2} && - \textit{linear} \\ f(\vec{pq}, \nabla u) &= e^{-\left(\frac{d_h \|\nabla u\|}{\lambda}\right)^2} && - \textit{nonlinear} \\ f(\vec{pq}, \nabla u) &= e^{-\left(\frac{d_o}{\lambda}\right)^2} \quad \textit{or} \quad = e^{-\left(\frac{d_o \|\nabla u\|}{\lambda}\right)^2} && - \textit{anisotropic} \end{aligned} \tag{3.9}$$

We use this second form, which in addition to applying the desired directionality also allows the mask (Fig. 3.3c) to vary its rate of falloff with distance from the modelled isosurface, so that weak gradients result in a less pronounced directionality in the mask over the neighbourhood covered by the filter, thus increasing the overall level of smoothing applied in low gradient

regions. In the limit, as the gradient strength tends to zero, this amounts to a simple unweighted mean filter, where the entire neighbourhood is covered by the broadening Gaussian peak. At the other extreme, when the gradient magnitude increases and the peak becomes sufficiently narrow, only neighbour points falling directly on the modelled isosurface make any contribution at all, thus enforcing the maximum directionality possible on the filtering results. In discretely sampled data it is of course very likely that none of the neighbouring voxels falls exactly on an arbitrarily oriented plane and so no smoothing happens at all.

The parameter λ controls the overall strength of the smoothing in each case. Large values broaden the Gaussians and result in greater contributions from neighbouring sample points. Small values contract the Gaussians and restrict the effect that neighbours have on the final value at each point. Substituting for d_o in the last expression and simplifying we get the form for the weighting function used in our initial experiments:

$$f(\vec{pq}, \nabla u) = e^{-\left(\frac{\vec{pq} \cdot \nabla u}{\lambda}\right)^2} \tag{3.10}$$

3.2 Elliptic Filter Model

We next examine how the three filter families so far described, the cylinder, sphere and slab merely represent three special cases in a single larger family of anisotropic filters, modelled as Gaussian ellipsoids. Consider taking the spherical mask that corresponds to isotropic filtering (Fig. 3.10c), and stretching it along one of its axes so as to form an egg shaped mask, Fig. 3.10b. Continue stretching this long axis and in the extreme you arrive at a cylinder, Fig. 3.10a.

Alternatively if you again take the sphere and instead compress it along this same axis (or equivalently stretch it uniformly along the other two orthogonal axes) you arrive at a saucer shaped mask (Fig. 3.10d), which in the extreme tends to the planer shape of Fig. 3.10e. This whole family of filter mask shapes can be described in terms of ellipsoids, with the three axes of the ellipsoid parameterising the system. In fact we need only analyse the 2-D equivalent, the ellipse, as we choose to keep two of the axes equal throughout.




Fig. 3.10: Five points along the ellipsoid based mask shape continuum. In all cases, in order to perform the smoothing, the ellipsoid’s vertical axis is aligned to the local gradient vector, and the matched pair of horizontal axes will thus lie in the plane of the modelled isosurface.

Allowing three independent axes would require the definition of a secondary orthogonal gradient metric such that you would effectively have a direction of highest gradient and an orthogonal direction of lowest gradient, with the third orthogonal axis being intermediate between these two, all in terms of the local greyscale gradient in each direction. This further generalisation, corresponding to a secondary gradient metric being defined in the plane of the modelled isosurface, would introduce additional complexity into the calculations for little gain, and would render the resulting filters more complex and difficult to control, so in this work we enforce the constraint that two of the ellipsoid axes remain equal in length at all times, and thus allow the development to proceed in two dimensions.

Consider the equation of an ellipse, $1 = \frac{x^2}{a^2} + \frac{y^2}{b^2}$, where a and b are the two half-axes as indicated in Fig. 3.11. The parameters a and b control the ellipse’s size and shape, which equates to the strength and directionality of the smoothing to be achieved by the corresponding anisotropic filter proposed by this model.




Fig. 3.11: The form of an ellipse

Now if instead of a single ellipse we next specify a whole family of ellipses, with the same centre, and the same shape (the ratio of a to b remains constant), but of varying sizes, we then arrive at the set of ellipses illustrated in Fig. 3.12, where each ellipse is represented by an iso-intensity contour, with the intensities falling off with distance from the centre. The rate of falloff is

again controlled by the familiar exponential model.

This family of ellipses can be represented as $C^2 = \frac{x^2}{\alpha^2} + \frac{y^2}{\beta^2}$, where each individual ellipse's characteristic parameters a and b are given by $a = \alpha C$ and $b = \beta C$. Thus we can specify that α and β are constants defining the family in question, and C selects the individual members within that family. As such C becomes the new (anisotropic) distance metric, used to define the weighting factors as before:




Fig. 3.12: A family of ellipses

$$wt = e^{-C^2} \tag{3.11}$$

Effectively what we are saying is that all the voxels that fall on the surface of a given ellipsoid within the family take the same weighting factor, based on the intensity we have assigned to that particular member of the ellipsoid family. Thus in general more distant voxels receive a lower weighting, but the falloff is anisotropic, achieving the desired directional smoothing effect.

Fig. 3.13 thus characterises the parameter space in α and β , indicating the manner in which the filter size and shape changes with the changing input parameter values, as α and β range over $[0, \infty]$. We can therefore identify the general filter shape that will be achieved for any given (α, β) pair. This parameter space, illustrated in Fig. 3.13, shows the general positions within the space of all the various filter mask shapes previously described.

Next we need to specify the form of C^2 in Eq. 3.11 above. Extending the model already presented in Section 3.1, we wish to define the function in terms of λ , controlling the size, or smoothing strength, and an additional parameter μ , which is used to control the shape of the mask, or degree of directionality of the smoothing. These two parameters are functions of α and β . In choosing a form for λ and μ so as to achieve the desired relationships, we consider two characteristics of the ellipse, its area \mathbf{A} , and its perimeter \mathbf{P} . As indicated in Eq. 3.12, while the area is easily calculated, no simple method for determining the perimeter of an ellipse exists, one must resort to

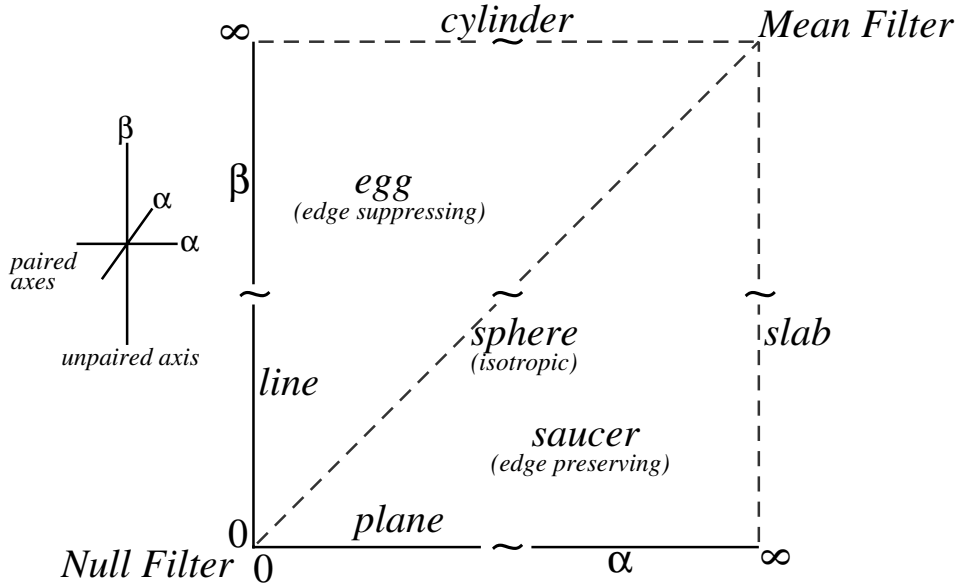


Fig. 3.13: 2-D parameter space for the family of anisotropic filters described. This represents the plane from the full 3-D parameter space where two of the ellipsoid axes (the α 's here) are held equal.

elliptic integrals. Commonly used lower and upper bounds for the perimeter are, however, given. Notice that if $a = b$ then, as expected, both reduce to $2\pi r$, the circumference of a circle.

$$\begin{aligned} \mathbf{A} &= \pi ab \\ 2\pi\sqrt{ab} &< \mathbf{P} < \pi\sqrt{2(a^2 + b^2)} \end{aligned} \quad (3.12)$$

These considerations allow us to design a number of possible formulations for λ and μ that lead to simple and intuitive parameter selection in the final filter, in terms of the familiar size and shape parameters. We examined two of these possible formulations in particular, as shown, along with their corresponding equations for α and β , in Eqs. 3.13 and 3.14.

The size parameter, λ , should relate to some *magnitude* metric in the underlying ellipse model, and should involve positive contributions from both α and β . Thus formulations based on either the area or perimeter of the ellipse would be plausible. The λ in Eq. 3.13 derives from the perimeter upper bound, while that in Eq. 3.14 can be seen as being related to both the perimeter lower bound and the area.

The shape parameter, μ , will clearly involve the ratio of α to β , as this ratio

defines the ellipticity⁹ of the ellipse. By introducing the logarithmic function we achieve symmetrical behaviour in μ . When the ratio of α to β equals one, the ellipse is a circle, the isotropic case. The edge suppressing case corresponds to a range in the ratio of zero to one, while in the edge preserving case, the ratio ranges over one to infinity. By employing the logarithmic function we relocate the isotropic case to the origin, and achieve symmetry in the edge suppressing and edge preserving cases, as illustrated in Fig. 3.14. This insures that the process of choosing and modifying values for the parameter μ remains a straightforward task.

We selected the set given in Eq. 3.14 as providing both the cleanest implementation and the best characteristics in the resulting filter model. This second form for λ in particular is preferable to the alternatives, with the product of α and β resulting in our size parameter λ being a function of the area of the corresponding ellipse in each case. The first form for λ , which is closely related to the upper bound for the ellipse perimeter, results in a less satisfactory relationship between α and β , especially in the case where one becomes much smaller than the other, and thus loses all influence over the final filtering result. The precise forms chosen for λ and μ , involving as they do, square roots, were in fact formulated so as to give the cleanest possible results when substituted into the equation for our final filter model, as shown in Eq. 3.16.

$$\begin{aligned}
 \lambda &= \sqrt{\alpha^2 + \beta^2} & \alpha &= \frac{\lambda}{\sqrt{1 + (e^\mu)^{-2}}} \\
 \mu &= \ln \frac{\alpha}{\beta} & \beta &= \frac{\lambda}{\sqrt{1 + (e^\mu)^2}}
 \end{aligned}
 \tag{3.13}$$

$$\begin{aligned}
 \lambda &= \sqrt{\alpha\beta} & \alpha &= \lambda e^\mu \\
 \mu &= \ln \sqrt{\frac{\alpha}{\beta}} & \beta &= \frac{\lambda}{e^\mu}
 \end{aligned}
 \tag{3.14}$$

The forms for λ and μ given in Eq. 3.14 results in the parameter space illustrated in Fig. 3.14, which again shows the positions of the various filter mask shapes discussed, now within this alternate space. This parameter space can be seen as a transformation of that given in Fig. 3.13, by which the α axis

⁹Ellipticity — The degree of deviation from a circle or sphere of an ellipsoidal shape or path, measured as the ratio of the major to minor axes.

becomes the positive μ axis, the β axis becomes the negative μ axis, and the line $\alpha = \beta$ (isotropic, spherical masks of varying size) becomes the λ axis.




Fig. 3.14: The 2-D lambda-mu parameter space corresponding to the parameter forms given in Eq. 3.14, yielding a pair of filter parameters cleanly controlling smoothing strength and directionality respectively.

From this we can develop the final form of our weighting function $f(\vec{p}\vec{q}, \nabla u)$ first introduced in Section 3.1.1. First, in Eq. 3.15 we define the expansion of C^2 , and then we substitute this back into the formula from Eq. 3.11, in order to give the final form of the anisotropic ellipsoid filter weighting function in Eq. 3.16. The optional $\|\nabla u\|$ term is included, as discussed in Section 3.1.4, adding an additional degree of gradient strength dependence to the operation of the filter. We can see that when $\mu = 0$ the anisotropic term disappears and we are left with an equation identical to the nonlinear, isotropic form first given in line two of Eq. 3.9, in Section 3.1.4.

$$\begin{aligned}
 C^2 &= \left(\frac{x\|\nabla u\|}{\alpha}\right)^2 + \left(\frac{y\|\nabla u\|}{\beta}\right)^2 \\
 &= \left(\frac{d_a\|\nabla u\|}{\lambda e^\mu}\right)^2 + \left(\frac{d_o\|\nabla u\|e^\mu}{\lambda}\right)^2 \\
 &= \left(\frac{\|\vec{p}\vec{q}\|\|\nabla u\|}{\lambda e^\mu}\right)^2 + \left(\frac{\vec{p}\vec{q} \cdot \nabla u}{\lambda}\right)^2 (e^{2\mu} - e^{-2\mu})
 \end{aligned} \tag{3.15}$$

$$f(\vec{p}\vec{q}, \nabla u) = e^{-\left(\left(\frac{\|\vec{p}\vec{q}\|\|\nabla u\|}{\lambda e^\mu}\right)^2 + \left(\frac{\vec{p}\vec{q} \cdot \nabla u}{\lambda}\right)^2 (e^{2\mu} - e^{-2\mu})\right)} \tag{3.16}$$

We now know how to calculate all the terms in Eq. 3.16 and we are thus in a position to calculate the new grey values. The updated value for the current sample point is the normalised mean of the weighted samples as in Eq. 3.17, where i indexes each neighbour (including the current point) from 1 to n , and v_i and wt_i are the sample value and weight respectively at each location. This corresponds to the calculation shown in the last line of Algorithm 3.1, presented in Section 3.1.1.

$$v = \frac{\sum_{i=1}^n (v_i wt_i)}{\sum_{i=1}^n wt_i} \quad (3.17)$$

3.3 Filter Characterisation and Performance

The filtering scheme described above has been utilised in the processing of a large database of test images and volumes. Our particular focus has been on the enhancement of fine branching structures in the data, where the narrowest branches are often significantly degraded by the high noise levels present. Fig. 3.15 illustrates the de-noising of homogeneous regions combined with the preservation of important boundaries.




Fig. 3.15: The unfiltered image and result after 10, 20, and 200 iterations

The image shown is an off axial body cross section demonstrating the upper portion of the liver to the left and the base of the subject's left lung to the right. The high intensity regions within the liver represent a section through the 3-D branching structure that we wish to enhance. Note in particular the small isolated bright (dark in these inverted images) regions that have been well preserved in the filtered images against the otherwise smoothed background of liver tissue. These are cross sections through thin branches in a 3-D tree structure. As the number of iterations applied is increased, the homogeneity within the regions continually improves but the identified boundaries remain strongly defined.

The edge strength level at which boundary preservation gives way to regional smoothing is governed by the value chosen for the scaling parameter λ . This value is selected interactively by the user, based on the particular characteristics of the dataset being processed. The application of a less vigorous smoothing would better preserve the less distinct outer boundaries, at the expense of a higher level of interior noise, requiring a greater number of iterations in order to lead to a comparable level of regional homogeneity.

In Fig. 3.16 we can see a closeup of two fine branches in the data demonstrating varying contrast, and in places being no more than two pixels across. Again the adaptive Gaussian filter has homogenised the background without destroying the integrity of the important boundaries with the neighbouring branch structures. The images in 3.16b show the data under median filtering and help to illustrate the superior edge retention and noise suppression characteristics of the adaptive filtering scheme.

The filter aims to achieve the best possible balance between two competing considerations: greylevel homogenisation within regions and gradient preservation at boundaries. Choice of the parameter λ is governed by the relative importance attributed to each of these two competing factors. The best filter in a particular situation can thus be seen as the one that achieves the highest degree of intra region homogenisation while maintaining a required level of greyscale gradient at the regional boundaries.

In order to characterise the performance of the filter we have graphed boundary preservation against the level of intra regional noise suppression achieved for a number of different filter types (Fig. 3.17). To calculate these metrics we first




Fig. 3.16: Closeup of two ducts under filtering. The adaptive filter has been applied for 20 iterations in this case.

partition the test data into three categories, based on local gradient strengths in the original data, assigning each data point as either strong, intermediate, or weak. This yields a set of typical boundary points in the strong category, and another set of points from regional interiors in the weak category. The intermediate class consists of sample points that are neither strongly boundary nor strongly body points.

The partitioning levels are selected on the basis of the level of edge strengths that we wish preserve in the test data. Typically the strong category consists of the top five percent of gradients present, with the weak category being set to the lowest fifty percent. This reflects the typical composition of a test dataset, with boundary points representing a relatively small fraction of the area covered, and regional interiors accounting for much more of the data. The average percentage change in the local gradients within each category after filtering, is recorded in order to quantify the noise suppression and edge retention achieved in each case.

Each filter is applied for between zero and twenty iterations, leading to twenty one data points on each line plotted. In all cases it can be seen that the early iterations have the most substantial effect, while later iterations produce




Fig. 3.17: Region smoothing versus edge retention for the four filter types, demonstrating the clear difference between adaptive and nonadaptive techniques. Gradient-weighted adaptive approaches are specifically designed to preserve strong boundaries and as such we would expect to see a large separation between the two pairs of approaches, as is shown in this plot.

increasingly little change to the results. Each of the lines originates at the graph origin indicating the unfiltered data with zero percent suppression on both axes. The point at (100,100) corresponds to a completely featureless image where all gradient information has been suppressed. The ideal filter stays as close to the x axis as possible, indicating that edge strengths at the important boundaries have not been degraded, while increasing amounts of unwanted variation within the regions have been eliminated.

In the mean filter we simply replace each greyvalue with the average value within a given neighbourhood. The performance of this technique is predictably poor. Measuring the filtering effect in terms of percentage suppression of the local gradients, as we do in Fig. 3.17, makes particularly clear the indiscriminate nature of this approach. High gradient regions, which thus have a lot to lose, demonstrate very rapid decay in their gradient magnitudes in this case. The linear filter characterised on the first line of Eq. 3.9 (of which the mean filter is just a special case, with $\lambda = \infty$) performs similarly. Reducing the value of λ mainly just slows the rate of decay, but ultimately edges are preserved no better than in the case of the mean filter.

The nonlinear form (the second line in Eq. 3.9) takes account of the gradients and as such produces far superior results. However, examination of the output images in Fig. 3.18 demonstrates the failing of this method. The regional boundaries, while well preserved remain very noisy. The anisotropic form (Eq. 3.9, line three) by contrast achieves similar levels of edge preservation and noise suppression but additionally manages to produce greatly smoothed regional boundaries, ideal for subsequent identification in a segmentation phase. Again the level at which preservation gives way to smoothing is controlled by the input parameter λ . In the case shown in Fig. 3.18 the aim was to preserve as much as possible of the visible ductal elements within the liver, representing branches of the pancreato-biliary system.




Fig. 3.18: Results after 20 iterations of the four filter types examined: mean, linear, nonlinear, and anisotropic filters

Further testing, alongside more advanced smoothing techniques, has also been performed, and is fully documented in two papers (Lynch et al., 2004, Ghita et al., 2005a). The results shown in Table 3.1 are taken from Lynch et al. (2004), and illustrate the excellent performance of our technique when compared with that of a number of the most widely used existing adaptive approaches.

The filtering scheme presented represents an effective and versatile approach to adaptive noise suppression where the preservation of weak region bound-

aries is required. By building on the framework of Gaussian filtering we have retained a simplicity and elegance that allows the filters described to be visualised and tuned simply and systematically and as such the adaptive Gaussian filter represents a particularly attractive option for application in large scale studies where automation and simple adjustment are crucial.

	Photographic Image			MRCP Image		
	SD	Edge Height	Edge Width	SD	Edge Height	Edge Width
Unfiltered	57.4	31	2.26	277.7	219	2.04
Gaussian	41.0	15	4.40	102.8	196	2.16
Savitzky-Golay	40.8	23	2.50	61.23	158	2.48
Chen	24.2	26	2.13	42.99	211	2.00
Perona & Malik	27.7	25	2.17	69.63	214	2.00
Robinson	31.9	30	2.17	35.05	219	1.99

Table 3.1: Results achieved applying five smoothing techniques to two test images. Lower standard deviations are better, indicating more smoothing within homogeneous regions of the images. Larger edge heights are better, indicating better preservation of edge strength at semantically important edges. Smaller edge widths are better, indicating less blurring or spread at semantically important edges.¹⁰

The algorithms offer good performance, coupling excellent filtered results with fast execution, especially when utilising lookup tables for the expensive mathematical operations required. They typically need only a very few iterations in order to achieve the desired levels of filtering. The general form is also open to significant further enhancement without an increase in algorithmic complexity, by extending the current treatment to include further scaling and shaping of the contributing masks. The results achieved demonstrate the approach to be robust in the presence of noise and to be capable of retaining small but semantically important features in the data. This thus represents an excellent preprocessing filter for the preparation of noisy data for further analysis.

¹⁰This table of results is taken from Lynch et al. (2004). Full details of this series of comparative tests can be found in that paper.

Chapter 4

Greyscale Reconstruction

This chapter describes a morphological approach to the reconstruction of fine branching structures in 3-D data (Robinson and Whelan, 2004*b*), developed from the basic procedures of reconstruction by dilation. We address a number of closely related questions arising from this reconstruction goal, including issues of structuring element size and shape, noise propagation, iterated approaches, and the relationship between geodesic and conditional dilation. We investigate and assess the effect and importance of these considerations in the context of the overall reconstruction process, and examine the effectiveness of the approach in addressing the task of reconstructing narrow branch features in noisy, non-isometric volume data. We also describe the *downhill filter*, a novel, optimal algorithm for the implementation of reconstruction by dilation first presented in Robinson and Whelan (2004*a*), and we examine the generalisation of this algorithmic pattern to the implementation of a class of related image processing procedures, which we call *directed filtering*. This work contributes to the development of a computationally efficient processing framework within which to construct an interactive and user friendly computer assisted diagnosis (CAD) environment suitable for the routine, realtime screening of MRCP data.

4.1 Reconstruction by Dilation

The classical reconstruction by dilation procedure (Vincent, 1993, Soille, 1999) is an effective and much-utilised image processing tool applied extensively in

the segmentation and classification of complex scenes (Salembier et al., 1996, Araujo et al., 2001, Metzler et al., 2001, Angulo and Serra, 2003). Seeded regions are retained while neighbouring unseeded regions are attenuated to the intensity level of the surrounding background data. The approach yields excellent results in isolating compact structures in noisy data. However, when the regions of interest are less compact, including fine branching elements, the approach performs less well, especially in the presence of noise. This behaviour is due to the geodesic growth properties at the heart of the definition of reconstruction by dilation. The geodesic dilations that constitute a reconstruction by dilation guarantee that there exists a connected, strictly downhill (in terms of pixel intensity) path to each sample point from one of the original set of seed points that were used to initialise the procedure. This property is what achieves the suppression of non-seeded high intensity regions.




Fig. 4.1: Narrow branch preservation in hybrid reconstruction. The detail in the original branch (a), is severely compromised by the standard geodesic reconstruction by dilation (b), while our hybrid reconstruction approach (c), preserves the detail more faithfully.

The difficulty arises when a narrow element is encountered in a seeded region. Any signal drop-off along the narrow branch, due to noise or to transitory signal reduction, can result in an undesirable attenuation of the intensity level along the entire remainder of the branch length, as illustrated in Fig. 4.1. This is not an issue in the reconstruction of more compact regions as there will exist some convoluted high intensity path to carry the signal past any blockage. As the features in the region to be reconstructed become more and more narrow, the chances of encountering signal drop-offs that can not be negotiated at the higher signal intensity level, increase. In the case of fine branches, where the high intensity path is only one or two pixels wide, the likelihood of undesirable

signal suppression becomes extreme, leading to incomplete reconstruction of the desired objects.

In order to counter this difficulty we have developed a non-geodesic extension to the reconstruction by dilation procedure aimed at bridging small gaps in the high intensity path, while still effectively suppressing the signal intensity in neighbouring non-seeded regions. The approach has the additional desirable property of preserving more fully the texture information in the reconstructed regions and suppressing the stepped contour effects and intervening flat zones that often otherwise manifest. These properties can be particularly beneficial in terms of both the analysis and the visualisation of the processed data.




Fig. 4.2: Illustration of the ductal tree whose isolation and analysis is the ultimate goal. Neighbouring high intensity regions, primarily gastrointestinal fluids, complicate the task.

The particular motivation for this investigation stems from the nature of the visualised biliary tree signal, as observed in the studies contained within our MRCP database. Our aim is the isolation of the ductal system from within this data and as such we wish to suppress the signal from neighbouring structures while at the same time faithfully preserving the signal in the tree itself. Fig. 4.2 shows an MRCP dataset in maximum intensity projection (MIP) rendering,

where the ductal tree is clearly visible along with a number of occluding high intensity structures that we wish to suppress. Successful isolation of the finer branches towards the periphery of the tree in particular can be dependant on the effective suppression of high intensity proximal structures in the scene.

4.1.1 Morphological Dilation Operators

First we briefly review the definition of reconstruction by dilation, and the differences between dilation, conditional dilation, and geodesic dilation (Serra, 1982, Soille, 1999). Reconstruction by dilation is a morphological procedure that can be defined similarly in the binary and greyscale cases in terms of their respective geodesic dilation operators. In the binary case, reconstruction by dilation can be seen as the application until stability of the geodesic dilation operation, $(\delta^{(1)}I) \cap C$, where $\delta^{(1)}$ is the fundamental dilation, I is the marker and C the mask image, and \cap is the boolean intersection operator. In terms of the greyscale morphology this becomes $(\delta^{(1)}I) \wedge C$, where \wedge , the point-wise minimum operator replaces \cap , the boolean intersection operator.

The definitions presented in Table 4.1 explain symbols used in the treatment to follow. Provided definitions are for the explicit variables used throughout and for operators whose meaning may not be immediately obvious. The mathematical notation used throughout this section is based on set theory and the predicate calculus, and follows the formal specification style described in Woodcock and Loomes (1988).

The size N dilation $\delta^{(N)}$ of an image I by a given structuring element defining a neighbourhood N_G , can be represented as in Eq. 4.1, where p and q are pixels and D represents the image domain under consideration.

$$\delta^{(N)}I \hat{=} \forall p : D \bullet \forall q : N_G(p) \bullet (q \leftarrow \max(p, q)) \quad (4.1)$$

Conditional and geodesic dilations are developments of the basic dilation principle that stem from the idea of using an additional image as a conditioning mask that selectively limits the dilation process (see Fig. 4.3). This is achieved by defining the final image as the point-wise minimum of the basic dilation with this conditioning mask, such that the conditional operator may only change a

Symbol	Definition
I	– The mask image
J	– The marker image
D	– The domains of I and J (which must be equal)
$I[p]$	– The p^{th} pixel in image I
$N_G(p)$	– The neighbours of pixel p on the grid G
$\delta^{(N)}I$	– Dilation of size N , applied to I
$\delta_C^{*(N)}I$	– Conditional dilation of size N , applied to I , conditioned on C
$\delta_C^{(N)}I$	– Geodesic dilation of size N , applied to I , conditioned on C
$\mathbb{R}_C I$	– Reconstruction by dilation of I , conditioned on C
$\mathbb{R}_C^{h(N)} I$	– Hybrid reconstruction of strength N , of I , conditioned on C
$\forall a : B \bullet P(a)$	– Universal quantification: for all a in B predicate $P(a)$ is true
$\{a : B \mid P(a)\}$	– Set Comprehension: the set of all a in B that satisfy predicate $P(a)$
\equiv	– Is equivalent to
$\hat{=}$	– Is equal to by definition
\leftarrow	– Assignment
$=$	– Equality
\Rightarrow	– Implication
$($	– List catenation
<i>head</i>	– Returns the first element in a list
<i>tail</i>	– Returns a list minus its head
\triangleright	– Range subtraction: removes a specified element from a list
<i>squash</i>	– Closes up holes in a list (e.g. left by range subtraction)

Table 4.1: Symbol definitions.

pixel in the final image to an intensity level less than or equal to that in the conditioning image.




Fig. 4.3: Using binary geodesic dilations to segment a brain image; (a) the original binary data, (which is also used as the conditioning mask), (b) after an erosion of size three, (c) the components of interest are selected by labelling of the largest regions, (d) a geodesic dilation of size four reconstructs the brain surface, (e) after 50 geodesic dilations much unwanted material is being brought back into the segmentation.

In Eq. 4.2 we give a definition for the process of conditional dilation $\delta_C^{*(N)}$ in terms of the standard dilation $\delta^{(N)}$, and in Eq. 4.3 we formalise the definition already given in the text for geodesic dilation $\delta_C^{(N)}$. In both cases, I is the input image and C is the conditioning image.

$$\delta_C^{*(N)}I \cong \delta^{(N)}I \wedge C \tag{4.2}$$

$$\delta_C^{(N)}I \cong \delta^{(1)}I \wedge C \dots N \text{ times} \tag{4.3}$$

Figs. 4.4 and 4.5 provide an illustration of the three kinds of dilation. Standard dilation as in Fig. 4.4c simply grows regions by the specified amount. We use the symbol $\delta^{(N)}$ to specify a standard dilation of size N . This is equivalent to N applications of the fundamental dilation operator, $\delta^{(1)}$.




Fig. 4.4: A comparison of standard, conditional, and geodesic dilation using an elementary 1-D, 2-connected structuring element, after five iterations (top row), and iterated until stability (bottom row), as applied to a 1-D greyscale signal; (a) the conditioning mask, (b) the starting signal, (c) standard dilation, (d) conditional dilation, (e) geodesic dilation. The conditioning mask is overlaid in grey onto all plots to clarify the results.

Conditional dilation, Fig. 4.4d, grows regions by the specified amount and then removes any growth that is not sanctioned by the conditioning mask. What differentiates geodesic dilation, Fig. 4.4e from conditional dilation is that in the geodesic case the intersection with the conditioning mask must be applied after each fundamental dilation whereas in the conditional case dilation is performed to the desired final level before the intersection with the mask is calculated. In other words, for a geodesic dilation of size N , the compound operation $(\delta^{(1)}I) \wedge C$ must be applied N times. This illustrates why it requires

multiple iterations in order to achieve a geodesic dilation of any given size, while a conditional dilation of the same size can be achieved in a single iteration by dilating once with the appropriate size of structuring element prior to calculating the intersection with the conditioning mask.




Fig. 4.5: A comparison of standard, conditional, and geodesic dilation using an elementary 2-D, 8-connected structuring element, iterated until stability, as applied to a greyscale image. Pixel intensities are inverted for clarity.

The image in Fig. 4.5e shows the geodesic dilation iterated until stability and as such is by definition the reconstruction by dilation of the image in 4.5b conditioned on the image in 4.5a. Likewise, the plot in Fig. 4.4e on the second row shows the reconstruction by dilation of the signal in 4.4b conditioned on the signal in 4.4a. Study of the plot reveals the ‘downhill’ nature of the process, which informed the approach taken in developing our optimal reconstruction by dilation algorithm and gave it its name of *downhill filtering*. Moving away from the maxima in the starting signal, the filtered signal falls with the conditioning signal, but once it has fallen it does not rise up again as the conditioning signal rises. This behaviour means that high intensity regions that are marked will be retained while those that are not marked will be attenuated to the level of the highest valley separating them from any marked regions. This is the characteristic behaviour that defines geodesic dilation.

4.2 Hybrid Reconstruction

As the caption in Fig. 4.5 states, the structuring element used in that example is an elementary 2-D, 8-connected structuring element. If the extents of the structuring element used in the dilation process reach beyond the innermost shell of sample points surrounding the origin, the filter is no longer geodesic

and cannot be used to perform reconstruction by dilation in the traditional sense of its definition. Dilation with a non-elementary structuring element is effectively equivalent to the application of more than one elementary dilation between each application of the minimum operator as in Eq. 4.4. The design and manipulation of structuring elements is an important topic in this field (Ji et al., 1989, van Droogenbroeck and Talbot, 1996, Park and Yoo, 2001), and proves valuable in the development of our procedure here.

$$\delta_C^{h(N,M)} I \cong \delta^{(N)} I \wedge C \dots M \text{ times} \quad (4.4)$$

The hybrid dilation of Eq. 4.4 has the potential to achieve the behaviour that we wish to utilise in our reconstruction approach, as it will allow the dilation to extend beyond small regions of intensity dropout, without breaching the more extensive low intensity valleys between disconnected neighbouring regions. The more dilations applied per application of the point-wise minimum operator, the wider the gaps that the reconstruction can cross.

Thus we can see that there exists a family of reconstructions for any given starting data, where the optimal solution can be chosen in terms of how much physical separation exists at the point of closest proximity between seeded and unseeded regions in the data. So long as this measure allows sufficient scope to bridge the gaps in the fine branch components of the seeded regions, the reconstruction goal can be successfully achieved. We can thus define a new hybrid reconstruction by dilation procedure as in Eq. 4.5, where the parameter N governs the size of the neighbourhood to use. This specifies which particular member of the family of reconstructions mentioned above is achieved, and thus how wide are the gaps across which the reconstruction will move in isolating the fine branch features within the seeded regions.

$$\mathbb{R}_C^{h(N)} I \cong \delta^{(N)} I \wedge C \dots \text{to completion} \quad (4.5)$$


This hybrid reconstruction allows us to preserve the finer ducts visible in the biliary trees of the coronal HASTE datasets that we are examining. As such, it provides an efficient approach to the selective suppression of non-relevant high intensity features in the processed data volumes, thus allowing us to focus more effectively on the true regions of interest within each scan.

In our experiments we applied both standard reconstruction by dilation using 6, 18, and 26-connected structuring elements (the three fundamental 3-D structuring elements leading to geodesic reconstructions), and we also applied a series of reconstructions utilising larger structuring elements. These larger elements were constructed so as to achieve approximately isometric reconstructions in the non-isometric volume data that we were analysing. The data is isometric in the x and y directions, with voxel dimensions of approximately $1.3mm$ each way, but in the z direction the voxel dimensions increase to $4.0mm$. Thus in order to achieve dilation more consistently in all directions, an anisotropic (in voxel terms) approach was preferred, so as to compensate for the non-cubic nature of the data. We found this to be the most effective approach, maximising the amount of unconstrained dilation we could use between applications of the minimum operator before the procedure starts to include unwanted structures in the reconstruction.

4.2.1 Reconstruction Results

We processed our database of volumes using both traditional geodesic reconstruction by dilation and our hybrid reconstruction approach applied at varying strengths, and assessed the reconstruction results achieved in each case. Fig. 4.6 illustrates the superior intensity preservation characteristics of the hybrid reconstruction approach in the processing of objects of interest that include fine branching features. The level of retention achieved increases with the strength of the hybrid reconstruction applied.

We performed the series of reconstructions and then measured the degree of intensity suppression in the neighbouring unseeded regions and in target branches of varying widths within the seeded regions. Fig. 4.6 shows the variations in signal drop-off observed at two different levels of our hybrid reconstruction, along with standard reconstruction by dilation. In this way we were able to demonstrate the enhanced level of reconstruction achieved using large anisotropic structuring elements. Eventually, as the size is increased beyond the optimal, the signal intensity in neighbouring regions begins to pick up until, in the extreme, the reconstruction approximates the original unfiltered data, with only the highest intensity peaks in the data being reduced to the level of the highest intensity sample points present in the original seed data.



(a) Original (b) Geodesic (c) 38 Anisotropic (d) 54 Anisotropic

Fig. 4.6: Closeups from a volume dataset showing branch tips at various levels of reconstruction demonstrating both fine and course branches: (a) original unfiltered data, (b) 6-connected geodesic reconstruction by dilation, (c) reconstruction using an anisotropic 38-neighbour structuring element, (d) reconstruction using an anisotropic 54-neighbour structuring element.

On the top row in Fig. 4.6 a single branch of the biliary tree is visible, originating at the top of the image section, along with signal from the subject's right kidney in the lower right quadrant of the image section. All three reconstructions have successfully suppressed the signal from the kidney, but the degree of preservation of the branch varies considerably, improving successively from left to right in the three reconstructions shown. In the case of the 54-element hybrid reconstruction, the duct is virtually unaltered from its unfiltered version. On the bottom row we see a number of ducts intersecting the plane of the image at various angles. Once again the degree of preservation improves successively from left to right across the three reconstructions, illustrating the expected operation of the hybrid reconstruction procedure.

In Fig. 4.7 we illustrate the relative suppression of unseeded and narrow branch seeded features by examining the greylevel change at each pixel before and after filtering. We wish features other than the biliary tree to show up strongly in the difference maps, indicating the greatest levels of suppression, while the branches of the tree, we wish to be absent, indicating no change from the original data. While unseeded regions are strongly suppressed in both cases, we can see that the hybrid approach has caused almost no deterioration in the

branches of the biliary tree, while in the case of the geodesic reconstructions significant suppression is evident.





Fig. 4.7: Reconstruction difference images after geodesic and hybrid reconstruction respectively: (a) the original data, (b) original minus geodesic reconstruction, (c) original minus hybrid reconstruction. Circles in (b) indicate unwanted suppression of biliary tree ducts in the geodesic case, which is not observed in the equivalent hybrid reconstruction difference images (c), to the right.

In Fig. 4.8 we can, in addition, observe the enhanced texture retention properties of our hybrid reconstruction approach, where the second row of images achieved using traditional reconstruction by dilation demonstrate excessive smoothing and the introduction of sharp graduations within the reconstructed tree, while the images on row three show superior preservation of the fine detail from the original data (shown on the top row). This can be of particular importance for the accurate interpretation of the final data by the radiologist.

We also observe the role that noise in the data plays in propagating the high intensity signal across background valleys. Once the approach departs from the geodesic scheme where a strict uphill intensity path is always retained between any point and an original seed region, isolated high intensity noise peaks in the background regions have the potential to piggyback the signal



(a) Volume Study 1

(b) Volume Study 2

Fig. 4.8: Maximum intensity projections of two of the datasets from our study, performed on the original (top row), geodesic reconstructed (middle row), and hybrid reconstructed (bottom row) data volumes.

across the valleys like a series of stepping stones. This effect makes strong salt and pepper noise particularly unfavourable in the application of our technique. The nature of the noise distribution typical to our data makes the approach more applicable in this data, as even with very strong dilations, the degree of the unwanted propagation is kept to a manageable level. This is due to the intensity and spatial spread present in the signal noise, which means that the maintenance of a high intensity steppingstone path across valleys of any significant width becomes extremely unlikely.

By extending the basic principles of reconstruction by dilation beyond the geodesic case, we have presented a hybrid reconstruction technique specifically designed to optimally reconstruct objects containing fine branching structures in the source data, while still effectively attenuating the signal from neighbouring unwanted high intensity structures. Through the application of these techniques we have developed an effective and efficient image processing procedure that yields superior reconstruction results as a precursor to both further automated segmentation, classification, and analysis, and enhanced and simplified manual review of the data by the radiologist.

4.3 Downhill Filter

A number of optimisations and algorithmic efficiencies have been detailed in the literature for reconstruction by dilation and similar procedures in both binary and greyscale morphology, including structuring element decomposition (Park and Yoo, 2001) and manipulation (Sivakumar et al., 2000, van Droogenbroeck and Talbot, 1996), flat zones (Salembier and Serra, 1995), interval coding (Ji et al., 1989), and the use of ordered pixel queues (Vincent and Soille, 1991, van Vliet and Verwer, 1988, Vincent, 1993). These enhancements notwithstanding, the procedure remains computationally expensive and highly data dependant. We have developed an alternative approach to the problem that achieves the same filtering effect in a single pass through the data and is as such both fast and linear time in its execution.

Our novel and efficient *downhill filter* algorithm is an optimal implementation of the reconstruction by dilation procedure. Its development grew out of our investigation into the topic, during which we assessed existing implemen-

tations and enhancements presented in the literature. As such, the *downhill filter* represents an elegant and efficient single pass equivalent to the more computationally expensive algorithms for morphological greyscale reconstruction by dilation that have been previously reported. In principle reconstruction by dilation consists of the iterative application of the geodesic dilation operator, applied until stability. Our algorithm achieves this filtering effect through a process of region growth by ordered aggregation of surface pixels onto an expanding shell. In this section we present the new algorithm, along with a review of the traditional reconstruction by dilation approaches, and assess the relative computational performance of the different methodologies. Results are presented on 2-D and 3-D, synthetic and real world data. We conclude this section with a brief discussion of the extension of this computationally optimal algorithmic pattern to a family of related image processing procedures. We call this algorithmic pattern *directed filtering*.

4.3.1 Existing Algorithms

In order to both confirm the operation and gauge the performance of our approach, we implemented the four algorithms described in Vincent (1993) for the 2-D 8-connected and 3-D 6-connected cases of greyscale reconstruction by dilation. The details of these four approaches can be found in that paper. Brief outlines are given here for reference:

A Standard Technique: This algorithm follows the “iterate to completion” scheme outlined earlier in this chapter while describing the basic procedure for reconstruction by dilation. Combined elementary dilation and point-wise minimum operations are iterated until stability.

B Sequential Reconstruction Algorithm: Here again we iterate until completion, but the number of iterations required is reduced significantly by alternating raster order and anti-raster¹¹ order scanings of the data, and by allowing changes so far in the current iteration to be included in the calculations as they are performed.

C Reconstruction Using a Queue of Pixels: In this case, a FIFO queue is initialised with all the boundary pixels from the regional maxima in

¹¹Starting from the end of the image data array, and working back to the beginning.

the marker image. Then, pixels are removed, their neighbours examined, changed, and added to the queue as required. Processing continues in this fashion until the queue is empty.

D Fast Reconstruction Algorithm: This algorithm combines features of the previous two. Its main core proceeds in the same fashion as algorithm C, but it avoids the costly initialisation step of calculating regional maxima. Instead it initialises the queue during a single pass of the raster/anti-raster scheme of algorithm B, where the anti-raster phase is modified to identify and record the necessary pixels in order to initialise the FIFO queue.

As with the third and fourth algorithms above, the downhill filter operates on a pixel queue. However, instead of a FIFO queue, a random access queue is implemented in order to allow the processing of pixels in an optimal order, thus guaranteeing that every pixel is addressed only once in the course of the algorithm's execution. This approach yields a filter that is extremely efficient, and perhaps more importantly, is data insensitive in its execution time, depending only on the size of the dataset being processed. This is in contrast to the above four algorithms all of which are highly data dependant. Folded or rolled up structures in the input image, for instance, seriously compromise the execution speeds achieved by all four approaches so that no guarantees can be given as to the processing time required in the general case. Our algorithm exhibits no such level of variability in its execution speed.

The underlying principle behind the algorithm can be stated in the following terms. Given that it is known which pixels have so far been finalised, it is possible to finalise the next pixels thus. Find the highest valued finalised pixel that has one or more non-finalised neighbours. Each neighbour may now be finalised to be the lesser of: its mask value and the value of the aforementioned finalised pixel. This process is iterated until no non-finalised pixels remain. Effectively this can be thought of as having a shell of pixels around the finalised regions that expands in a controlled fashion, resulting in the desired reconstruction.

We present the development of the algorithm in two stages. In the first case we restrict the nature of the marker image provided as input to the filter. This simplifies the formulation of the algorithm and represents the form in

which we originally developed it, reflecting the particular requirements of our application in the reconstruction of the biliary tree in MRCP data. Our initial approach is to strengthen the usual precondition by stating that each pixel in the marker is either equal to the corresponding pixel in the mask or it is equal to zero: Eq. 4.6.

$$\forall p : D \bullet ((J[p] = I[p]) \vee (J[p] = 0)) \tag{4.6}$$

In the general formulation of reconstruction by dilation the precondition is presented in a weaker form (Eq. 4.7) so as to only insist that the marker image is pixel-wise less than or equal to the mask image.

$$\forall p : D \bullet (J[p] \leq I[p]) \tag{4.7}$$

Our strengthened form covers a very useful subset of the general reconstruction problem where some seed regions are extracted directly from the mask image in order to initialise the reconstruction, and everything else in the marker is initialised to zero. Later we relax this precondition to its more general form and extend our algorithm correspondingly, allowing the marker image pixels to take intensity values less than or equal to the corresponding pixels in the mask. This latter more general form is sometimes useful or necessary, for instance in the calculation of the h -domes¹² of an image (Vincent, 1993).

¹²The h -dome image $D_h(I)$ of input image I is given by $D_h(I) = I - \mathbb{R}_I(I - h)$. That is, I minus the reconstruction by dilation of $I - h$, conditioned on I . Thus the scalar value h governs the minimum height of the maximal regions which are isolated in the process.

4.3.2 The Strengthened Form

In Algorithm 4.1 below, the reconstruction is determined directly in the marker image J . Initially each pixel in J is either equal to the corresponding pixel in I or it is equal to zero, (therefore by definition all non-zero pixels in J are at their final value). Let m be the maximum value of the pixels in J . We note that no pixel in the final image can take a value greater than m . We maintain m lists. As a pixel is set to its final value it is placed in the list corresponding to that value. An examination of the efficient handling of the lists and the tracking of which pixels are finalised at any point is beyond the scope of this discussion. The implementation of the algorithm is provided in Appendix B. Examination of this code will illustrate our solution. For now we proceed in general terms. Initialisation consists of placing each non-zero pixel in J into its appropriate list. Following initialisation, processing proceeds starting with list m and proceeding down towards list 1. While the current list is not empty, the next element is removed and its neighbourhood is examined. For each neighbour that has not already been finalised, J is set equal to the lesser of the current list number and the value in I at this location, and the neighbour is added to the corresponding list.

Algorithm 4.1: *reconstruction by downhill filter (strengthened precondition):*

- Find m , the maximum pixel value in image J

$$\{m : \mathbb{N} \mid m \in J \wedge \forall p : D \bullet (m \geq J[p])\}$$
- Place each non-zero pixel in J into its appropriate list
$$\forall p : D \bullet (J[p] \neq 0 \Rightarrow L[J[p]] \leftarrow L[J[p]] \dot{\cup} p)$$
- Process the m lists from high to low:
 - For* $n = m..1$
 - While* $L[n] \neq \emptyset$
 - $p \leftarrow \text{head}(L[n])$
 - $L[n] \leftarrow \text{tail}(L[n])$
 - For* $q \in N_G(p)$
 - If* $I[q] > 0 \wedge \text{!finalised}(J[q])$
 - $J[q] \leftarrow \min(n, I[q])$
 - $L[J[q]] \leftarrow L[J[q]] \dot{\cup} q$

4.3.3 Relaxing the Precondition

In Algorithm 4.2, with the weakened precondition on J it can no longer be stated that all nonzero pixels in J are at their final value to start with. We still initialise in the manner described above, but we must now cater for the fact that we may need to remove pixels from one list (where they were initialised) and place them onto another list (where they are finalised). As before, the reconstruction is determined directly in the marker image J . Initially each pixel in J is less than or equal to I . Let m be the maximum value of the pixels in J and as before we maintain m lists. Initialisation consists of placing each non-zero pixel in J into its appropriate list. Following initialisation, processing proceeds starting with list m and proceeding down towards list 1. While the current list is not empty the next element is removed and its neighbourhood is examined. For each neighbour that has not already been finalised, J is set equal to the lesser of the current list number and the value in I at this location, and the neighbour is added to the corresponding list. If it was already in a list it is also removed from that location.

Algorithm 4.2: *reconstruction by downhill filter (standard precondition):*

- Find m , the maximum pixel value in image J
 $\{m : \mathbb{N} \mid m \in J \wedge \forall p : D \bullet (m \geq J[p])\}$
- Place each non-zero pixel in J into its appropriate list
 $\forall p : D \bullet (J[p] \neq 0 \Rightarrow L[J[p]] \leftarrow L[J[p]] \dot{\cup} p)$
- Process the m lists from high to low:
For $n = m..1$
While $L[n] \neq \emptyset$
 $p \leftarrow \text{head}(L[n])$
 $L[n] \leftarrow \text{tail}(L[n])$
For $q \in N_G(p)$
If $J[q] < \min(n, I[q])$
If $J[q] \neq 0$
 $L[J[q]] \leftarrow \text{squash}(L[J[q]] \dot{\cup} \{q\})$
 $J[q] \leftarrow \min(n, I[q])$
 $L[J[q]] \leftarrow L[J[q]] \dot{\cup} q$

4.3.4 Algorithm Performance Tests

We implemented the five algorithms discussed and conducted a number of tests using a wide range of image and volume datasets. A 2-D version of each algorithm was written for tests conducted using the image data, and a 3-D version of each was prepared for use with the volumes. The first set of implementations utilised the 2-D 8-connected neighbourhood, while the second set used the 3-D 6-connected neighbourhood.

We also went on to develop a general purpose version of the downhill filter that accepts an arbitrarily sized and shaped structuring element to allow for maximum flexibility. As described earlier in this chapter, we conducted tests with larger non-symmetrical neighbourhoods in order to examine the potential for stepping over narrow regions of signal void in the images and volumes, which might otherwise fracture the connected regions that we wish to isolate. These subsequent experiments culminated in the development of the hybrid reconstruction approach already described in Section 4.2.

Validation tests

In assessing the five algorithms, our first step was to perform a series of validation tests to confirm that all five were both well-formed and correctly implemented. We applied each function to a range of test images and volumes and confirmed that all produced identical results pixel-for-pixel. Algorithm A, which directly implements the definition of reconstruction by dilation using iterated geodesic dilations, was taken as providing the baseline results for our comparisons. The output generated by each of the other algorithms was compared against that of algorithm A and we thus confirmed that all algorithms performed correctly.

Fig. 4.9 shows before and after images for three of the datasets used in our tests. Seeded regions are retained while other high intensity regions are attenuated. In Fig. 4.9a the seed was placed at the outermost end of the spiral arm in the top right hand sector of the image, and the four unseeded regions (the border, the unmarked interlaced spiral, and two small regions at the centre of the image) can be seen to have been eliminated. In Fig. 4.9b the seed was placed in the bright band half way across and a third of the way down

the image, in the U-bend of the looped region that is the target object. The three unconnected horizontal bars are thus removed in the filtered result. In Fig. 4.9c the seed was placed in the broad descending trunk of the tree structure we wished to isolate and again the unconnected high intensity regions in the image are eliminated as required.




Fig. 4.9: Filtering results on three test images demonstrating the elimination of unseeded regions of high intensity signal in the images while retaining signal connected to the original markers. Arrows in the reconstructed images indicate the locations of the seed points used in these reconstructions.

Looking back to the top two rows of Fig. 4.8 we can see two volumetric datasets rendered in maximum intensity projection, again each one before and after filtering. As the data is 3-D, some regions of high intensity signal that appear in the renderings to be in contact with the seeded region but that are in fact in a different plane can be seen to have been eliminated. This serves to illustrate how the filter can be used to clear the field around the objects of interest in a volume allowing subsequent examination to better focus on the true regions of interest.

Execution timing tests

Once the basic implementation validation procedures had been completed, we moved on to perform execution timing tests in order to characterise the relative performance of the five approaches on a range of data. Tables 4.2 and 4.3 document the results of these tests for five image and five volume datasets respectively. The methods are labelled alphabetically in correspondence with the labelling given in Section 4.3.1, with the downhill filter being added to the classification as method E. For comparison purposes we note that all tests were performed on a 1.8GHz Pentium 4 with 512MB of RAM.

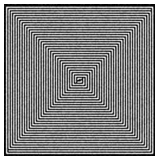

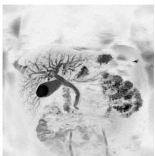
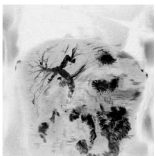

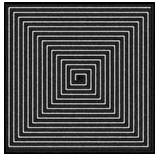

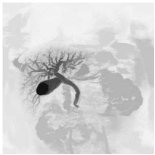
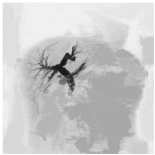

Before	Spiral	Loop	Study 1 MIP	Study 2 MIP	Blank
Method					
A	47898.0 (7191)	2884.3 (409)	1892.1 (270)	1735.0 (247)	700.0 (130)
B	59.4 (17)	29.7 (8)	37.5 (9)	28.6 (7)	6.2 (2)
C	175.0	42.2	40.6	37.5	14.0
D	139.1	43.5	28.2	21.8	4.6
E	10.9	9.4	9.3	9.4	7.8
After					

Table 4.2: Execution timings for the five algorithms tested, as applied to five test images. All images are of dimensions 256×256 . Times are in milliseconds. Bracketed numbers indicate iterations required in methods A and B. Methods C, D, and E are queue based rather than scan based and as such do not yield an iteration count.

In Table 4.2 we present execution times in milliseconds for each of the five algorithms as applied to each of five test images. As expected the standard technique, method A demonstrates extremely poor performance relative to all the other methods. Method E, the downhill filter demonstrates a clear superiority in execution times across a range of data, with only the blank image yielding better times in any of the other algorithms due to the higher initialisation overheads involved in method E. In all cases the image size is 256×256 pixels, and the data is eight bit.

Note also that the times for method E vary by only a couple of milliseconds across the entire range of test images (a variation of less than 40%), while the most consistent of the rest (method B) shows variations of close to a factor of ten in its execution times, and the largest intra-method variation approaches a 70 times difference between best and worst. This reflects the guaranteed single pass nature of the downhill filter in contrast to the effectively unbounded nature of the other approaches. We attribute the small variations that do exist in the execution times for method E to varying caching-related overheads incurred on the hardware platform employed, based on the differing order in which pixels are addressed throughout the image from one dataset to the next.

The first test image consists primarily of two interlaced spirals and was intended most especially to stress methods A and B, which both address the image pixels in scan order and as such are particularly prone to poor performance when faced with any kind of wrapped-up structure in the input image. Methods C and D also fare particularly badly in this case due to the indiscriminate way in which pixels are added to the FIFO queue for processing, which can result in pixels being processed several times before their final value is reached. Method E, by contrast, demonstrates little variation in its execution time for this image as compared with the other test images illustrating its stability across all data configurations.

The next three images, the synthetic loop and the two maximum intensity projections, represent data in a more usual range of image structure, illustrating the typical performance to be expected of each of the tested methods. Once again method E performs most efficiently, while method A, as expected, fails to approach the performance of the other algorithms.

The final image is a constant mid-grey and represents the most trivial form of input. With a single pixel at the centre of the image used as a marker, method A still requires as many iterations as the most distant pixel in the image in order to reach completion. Methods B, C, and D all demonstrate their ability to take advantage of the simple configuration and all achieve low execution times, with B and D actually outperforming E for the first time. Method E itself continues to perform well, with its single pass execution path once more in evidence, maintaining a consistent performance.

Table 4.3 shows results, recorded in seconds, for five of the test volumes used in the study. Here the variations are even more pronounced than with the image data, as once again timings on the standard technique, method A, illustrate how slow this method can be, while the downhill filter again achieves the best execution times recorded, across a broad range of datasets.

Method	3-D Spiral	3-D Roll	Study 1	Study 2	Blank
A	104868.9 (255003)	6051.1 (13643)	403.0 (742)	227.5 (428)	137.0 (370)
B	189.0 (542)	7.5 (20)	8.9 (17)	12.5 (24)	0.7 (2)
C	58.4	53.7	23.6	21.5	2.6
D	54.2	54.1	14.7	12.6	0.5
E	2.5	1.9	3.4	3.3	1.2

Table 4.3: Execution timings for the five algorithms tested, as applied to five test volumes. All volumes are of dimensions $298 \times 298 \times 60$. Times are in seconds. Bracketed numbers indicate iterations required in methods A and B. Methods C, D, and E are queue-based rather than scan-based and as such do not yield an iteration count.

Again the five datasets chosen for the study cover a number of levels of complexity. All volumes were of the same size ($298 \times 298 \times 60$ voxels), and all consisted of eight bit data. The first two were synthetic volumes consisting of interlaced spirals of different configurations designed to stress the algorithms. Method A takes close to 30 hours, and over a quarter of a million iterations to fully process the first of these datasets. Method E does the same job in just two and a half seconds.

Volumes three and four were the same volumes whose maximum intensity projections were used in the image tests, and were included to represent data within the usual expected input range for the filters. As before, a blank, constant mid-grey dataset was used to test the most trivial case of input data, resulting again in methods B and D outperforming the downhill filter for the only time.

We have thus shown that the downhill filter offers an efficient alternative to existing reconstruction by dilation approaches to region isolation. It exhibits a number of clear advantages over the other approaches to reconstruction by dilation examined in this study. The most important of these advantages may be summarised as follows:

- 1 **Short execution time:** The downhill filter proved to be faster than any of the other four algorithms tested in all but the most trivial of cases (reconstruction on a blank image), where the algorithmic overheads dominate the execution time.
- 2 **Consistent execution time:** The downhill filter is insensitive to extremes of structure in the input data. Unlike the other algorithms tested, execution times are not dependant on the nature of the data processed. This makes the downhill filters behaviour predictable as well as fast.
- 3 **Simple formulation:** The concise nature of the downhill filtering algorithm means that it is implemented in fewer lines of code than are required for the other algorithms discussed. The implementation given in Appendix B occupies just 40 lines of C code.

This demonstrates the downhill filter’s excellent performance, and in that context shows reconstruction by dilation to be an attractive and highly usable filtering tool ideal for many image processing tasks, and suitable for use in time-sensitive, real-time applications.

4.3.5 Directed Filtering

The algorithmic pattern first implemented in the form of the downhill filter can be generalised to form an optimal solution to a class of algorithms that share a specific set of characteristics. We call this generalised algorithmic pattern *directed filtering*. In concrete terms the following is required:– given the final values for some subset of the sample points in a dataset, it is always possible to determine the final values for at least some of the as yet unfinalised “neighbours” to these sample points. If this is the case, then so long as the neighbourhood used in defining a sample point’s neighbours is exhaustive in its coverage (see Fig. 4.10), and some starting set of finalised sample points can be determined in order to initialise the process, the pattern of directed filtering is applicable.

Image processing techniques that are suitable for the application of this algorithmic pattern include reconstruction by dilation (as investigated extensively above), the grassfire distance transform, the watershed segmentation algorithm and an object skeletonisation procedure. The basic algorithm consists of an

initialisation phase followed by a single, ordered pass through the data, during which each sample point is finalised.




Fig. 4.10: Exhaustive and non-exhaustive neighbourhoods. The black dot indicates the structuring element's origin in each case. Using the neighbourhoods illustrated in (a) any pixel can be arrived at from any other. Thus the whole image can be covered starting with any initial subset of pixels. With the neighbourhoods in (b) full coverage can not be guaranteed.

The term *directed filtering* refers to the ordered fashion in which sample points are addressed as the algorithm proceeds through this main phase of processing. Processing commences at the highest (or lowest) value and proceeds down (or up) until the farther intensity extreme in the data has been reached. In effect, we utilise a random access queue, but the important feature is the controlled manner in which we add new elements to the queue. This is what achieves the directed (downhill or uphill) nature of the filter's operation and guarantees completion in a single pass through the data, as explained in Section 4.3.2.

Thus as we have already seen, in the case of reconstruction by dilation the marker image dictates the highest value that can appear in the final output. From here we proceed down through successive intensity levels finalising pixels to each level as we go (see Algorithm 4.1). Similarly in the case of the grassfire distance transform for instance, we start with a binary image and we can immediately determine that all *figure* (or *foreground*) pixels with a *ground* (*background*) neighbour must take a distance value of one. This is the initialisation phase. Then in the processing phase we continue, increasing the distance incrementally and finalising at each stage those unfinalised pixels contained in the neighbourhood of the previously finalised pixels. The result is a distance map as illustrated in Fig. 4.11, indicating the distance for each *figure* pixel to the nearest *ground* pixel in the image. If a 4-neighbourhood is used, then a city block distance map is achieved, and if an 8-neighbourhood is used, then a chessboard distance map is achieved (see Table 4.4).

<i>Chessboard distance</i>	=	$Max(dx, dy)$
<i>Euclidean distance</i>	=	$\sqrt{dx^2 + dy^2}$
<i>City block distance</i>	=	$dx + dy$

Table 4.4: Given two points A and B such that $\vec{AB} = [dx, dy]^T$, then we define the three standard distance metrics, *Chessboard*, *Euclidean*, and *City block*, such that the *Chessboard* and *City block* distances provide lower and upper bounds respectively for the true *Euclidean* distance from A to B .




Fig. 4.11: The grassfire transform yields a distance map indicating the chessboard distance from each pixels to its nearest background neighbour. The images above show a segmentation taken from a coronal MRI brain scan and the greylevels in the distance map have been inverted for greater clarity.

An implementation for the grassfire distance transform, which was used to generate the results shown in Fig. 4.11, is included in the code section of Appendix B. This provides a second concrete implementation (along with the implementation for reconstruction by dilation, also included in Appendix B) for the algorithmic pattern of *directed filtering*. As can be seen in the code, this implementation examines the 8-neighbourhood of each pixel and as such calculates a chessboard distance map, as described above. The implementation provides a compact and efficient rendering of the transform and illustrates further the general form taken by the directed filtering algorithmic pattern.

Chapter 5

Implementations

In this chapter we provide an overview of the implementation framework within which the preprocessing techniques described in the main body of this thesis have been applied. We present an outline of the complete MRCP processing pipeline developed during the course of this work, in order to provide a well formed context from which to view the application of the specific algorithms and techniques that have been examined in detail in the preceding chapters. To this end we give a step-by-step description of the processing applied to an MRCP dataset from initial DICOM input to final processed results as presented for review to the end user. We highlight the place and importance within this framework of the main processing steps whose investigation have formed the body of this thesis. We also highlight where appropriate additional areas of application that we have investigated in the course of our research, applying our tools and techniques to the processing of data from sources other than MRCP, in order to illustrate the wider applicability of these data processing approaches. Beyond our main focus of pre-segmentation filtering for MRCP data, we have already mentioned in previous chapters the application of these techniques to the processing of whole body MRI data. In addition to this, the work on adaptive smoothing has been applied to a wide range of data types, taken from both medical and non-medical sources, where it has demonstrated impressive levels of performance (Lynch et al., 2004, Ghita et al., 2005*a*). While a full presentation of our work conducted in the analysis of whole body MRI for the measurement and localisation of body fat tissue is given in Appendix A, here we simply illustrate the application to WB-MRI data, of the specific techniques that form our main focus in this thesis.

5.1 Processing Framework

We have examined a number of key elements required in order to implement an effective data preprocessing scheme so as to address the particular difficulties encountered in the analysis of MRCP data, as a precursor to the application of computer assisted diagnosis techniques. In order to achieve our goal, we have implemented a data processing pipeline that takes raw MRCP data and generates a filtered dataset suitable for further analysis and visualisation. This sequence of preprocessing steps is specifically designed to overcome the difficulties associated with the particular data characteristics demonstrated by MRCP data.

Volumetric reconstruction

The data that initially comes from the MRI scanner is presented as a large collection of individual data files saved in the DICOM format. Each file contains the data for a single slice and typically all the files from several studies¹³, each normally composed of a number of multi-slice series are saved into a single location. Thus the first task is to sort through the files and recombine all the studies and series present. To this end we must examine the DICOM header information stored in each file so as to be able to determine how to group the individual slices into volumes. The DICOM fields indicated in Table 5.1 are read from each file and used to correctly reconstitute all the volumes present. A database is thus built up that contains a list of all the studies available, and within each study, the individual series that were acquired. We are then in a position to load, review, and process any series from the DICOM database that has been constructed.

¹³A **study** refers to all the MR data acquired during one visit by a subject to the scanner. A study consists of one or more series. A **series**, in turn, consists of one or more individual MR images, acquired in one go. A series can consist of a single image (e.g. a RARE series), a sequence of images forming a volume (e.g. a HASTE or TRUFI series), or a set of individual images (e.g. a scout series, typically of three axial and three coronal views, widely spaced, and used to guide the acquisition of further series).

Field Name	Field Description
StudyUID	Identifier string unique to this study
SeriesUID	Identifier string unique to this series
Columns	The slice width in pixels
Rows	The slice height in pixels
InstanceNumber	The slice number within the volume
PixelSpacing[0–1]	Pixel x and y dimension in millimetres for this slice
SliceThickness	Thickness of the slice in millimetres
ImageOrientationPatient[0–5]	Two 3-D vectors parallel to slice rows and columns
ImagePositionPatient[0–2]	The x, y, z coordinates of series origin in global space

Table 5.1: DICOM header fields used by our reconstruction routines.

Initially we examine the raw data volumes and identify those volumes (typically the coronal HASTE data series present) that we wish to process further. The data as collected is invariably non-isometric in the inter-slice direction. At this stage it can be resampled onto an isometric grid for easier review and interpretation, typically using a standard tri-cubic interpolation scheme (Fig. 5.1), which locally fits a cubic interpolation model at each position in order to perform the data resampling required. The results of such a data interpolation process are illustrated in Fig. 5.2, which shows an axial view through a coronal HASTE dataset before and after interpolation.




Fig. 5.1: The cubic interpolation model fits a cubic curve ($y = ax^3 + bx^2 + cx + d$) through four consecutive sample points in the data and thus allows an interpolated value y to be calculated for any given x .

While such interpolation procedures are useful for visualisation purposes, their application will usually be left until after any other processing steps have been performed. While such processing will often therefore have to take account of the non-isometric nature of the data, this complication is outweighed by the smaller size of the input dataset and the consequent faster execution times achieved using the un-interpolated data. Since any interpolation scheme can only *estimate* intensity values for the missing data based on the original, a well designed non-isometric algorithm applied to the original data will always perform at least as well as an equivalent isometric algorithm operating on interpolated data.

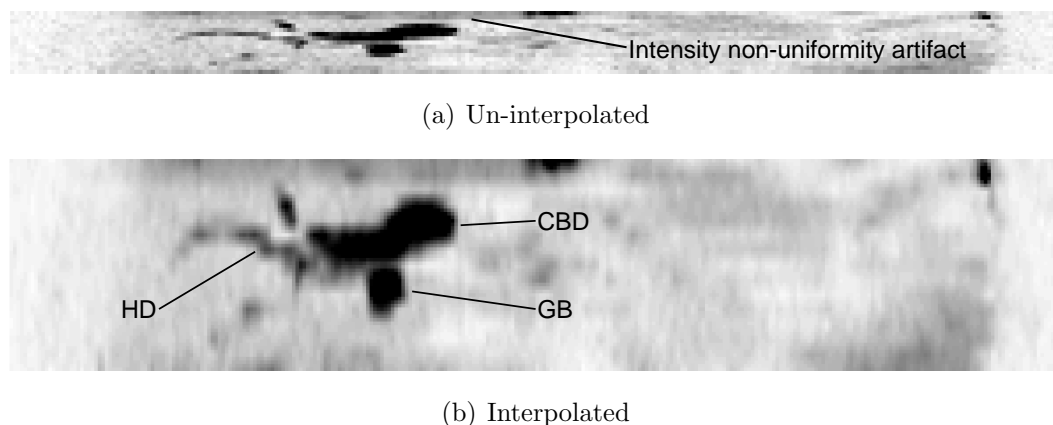


Fig. 5.2: The effect of tri-cubic interpolation in review data. The same axial view through a coronally acquired fifteen slice HASTE MRCP dataset is shown before (a), and after (b), the data has been interpolated in order to achieve iso-metric voxels. The common bile duct (CBD), gall bladder (GB), and hepatic duct (HD) are all depicted in axial cross-section. Notice the intensity non-uniformity effect visible at the top of both views.

Intensity non-uniformity correction

The first phase of data preprocessing involves the correction of greyscale inhomogeneities in the data. As discussed in Chapter 2, coronal HASTE MRCP data often demonstrates an inter-slice greyscale non-uniformity effect. This effect can be clearly observed in Fig. 5.2, which depicts an axial cross-section through the fifteen coronal slices in a typical coronal HASTE dataset. Elevated intensity levels, particularly in the region of the liver tissue, can be seen in the top several rows of these cross-sectional views, corresponding to an intensity non-uniformity effect across the first several slices of the coronal data, exactly as described in Chapter 2. As previously discussed, this step is also required in the processing of WB-MRI data. The effect is more clearly illustrated in this type of data as demonstrated in Fig. 5.3 where uncorrected and corrected data is shown for comparison.

Thus, after the initial reconstruction phase, the next step in the processing pipeline is to apply the greylevel homogenisation procedure of Chapter 2, matching the data histograms between successive slices in the data volume. Due to the characteristics of the MRI acquisition protocols utilised in the collection of the MRCP data, intensity non-uniformities appear and need to be compensated for. Thus we developed this histogram-based intensity non-uniformity correction scheme, in order to minimise the effect of these greyscale

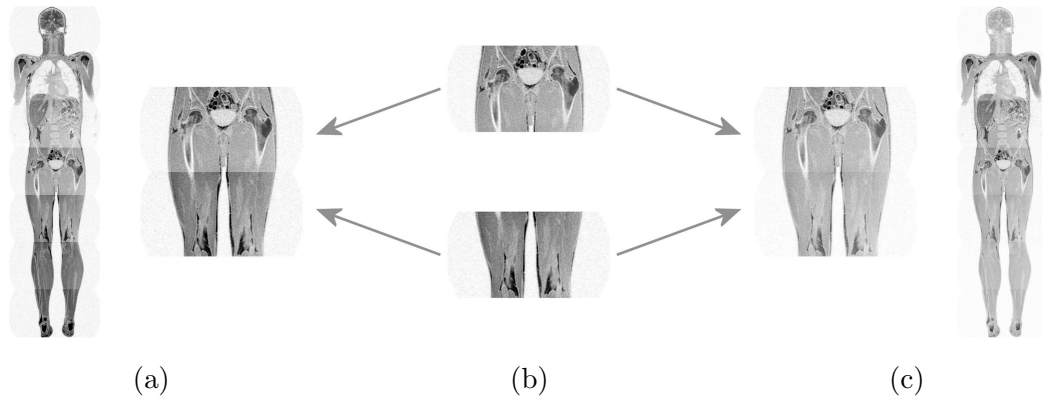


Fig. 5.3: Intensity non-uniformity correction in multi-section whole body MRI data. The original coronal sections (b) may be reconstructed into the final whole body dataset either with (c) or without (a) the application of an intensity non-uniformity correction procedure.

inhomogeneities and thus ensure the best possible intra-regional homogeneity across the full extents of each dataset.

The principal novelties that distinguish our approach to this inter-slice and inter-volume intensity non-uniformity correction task lie in the robust mode detection procedure and the histogram preserving nature of the sample redistribution phase. We designed the robust global mode detection approach in order to provide superior noise immunity in the process of locating the two most significant modes in a noisy multi-modal data histogram, resulting in the stable and consistent operation of our correction procedure. Then, by avoiding the introduction of extraneous new maxima and minima into the data histogram during the process of sample redistribution, we assist in the operation of subsequent histogram-based analysis operations.

Adaptive smoothing

The intensity non-uniformity correction phase of the processing pipeline is followed by a gradient-weighted adaptive Gaussian smoothing step designed to reduce noise, while at the same time preserving semantically important boundaries in the data. In this phase of the process we focus attention on the goal of achieving a significant reduction in the considerable noise present in the MRCP data, and to this end we developed a 3-D, gradient-weighted, adaptive filtering scheme based on the Gaussian smoothing model. Fig. 5.4 demonstrates the operation of this smoothing filter on one slice from an axial MRCP dataset.

This approach is designed to operate in non-isometric data, and attenuates signal noise, while at the same time preserving well the semantically important discontinuities present in the data.




Fig. 5.4: An image from a thirteen slice axial HASTE dataset, unfiltered and after twenty iterations of the smoothing filter have been applied.




Fig. 5.5: Adaptive smoothing applied to three types of MRI data. The filter is easily tuned to the particular characteristics of the data being smoothed, taking into account the levels of noise present and the boundary strengths observed at the interfaces between important regions in the data.

In Fig. 5.5 we illustrate the operation of this filter on whole body, heart, and brain MRI data respectively, highlighting the wider application of the technique. The versatile nature of the oriented, scaled and shaped ellipsoid mask used in this smoothing filter provides a flexible smoothing operator that can achieve the desired levels of regional smoothing and boundary preservation

under a wide range of data conditions. The two tuning parameters, λ and μ provide a simple, comprehensive, and intuitive method of control, with λ governing the overall smoothing strength while μ controls the degree of directionality to be applied, from edge suppressing through isotropic and on to edge preserving at the other extreme.

Thus the gradient-weighted adaptive Gaussian smoothing filter provides a powerful de-noising operator based on an elegant and easily tuned anisotropic ellipsoid mask. Its operation is intuitive and offers a high degree of flexibility and performance within a readily visualised and applied filtering framework.

Greyscale reconstruction

Next a targeted signal suppression technique is applied in order to remove extraneous areas of high intensity signal from surrounding anatomical regions (Fig. 5.6), leading to a considerably cleared field of view in the data volume. Our hybrid reconstruction approach applies a non-geodesic morphological reconstruction procedure to the data in order to suppress the signal originating from neighbouring structures in the scanned volume, while at the same time preserving the signal due to the finely branching tree structure that we wish to segment.



Fig. 5.6: Maximum intensity projections of one of the datasets from our study, performed on the original, geodesic reconstructed, and hybrid reconstructed data volumes.

This goal of retaining the narrow branch features during the morphological processing is addressed by our non-geodesic hybrid reconstruction approach, which was developed as a generalisation of the traditional reconstruction by dilation procedure. The geodesic growth properties that define reconstruction

by dilation are relaxed in order to allow the degree of greylevel connectivity required for a particular reconstruction task to be specified during the reconstruction process. As a result we can avoid the signal drop-off in long narrow features, which is a characteristic feature of the standard approach to morphological greyscale reconstruction.

Again we found useful wider application of these techniques for the morphological reconstruction of seeded regions in MRI data. In Fig. 5.7 a view from an whole body MRI dataset is once again given, providing an illustration of the application of this procedure to that class of MRI data. In this case the primary goal is to suppress the signal originating from the region of the liver, while preserving the widespread and in places weakly-connected body fat signal regions. By seeding the process with the highest intensity portions of the body fat signal, which can be arrived at using a simple threshold, we then preserve connected lower intensity portions of body fat signal, while effectively suppressing the disconnected liver region, whose highest intensities overlap significantly with the lowest intensities observed in the widely distributed body fat regions. In this way we achieve an effective separation of the two classes of tissue within the whole body dataset.

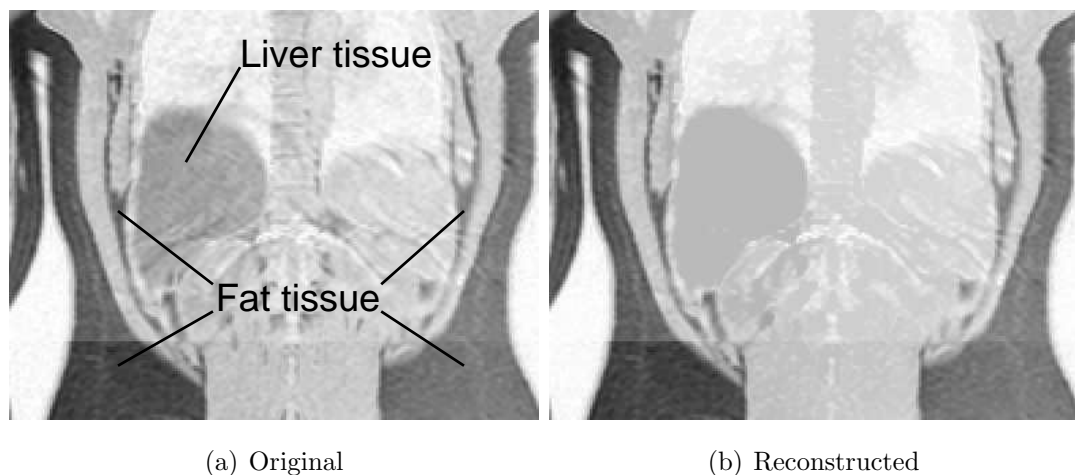


Fig. 5.7: Morphological reconstruction as applied to whole body MRI data so as to suppress signal from the region of the liver while preserving body fat signal. Notice the fat region to the right side of the image, at an intensity level very close to the liver tissue in the original data. After filtering the distinction between the two regions is much clearer.

A similar observation has also been made in relation to the signal observed within the brain regions in these WB-MRI datasets, which once again demonstrate a range of intensities that overlap with those observed in the fat regions that are the target of our processing and analysis.

Biliary tree isolation

The data preprocessing techniques described above have been developed and tested using a large database of MRCP studies covering a wide range of both well-visualised and poorly-visualised pancreato-biliary tracts. The three datasets shown in Fig. 5.8 illustrate coronal HASTE volumes in which the pancreato-biliary system is generally well visualised. The pancreatic duct is only clearly visible in the middle dataset, while the gall bladder is only in evidence in the top volume, but the biliary tree is exceptionally well depicted in all three. Such missing features often exist due to a lack of fluids in the systems, as the protocols on which MRCP studies are based are designed to highlight stationary fluids in the scanned volume. The degree of visualisation of the finer branches in the tree is extremely variable, the three examples here all show good visualisation.

By contrast, in the case of the two datasets illustrated in Fig. 5.9, the pancreato-biliary systems are not much in evidence at all, and what little is visible is all but obscured by the gastrointestinal fluid signal present in both datasets. In the top volume the common bile duct and common hepatic duct are visible. But the hepatic duct disappears above the initial bifurcation. The head of the pancreatic duct is also faintly visible where it approaches the common bile duct at the ampulla of Vater. In the case of the bottom dataset, slightly more of the tree is visible (although the pancreatic duct is not in evidence). The strong proximal signal due to gastrointestinal fluid has, however, presented an additional difficulty, and as can be seen in the surface rendered segmentation results on the right, it has proven impossible to successfully separate the signal from these two sources around the body of the common bile duct, lending extra difficulty to the task of assessing the condition of the common bile duct in this case.

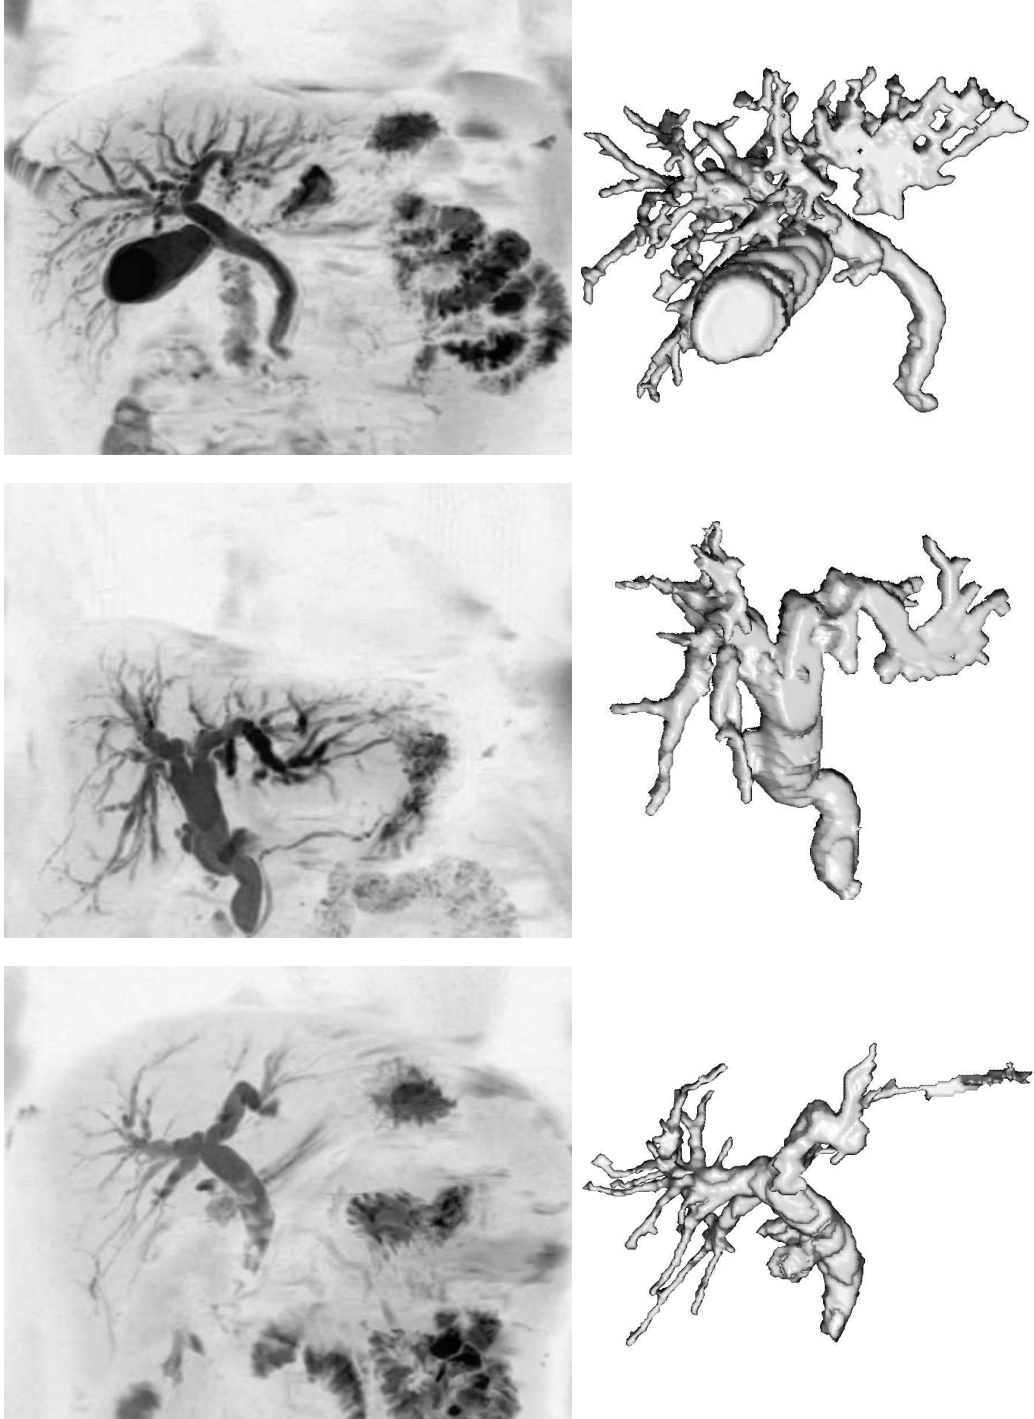


Fig. 5.8: MIP renderings and the corresponding surface rendered watershed segmentations in three datasets with well delineated biliary trees.




Fig. 5.9: MIP renderings and the corresponding surface rendered watershed segmentations of two datasets with poorly delineated biliary trees.

The surface renderings on the right in Figs. 5.8 and 5.9 were generated by first segmenting the data using a marker-based watershed segmentation procedure and then applying an efficient surface extraction algorithm developed from the traditional marching cubes approach. This technique generates a consistent, closed, and oriented triangle mesh (Fig. 5.10) suitable for fast, interactive examination using parallel and perspective projective rendering techniques.

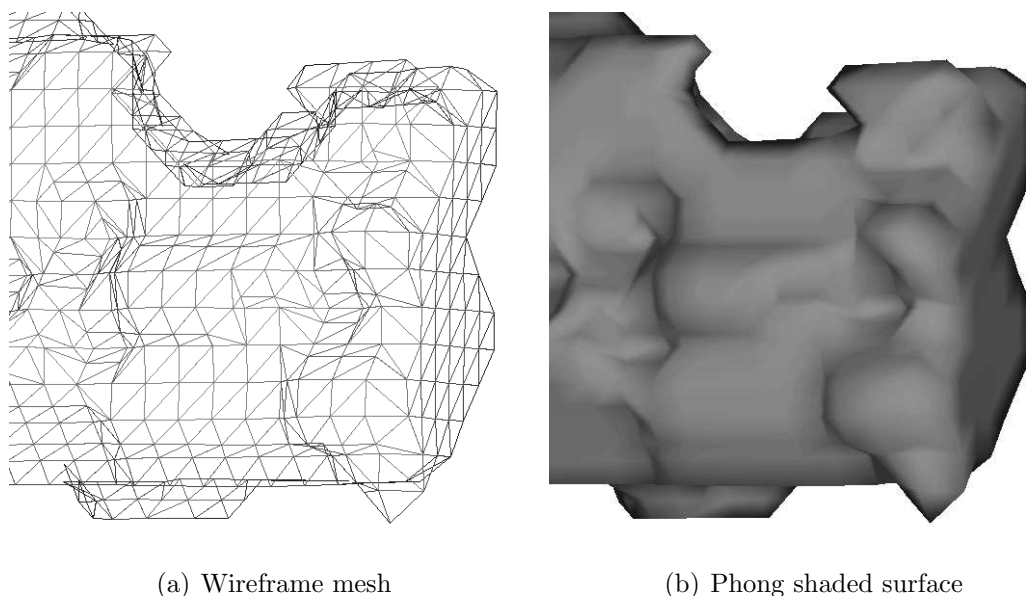


Fig. 5.10: Closeup detail showing a portion of a triangulated surface, as a wireframe mesh with back-face culling (a), and surface rendering using a Phong shading model (b).

Segmented data review and analysis

Once final processing and isolation steps have been performed, tools such as those illustrated in Appendix B can be used in order to visualise and interrogate the data further. An orthogonal sections viewer allows the user to browse the data in the axial, coronal, and sagittal planes, and to identify corresponding locations in these and the surface rendered image views in the various views provided by the tool. This allows for the comparison of different viewpoints onto the same data, as for instance illustrated in Fig. 5.11, where a common bile duct blocked with numerous stones of differing sizes is shown, first in a 2-D coronal slice, and then in a 3-D surface-rendered view, where voids left in the bile signal indicate clearly the positions of the multiple stones that are lodged in the duct. Fig. 5.12 shows the same reverse view of the dataset, this

time in maximum intensity projection. The stones in the common bile duct result in significant shadows in this type of rendering, but it is more difficult to visualise clearly what is causing the obvious shadows within the common bile duct in this type of MIP-rendered view.

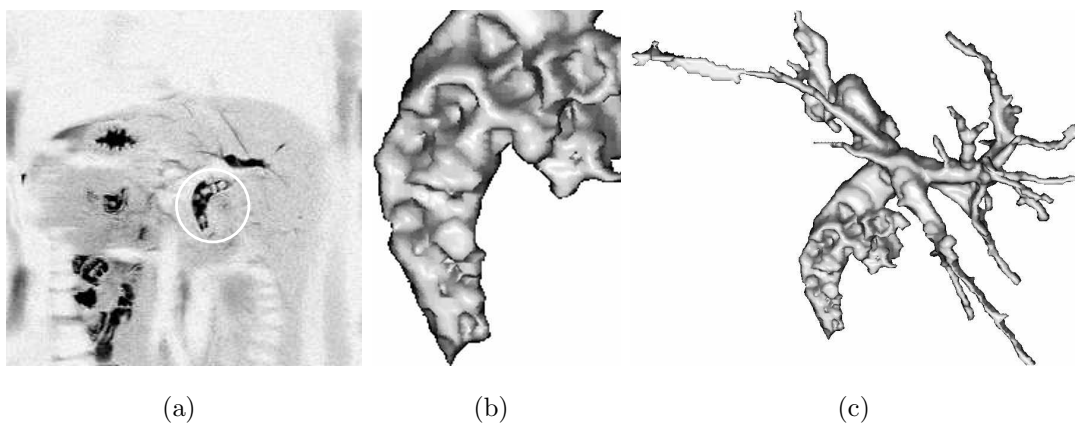


Fig. 5.11: Stones in the common bile duct (circled in (a)), as demonstrated in three dimensions, in posterior surface rendered view, (b) and (c).




Fig. 5.12: Stones in the common bile duct are less well visualised and delineated using a maximum intensity projection rendering scheme. Their presence is still well indicated in this quite extreme case, but in general MIP provides a poor approach for visualising such features.

Interactive feedback can also provide further details to the user as they select regions within the data. In this way the maximum of useful information can be extracted from the acquired data. Volume visualisation tools provide the user with a versatile volume and surface rendering interface to allow for the generation of arbitrary 3-D rendered views from any location and at any angle

situated either outside of or within the segmented tree structure. This provides for the possibility of generating external and virtual endoscopic views, and simulated fly-over and fly-through image sequences, allowing the operator to gain a better appreciation of the condition of the various ducts under examination. In this way a comprehensive set of computer assisted diagnosis (CAD) tools can be developed, based on the enhanced data provided through this set of data preprocessing procedures, leading to an improved level of utilisation of the basic information furnished by an MRCP examination.

Chapter 6

Conclusion

In this chapter we provide a high-level summary of the research work conducted to date and documented in the rest of this thesis, and an overview of the results that have been achieved. We present a discussion of the progress made and the challenges that remain outstanding in the overall task of developing a robust and effective computer assisted diagnosis (CAD) system for the rapid and comprehensive assessment of the pancreato-biliary system through the assistive review and analysis of MRCP data.

Finally we include a complete list of the publications stemming from this work, including journal, conference, and non-peer reviewed material. All of these publications, along with supporting materials can be viewed in pdf format at www.eeng.dcu.ie/~robinsok.

6.1 Summary

In this thesis we have examined the design, development, characterisation, and testing of novel and efficient pre-segmentation filtering procedures designed to achieve increased robustness and precision in the subsequent segmentation and analysis of the biliary tree from MRCP data. Our initial analysis of the MRCP data with which we would be working, identified the key characteristics of this data that would present difficulties in the processing of such MRCP volumes. This examination revealed MRCP to consist of cluttered and noisy, low resolution, non-isometric volume data, often with significant intensity non-

uniformities. We noted the biliary tree to be a fine branching structure that demonstrates extreme variability in the level of its visualisation within the data. These observations resulted in the development of the data preprocessing approach presented in this thesis.

Based on these observations we have designed and developed a focused set of image preprocessing algorithms, so as to facilitate the subsequent operation of non-complex segmentation and computer assisted diagnosis (CAD) procedures. Most notable in this regard are a number of techniques specifically designed to address the key elements of this image processing task, relating to greyscale inhomogeneity, noise, resolution, and signal clutter. The most important of these techniques consist of the following elements.

Firstly a new histogram-preserving approach to inter-image intensity non-uniformity correction addresses the observed lack of greyscale homogeneity in typical HASTE MRCP data without compromising the data histogram of the resulting, intensity corrected MRCP volumes. This improvement in data uniformity results in superior performance of subsequent region classification and analysis procedures.

Next we have provided a highly versatile adaptive smoothing filter implemented as an oriented, scaled and shaped ellipsoid filter mask. This image de-noising filter performs well when compared against similar techniques, and offers a simple and flexible framework within which is implemented a powerful and comprehensive adaptive data smoothing tool.

In addressing the problems encountered due to the widespread high intensity clutter in the data, primarily due to signal from gastrointestinal fluids and often extremely proximal to the signal from the pancreato-biliary system, an efficient new algorithm for morphological reconstruction by dilation called the *downhill filter*, and associated with this, a novel approach to the reconstruction of fine branching structures in noisy volume data, were developed.

Through this combination of flexible and efficient preprocessing algorithms, an effective route towards robust MRCP segmentation and analysis, and routine CAD in the assessment of the biliary tree from MRCP has been developed.

In addition to these major elements, we have also highlighted the extensive development of processing and visualisation tools and utilities, which has ac-

accompanied this work, and the broad range of additional application areas that have benefited from the processing techniques that we have developed, and that have in turn assisted in testing and validating the approaches and the procedures presented throughout this thesis.

6.2 Discussion and Further Work

We have presented in this thesis a data preprocessing approach that is designed to prepare MRCP examinations for the subsequent application of analysis, segmentation, and computer assisted diagnosis procedures. By addressing a number of important issues encountered due to the undesirable features demonstrated in the data produced by current MRCP protocols we have provided a set of tools that enhance the ease of use and the level of usefulness of MRCP for the assessment of the pancreato-biliary system.

In order to facilitate the application of CAD techniques in MRCP data, we have aimed to reduce noise and interfering signal regions in the data, and to increase the overall greyscale homogeneity across the slices of the MRCP volumes. These objectives address the difficulties that traditional CAD procedures encounter in the kind of noisy and low resolution data that MRCP provides. By cleaning and correcting the data in this way, we allow traditional segmentation, classification, and analysis procedures to operate more robustly in the difficult environment presented by MRCP. Only when such standard CAD tools can be applied with confidence, will this kind of data be routinely assessed with the aid of image processing techniques.

The next phase of the work required is therefore the adaption of existing CAD approaches to the specific tasks and challenges presented by MRCP. Colon CAD routines, for instance, typically address such tasks as the search for polyps protruding from the colon wall. However, the high resolution that is achieved in current virtual colonoscopy procedures means that such approaches need significant redesign if they are to be applied to the identification of gall stones in MRCP for instance, where the level of detail resolved in the scanned data is far lower.

The typical role envisioned for CAD techniques in the assessment of MRCP data is quite different to those currently being pursued in the more mature field

of colon oriented CAD research. The much smaller volumes of data generated in MRCP examinations means that the radiologist's main concern is not so much in the thorough coverage of all the acquired data, as it is in the case of CTC, where datasets can run to hundreds of slices, but rather in the enhancement of, and superior visualisation of noisy, and poorly visualised interesting features present within the scanned biliary tree.

In this context it is clear why the maximum levels of data cleaning and homogenisation possible are required in order to make progress in MRCP CAD applications. As such the work presented in this thesis represents an important and significant step along the way towards the goal of effective, routine biliary tree assessment from MRCP.

6.3 Publications Arising

The following list of publications stems directly from the work conducted during this project. All published papers are available for download as pdf documents at: www.eeng.dcu.ie/~robinsok, along with additional code and supporting materials where appropriate. Publications are presented in three groupings comprising journal, conference, and non-peer-reviewed material. Within each group, citations appear in ascending chronological order.

Peer Reviewed Journal Level Publications

Efficient morphological reconstruction: A downhill filter K Robinson, PF Whelan. *Pattern Recognition Letters* **25**(15), 1759–1767, 2004.

MRI diffusion-based filtering: A note on performance characterisation O Ghita, K Robinson, M Lynch, PF Whelan. *Computerized Medical Imaging and Graphics* **29**(4), 267–277, 2005.

Rapid automated measurement of body fat distribution from whole body mri D Brennan, PF Whelan, K Robinson, O Ghita, R Sadleir, J O'Brien, S Eustace. *American Journal of Roentgenology* **185**(2), 418–423, 2005.

Automatic seed initialisation for the expectation-maximization algorithm and its application in medical imaging M Lynch, D Ilea, K Robinson, O Ghita, PF Whelan. *Electronic Letters on Computer Vision and Image Analysis*, Under review.

Peer Reviewed Conference Level Publications

Segmentation of the biliary tree in mrcp data K Robinson, PF Whelan, J Stack. In *Proc. Irish Machine Vision and Image Processing Conference*, Pages: 192-200, Galway, Ireland, 2002.

Non-linear diffusion filtering for the enhancement of mrcp data K Robinson, PF Whelan, J Stack. In *Proc. IPEM Volumetric Analysis of MR Images*, Pages: 16-17, London, UK, 2002.

Identification of body fat tissues in mri data D Ilea, O Ghita, K Robinson, R Sadleir, M Lynch, D Brennan, PF Whelan. In *Proc. Optimization of Electrical and Electronic Equipment*, Pages: 227-232, Brasov, Romania, 2004.

Volumetric reconstruction: Matching and merging in dicom data K Robinson, O Ghita, PF Whelan. In *Proc. Computer Assisted Radiology and Surgery*, Page 1236, Chicago, USA, 2004.

Narrow branch preservation in morphological reconstruction K Robinson, PF Whelan. In *Proc. Irish Machine Vision and Image Processing Conference*, Pages: 42-49, Dublin, Ireland, 2004.

A performance characterisation in advanced data smoothing techniques M Lynch, K Robinson, O Ghita, PF Whelan. In *Proc. Irish Machine Vision and Image Processing Conference*, Pages: 123-128, Dublin, Ireland, 2004.

Measurement and localisation of body fat in whole body mri K Robinson, PF Whelan, O Ghita, D Brennan. In *3rd Annual IEI Biomedical Engineering Research Award*, Dublin, Ireland, 2005 — *Shortlisted paper*.

A 3d cad tool for body fat identification D Ilea, O Ghita, K Robinson, M Lynch, PF Whelan. In *Proc. European Congress of Radiology*, Vienna, Austria, 2005.

Intensity non-uniformity correction in multi-section wb-mri K Robinson, O Ghita, PF Whelan. In *Proc. SPIE OPTO-Ireland*, Dublin, Ireland, 2005.

Analysis of the pancreato-biliary system from mrcp K Robinson, PF Whelan. In *Proc. 18th IEEE International Symposium on Computer-Based Medical Systems*, Dublin, Ireland, 2005 — *SFI Best Student Paper Award*.

Invited and Non Peer Reviewed Publications

Computer aided diagnosis of medical image data in gastrointestinal procedures PF Whelan, R Sadleir, K Robinson, H Fenlon, J Stack, P MacMathuna. In *7th Annual Conference of the Healthcare Informatics Society of Ireland*, Dublin, Ireland, 2002.

Segmentation and analysis of the biliary tree in mrcp K Robinson. In *Biomedical Diagnostics Research Seminar*, Dublin, Ireland, 2003.

Morphological reconstruction: A downhill filter K Robinson. *Research Monograph: RM-03-03*, Vision Systems Laboratory, Dublin City University, Ireland, 2003.

Non-linear noise suppression strategies in medical image processing K Robinson. In *RINCE Research Seminar Series*, Dublin, Ireland, 2004.

Computer assisted diagnosis (cad) for the rapid automated measurement of body fat tissue from whole body mri PF Whelan, O Ghita, K Robinson. In *12th Seminar on Theoretical Foundations of Computer Vision (Seminar No 04251: Imaging Beyond the Pin-hole Camera)*, Dagstuhl, Germany, 2004.

A 3d cad tool for body fat identification O Ghita, PF Whelan, K Robinson, D Ilea. In *Joint Imaging Seminar, Vision Systems Group, Dublin City University, and Trauma & Rehab RRG, Imaging Group, University of Ulster*, Belfast, UK, 2005.

Appendices

Appendix A

Body Fat Analysis

This appendix is adapted and expanded from the paper ‘*Measurement and Localisation of Body Fat in Whole Body MRI*’, shortlisted for the 3rd Annual IEI Biomedical Engineering Research Award, 2005. It details research conducted alongside the MRCP project described in the main body of this thesis and illustrates the wider applicability of a number of the techniques described there.

Abstract—We examine the technical challenges relating to the application of computer assisted diagnosis (CAD) techniques to the quantification of body fat content and distribution in whole body MRI (WB-MRI), addressing in particular the questions of fat as a percentage of total body mass and the medial/lateral fat deposition ratio. Initially the overall body fat content is calculated using a fully automatic four-stage adaptive segmentation procedure. We illustrate the use of interactive visualisation and analysis tools to examine in closer detail the fat distribution in any given area of the body, and look at the task of further classifying the segmented fat volume into the four primary subclasses of body fat: subcutaneous, visceral, intermuscular, and bone marrow. Quantifying these subdivisions can be of particular importance as an indicator in certain medical conditions. We discuss the significance of the presented results in terms of general fitness metrics, in particular for the assessment of obesity, and as an important diagnostic indicator in other conditions, and we highlight the potential for the application of these techniques in routine medical screening and assessment procedures.

I. INTRODUCTION

The ability to quickly and accurately measure and localise body fat content presents significant opportunities in medical diagnostics and screening. The assessment and monitoring of obesity is increasingly recognised as a pressing concern as its incidence and impact grows. Its causes and its effects have been extensively studied and reported in the medical literature where it has been linked to a wide range of conditions including hypertension and coronary heart disease, diabetes, and various types of cancer, as well as to psychological disorders such as depression. As such, a fast and accurate route to routine and comprehensive body fat analysis is an important goal for the medical community, and medical imaging, through a combination of CAD techniques and WB-MRI examinations, offers the potential to achieve this objective.

Many alternative approaches to total body fat estimation exist ranging from the use of calipers, through water or air displacement methods, to x-ray and electrical impedance based techniques. In addition, results are often inferred from indirect metrics such as body mass index (BMI) and waist-hip ratio. However, all of these suffer from one or more of a number of major shortcomings. Many lack accuracy, are time consuming and difficult to perform, require expensive specialised equipment, or involve exposure to ionising radiation. In addition, most fail to offer the ability to localise fat deposition, which is increasingly recognised as an important diagnostic measure, as the ratio of intra-abdominal fat to total fat is an indicator in many conditions. MRI-based techniques also have their drawbacks, the most obvious being cost, as the expense of commissioning, operating, and maintaining an MRI facility remains high. On the other hand the WB-MRI based approach presented here achieves excellent accuracy and localisation, involves short scanning and processing times, and has the potential for use in routine screening as it does not involve exposure to ionising radiation and is contraindicated in only a very few individuals, subjects with heart pacemakers for example. As the availability and use of MRI for medical screening applications continues to grow, procedures such as MR-based whole body fat assessment will become increasingly commonplace, and we believe that the potential benefits to be gained in the widespread use of such a system will only help to accelerate this trend.

In the rest of this appendix we present a technical review of our CAD system for whole body MRI based body fat assessment, and discuss the results achieved so far in the analysis of a database of 42 subjects (21 male and 21 female). We examine the relationship between percentage body fat and the widely used body mass index (BMI), and discuss the shortcomings of the latter as a metric for the quantification of total body fat. We close with a brief outline of the next phase planned for this project, to address the automation of the final fat subclassification task, refining the initial segmentation so as to differentiate between the various types of fat deposits that have been identified. We mention the difficulties that must be overcome in achieving this goal in the framework of a fully automatic CAD system, and suggest a route to tackling these challenges.




Fig. A.1: Five coronal slices from a thirty two slice whole body dataset. Greyvalues have been inverted for clarity, so that fat tissues appear dark and air light, while lean tissues occupy intermediate greylevel intensities.

II. METHOD

Images are acquired in either six or seven slightly overlapping coronal sections depending on the height of the subject. For each section 32 coronal slices are generated covering the full thickness of the body from front to back with a slice thickness of $8mm$. The voxel dimensions are thus $2.02 \times 8.00 \times 2.02mm^3$ in x , y , and z respectively. The standard coordinate system is used, with x running from the subjects right to left, y running from the chest towards the back, and z increasing from feet to head, with the subject lying head first supine (face up) on the scanner table. The slices shown in Fig. A.1 are from a typical study, illustrating the case of a 28 year old normal weight female subject. In this case a voxel matrix of $256 \times 32 \times 937$ has resulted, achieving complete coverage of the body (the volume covered is approximately $517 \times 256 \times 1893mm^3$).

A. Histogram Matching & Data Smoothing

Before the individual coronal sections are recomposed into a single dataset (see Fig. A.4 for details) the greyscale intensity maps for each of the sections must be matched. Greyscale matching is necessary due to the fact that there can be significant intensity offsets between successive coronal sections due to the nature of the MRI acquisition process, see Fig. A.2. It is necessary to minimise these effects in order to optimise the performance of the automated analysis procedures to follow. We achieve this using an approach based on the histogram matching scheme introduced in Chapter 2.

An intensity histogram is constructed for each coronal section (Fig. A.3), the characteristic peak representing soft tissue is algorithmically identified in each case, and the peaks are aligned so as to achieve a matched greyscale distribution across all sections. Identifying the appropriate histogram peaks is a relatively simple matter. More complex is the question of how to scale the individual grey maps so as to preserve a smooth histogram for the overall dataset without stray peaks or voids. This is an important topic, which ensures that subsequent histogram based calculations, for instance in the automatic selection of threshold bands, can be performed robustly.




Fig. A.2: Initial greyscale matching is needed as a preprocessing step in order to facilitate the subsequent segmentation process. Here adjacent coronal sections exhibit a significant grey mismatch on the left, which has been largely eliminated in the normalised right hand image.




Fig. A.3: Unnormalised and normalised histograms for the seven coronal subsections in a whole body dataset. In the lower set of seven histograms the soft tissue peaks have been aligned to achieve the desired normalisation.




Fig. A.4: After greyscale matching is performed, the global coordinate system defined in the DICOM image headers is used to perform vector based volume reconstruction of the whole body MRI data. A transformation is applied mapping each slice into its correct location within the final volume dataspace.

We follow the histogram matching phase with the application of an adaptive smoothing procedure to improve local homogeneity in the data. The aim is to remove image noise while preserving the semantically important boundary information. An anisotropic gaussian filtering is applied, modified by the magnitude and direction of the local greyscale gradient vector. This can be formulated for each voxel as the weighted mean of that voxel and its neighbours. The adaptive weighting selection function (Eq. A.1) governs how much each neighbour contributes to the weighted mean as the current voxel is smoothed. The weight at each neighbour is calculated as a function of the dot product ($a \cdot b = a_x b_x + a_y b_y + a_z b_z$) of the offset vector to that neighbour ($\vec{p}\vec{q}$) and the local greyscale gradient vector (∇u) calculated at the current location. The parameter λ controls the overall strength of the smoothing. This is the initial form of our boundary preserving smoothing approach (c.f. Eq. 3.10), as noted at the end of Section 3.1 in Chapter 3.

$$wt(\vec{p}\vec{q}, \nabla u) = e^{-\left(\frac{\vec{p}\vec{q} \cdot \nabla u}{\lambda}\right)^2} \quad (\text{A.1})$$

This leads to an updated value for the current voxel as being the normalised mean of the weighted neighbours as in Eq. A.2, where i indexes each neighbour (including the current point) from 1 to n and v_i and wt_i are the voxel value and weight respectively, at each location.

$$v = \frac{\sum_{i=1}^n (v_i wt_i)}{\sum_{i=1}^n wt_i} \quad (\text{A.2})$$

This formulation results in an anisotropic smoothing of the data where the local gradient vector controls both the strength of the smoothing applied at each voxel and the degree to which that smoothing is constrained to operate perpendicular to the gradient vector (along the local isosurface or boundary). In this way we achieve the desired edge preserving effect in conjunction with good regional smoothing, in preparation for the effective operation of the segmentation phase to follow.




Fig. A.5: Data smoothing improves regional homogeneity in the data, resulting in superior segmentation results in the fat classification stage of the process.

B. Whole Body Height, Volume, & Mass Estimation

Once the matched dataset has been recomposed the first step in processing is to segment the entire body volume from the background. This is achieved by thresholding the data at a level chosen based on the location of the histogram local minimum between the background and lean tissue peaks, and then applying a 2-D hole filling procedure along the axial direction in order to arrive at a solid whole body object in the data. This gives us a starting point for the fat segmentation procedure to follow and since we know the voxel dimensions also allows us to estimate the subjects height and volume. Once we have completed the fat segmentation step that follows we can additionally estimate the subjects weight, based on the standard known densities for fat and lean tissue.

C. Body Fat Segmentation

The most important part of the process is the four phase body fat segmentation procedure. We can immediately see that the fat tissues in which we are interested are located primarily in the smaller third peak and high intensity tail of the data histogram. However, closer examination of the data reveals a considerable degree of variation in the greylevels that correspond to fat. Even after the histogram matching process the greyscale values for fat still overlap those associated with other nominally lower intensity tissue types such as those representing liver and brain. Therefore accurate segmentation can not be achieved by applying simple methods based on thresholding alone. In order to cope with these issues we have devised a four step segmentation algorithm.

- 1 An initial threshold level is calculated based on an analysis of the data histogram. The peak representing soft tissue is located and voxels whose values fall above the end of this peak are marked as potential fat voxels in the initial fat estimate.
- 2 We then apply a boundary enhancement procedure to compensate for signal drop-off in some peripheral regions of the data. This is based on a chamfer distance measure and applies positive weighting values in areas around the boundaries of the region of interest. This process helps to alleviate the undesirable consequences of partial volume effects at the fat boundaries and in areas where the fat layer is very thin.

- 3 Next we apply a 3-D region growing operator based on hysteresis thresholding, which is applied using a form of conditional dilation. Background voxels in the vicinity of fat voxels are conditionally added to the defined body fat region based on a banded comparison of their greylevels to the expected greylevels for fat voxels, derived from a statistical analysis of the fat estimate so far.
- 4 Finally we apply a region refining process whereby the candidate voxels are grouped into connected regions, which are then processed based on their regional mean greyvalues. Regions with a higher mean value, indicating a stronger fat signal are retained, while those with a lower value are subdivided into accepted and rejected subregions.

Through this process we arrive at a robust segmentation of the signal due to fat tissue within the data volume. See Fig. A.6, which shows a graphical display produced from a segmented dataset. The defined fat region as shown in this coronal slice includes components from all of the four major fat subdivisions, and demonstrates an accurate representation of the fat deposits that have been imaged in this subject.




Fig. A.6: Fat segmentation

For comparison Fig. A.7 illustrates the results achieved using a simplistic thresholding approach. Even with the optimal threshold level manually selected in this case much of the lighter tissue in the liver region has been included in the segmentation in A.7a while our approach, shown in A.7b, demonstrates accurate differentiation between fat and nonfat tissues. It is especially noticeable in the subcutaneous fat low down on the subjects left, at the far side of the abdomen from the liver, that the voxel intensities rep-

representing fat are drifting down towards the intensities seen in the liver. This slow variation throughout the dataset is one of the undesirable characteristics of the data that our approach specifically aims to overcome.




Fig. A.7: A straightforward thresholding approach fails to correctly classify some brighter regions of non-fat tissue as with the liver in (a). Our multistage adaptive classifier (b) correctly handles the liver tissue, eliminating it from the fat classification.

III. RESULTS & DISCUSSION

The analysis procedure that we have described was applied to a database of 42 whole body datasets. Initially numerical results were compiled and the body fat percentages were analysed. The segmentations achieved were reviewed visually by assessing the axial, coronal and sagittal slices in order to confirm that the voxels specified as being body fat matched the true fat regions in the data. 3-D renderings were also used to provide the best possible assessment of the data.

A. Numerical Results

Once a dataset has been processed a wide range of results are available to the user in both numerical and graphical formats. The initial system output display is as in Fig. A.8 where a range of numerical results are presented to the user. The two methods indicated in the results table correspond to the simple thresholding and multi-phase segmentation approaches discussed. As such the method one results are included in this research tool for reference and comparison only. The method two results represent the final system analysis of the input data. In this example the two methods show a fair degree of agreement

in their calculated results. Across the complete range of datasets analysed the simplistic approach often yields wildly inaccurate results, including much soft tissue in its fat estimate. Our method on the other hand performs consistently well on all the analysed datasets.

Series #1: 250 x 32 x 937, 2.022059mm x 8.0mm x 2.022059mm		
Name:		
Weight:	90.0kg	
Sex:	M	
Date of birth:	20/10/1972	
Estimated full body volume = 98,422.84cc		
Estimated full body height = 1.89m		
Calculated Body Fat Table:-		
	Method 1	Method 2
Body weight DICOM	90kg	90kg
Body weight calculated	103.14kg	102.91kg
Estimated BMI	28.73	28.67
Fat by volume	31,803.91cc	33,016.37cc
Fat by weight	29.2kg	30.31kg
% by volume	32.31%	33.55%
% by DICOM weight	32.44%	33.68%
% by calculated weight	28.31%	29.45%

Fig. A.8: System display for an overweight 32 year old male subject illustrating results calculated by the simple and adaptive fat segmentation procedures.

In the results table (Fig. A.8) estimates are made of the subjects height and weight, and measurements are performed to calculate values for the actual and percentage body fat detected, measured by volume and by weight. All of these calculations are performed automatically without any initialisation or subsequent intervention being required on the part of the user. The entire analysis process from raw input data to final results takes less than two minutes, executing on a standard desktop computer.

B. Graphical Results

In addition to these immediately available numerical results a number of visualisation tools can provide for 2-D and 3-D graphical interrogation of the segmented fat volume. Simple orthogonal review allows axial, coronal, and sagittal sections to be examined as illustrated in Fig. A.9. The images can be marked up in various ways to enhance the utility of the tool by highlighting the regions classified as fat tissue.

Male Subjects	Age	Measured BMI	Calculated BMI	Calculated % Body Fat	Female Subjects	Age	Measured BMI	Calculated BMI	Calculated % Body Fat
1	31	32.6	28.67	29.45	1	27	19.3	18.11	21.06
2	33	24.1	23.54	17.27	2	29	24.8	23.53	30.54
3	32	25.4	23.73	27.79	3	20	20.3	19.72	19.61
4	51	27.3	24.43	16.85	4	53	26.8	24.07	34.36
5	29	25.0	22.76	17.36	5	56	22.5	20.77	28.66
6	29	24.8	22.59	6.97	6	27	20.8	20.23	25.21
7	29	25.0	23.68	13.06	7	21	24.6	21.73	30.29
8	28	21.6	19.58	4.77	8	43	21.9	20.59	29.02
9	49	31.1	28.42	30.07	9	41	25.6	26.69	37.46
10	52	34.7	32.58	25.98	10	19	18.0	17.68	15.89
11	24	28.7	23.18	6.24	11	26	21.8	22.21	16.39
12	28	28.7	26.91	31.79	12	18	25.3	26.83	28.89
13	23	26.3	22.62	10.32	13	26	30.5	30.96	38.05
14	19	21.5	19.63	11.63	14	33	22.8	24.11	17.99
15	28	26.3	25.58	20.04	15	23	20.5	22.10	23.71
16	25	18.9	22.30	11.57	16	25	20.3	20.89	19.97
17	26	26.9	24.80	12.21	17	25	29.4	30.74	41.67
18	24	20.9	21.12	6.24	18	27	25.7	25.86	31.78
19	26	27.7	27.58	21.19	19	55	25.1	28.52	35.20
20	30	26.0	25.20	14.96	20	53	20.8	24.27	29.66
21	45	29.9	32.70	28.3	21	42	23.7	23.65	25.16

Table A.1: Body fat results from 42 whole body MRI datasets

3-D volume rendered views are provided by way of a versatile data rendering system, Fig. A.10, which provides an excellent overview of the data. When this tool is used in conjunction with data space clipping as demonstrated in Fig. A.11 it can be applied to effectively visualise the body fat distribution within a given volume of interest, providing a more complete and detailed view of the distribution of fatty tissue within the body.

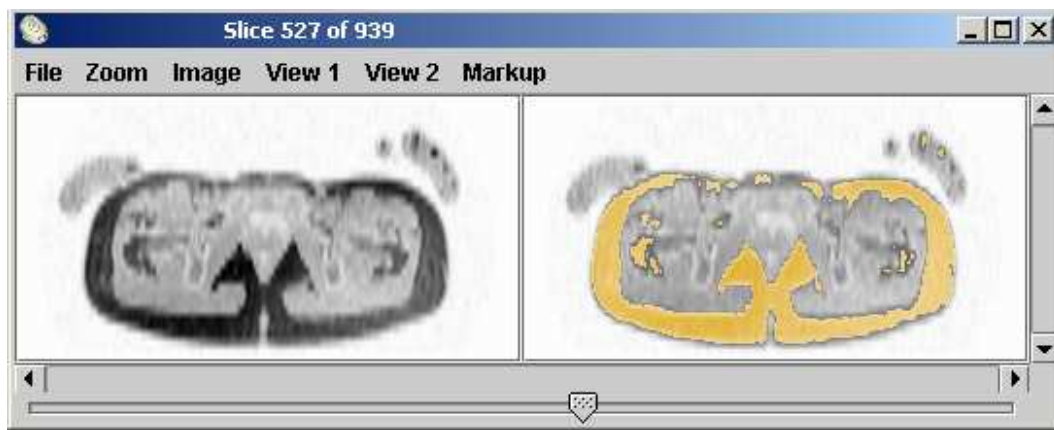


Fig. A.9: Orthogonal section viewer showing an axial slice. The raw data is displayed in the left hand pane while the identified fat regions are highlighted in yellow in the right hand pane. Options for what is displayed in each pane and how it is marked on the images are available through the menus.

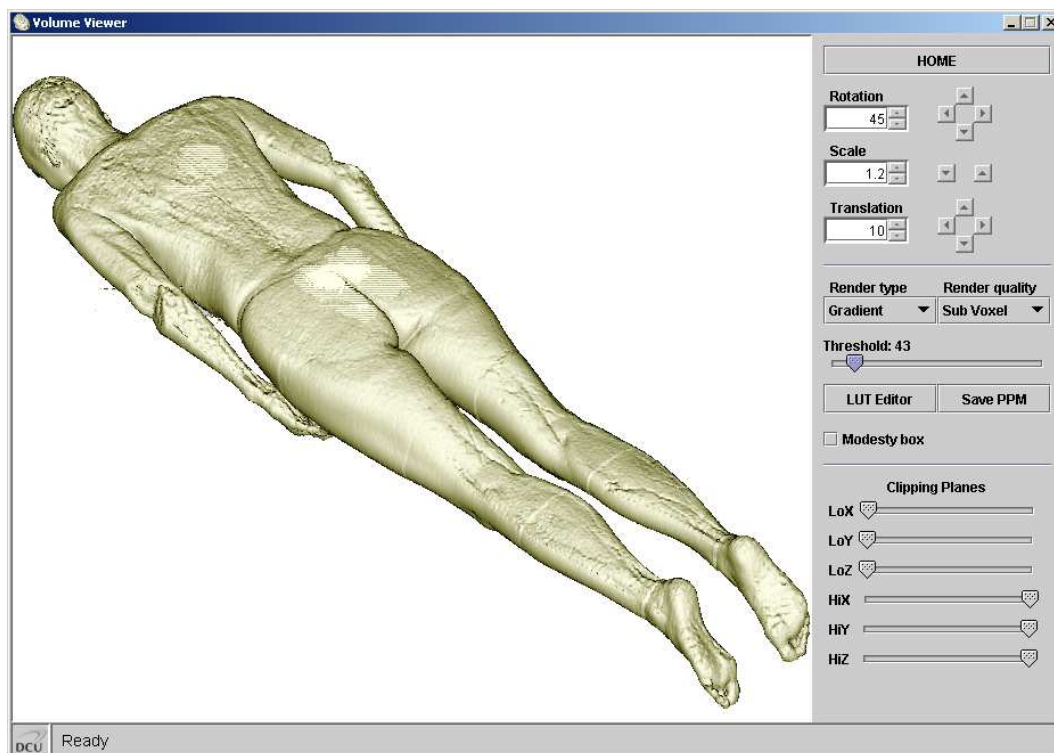


Fig. A.10: A rendering tool allowing flexible manipulation and visualisation of the dataset using a number of volume rendering techniques.




Fig. A.11: Medial cutaway view depicting the left side of the body, from abdomen to thigh, in an overweight 26 year old female subject demonstrating the four major fat subclassifications.

Using these tools it is possible to focus attention onto any chosen subregion of the body and having thus defined a region of interest, detailed analysis can once again be performed on just this targeted zone, generating an analysis of the body composition in terms of fat mass and distribution.

C. Body Fat & BMI

One of the most commonly used body fat indicators is an indirect metric called the body mass index (BMI), which is calculated as a person's weight in kilograms divided by the square of their height in metres. Table A.2 shows the standard BMI bands and the weight level categories that they correspond to.

The BMI is a crude measure of a person's level of body fat, having a number of well recognised failings. It is, none the less, widely used as a rough indicator as to a subject's state of health in terms of their body weight. As an indicator of health and fitness it is particularly ill suited in respect of two

BMI	Category
Less than 18.5	Underweight
18.5 to 24.9	Normal weight
25 to 29.9	Overweight
Greater than 30	Obese

Table A.2: Standardised Body Mass Index (BMI) categories used in the routine assessment of individual body mass: $BMI = \frac{weight}{height\ squared} kg/m^2$

groups, athletes and the elderly. In the case of highly trained athletes their extremely low levels of body fat are counteracted by an unusually high muscle mass that results in a misleadingly elevated BMI reading. On the other hand, in elderly subjects especially, moderate or high fat levels can be disguised by the effects of muscle wastage bringing the total body mass down and leading to an artificially depressed BMI.




Fig. A.12: Graph of BMI against percentage body fat illustrating the distinct gender difference and the difficulty in extrapolating reliably from BMI to actual body fat percentage. Men in general have a lower percentage body fat than women for the same BMI. Subjects with very low body fat but considerable muscle mass, such as athletes, exhibit elevated BMI values, while muscle wastage in older subjects can have the opposite effect.

In this study we examined the correlation between percentage body fat as measured using our approach and the equivalent BMI levels for each subject. The graph in Fig. A.12 illustrates our findings. By separating the data into male and female subjects we can clearly demonstrate the gender difference, showing that men tend to have a lower percentage body fat than women for the same BMI value. This is the expected result, reflecting the differing body compositions observed in men and women.

While the graph confirms that the general trends of body fat and BMI values match, it indicates the complex nature of the relationship between these two measures and highlights the limited degree to which BMI can be used in order to extrapolate an estimate for percentage body fat. As previously mentioned it is also the case that BMI yields no information as to the distribution of fat within the body and as such its usefulness is limited as compared to the richer source of data offered by a fully volumetric fat assessment approach such as WB-MRI.

The BMI was calculated directly from the patients' mass and height, and indirectly using the segmentation results. The directly measured BMI values of these datasets range from 18 to 35. Comparison of the direct and indirect methods was used to validate the segmentation procedure and a good correlation was found, see Fig. A.13. Much of the data spread in the graph in Fig. A.13 is due to incomplete body acquisitions in the MRI scanner.

Restricting the data to the 21 most complete datasets, where data loss at the extremities, especially at the feet of tall subjects, is kept to a minimum, results in a far superior straight line fit with much reduced deviations from the fitted line in the restricted data subsample. This is shown in the broken line and its associated subset of data points indicated in Fig. A.13. The sample points representing incomplete acquisitions all fall above the broken line and pull the full sample trend line (the solid line) in an upward direction. This reflects the relative effect of the missing data on the BMI and % body fat calculations. The missed data represents a relatively small volume, and hence weight (the feet and ankles), but can result in a disproportionately large variation in the estimated height of the subject. With down pointing toes the change can be as much as 30 centimetres in some cases. These two factors result in little difference in the estimated % body fat but can cause a significant elevation in the estimated BMI value since the subject is estimated to be shorter but not much

lighter than would be the case were the measurements based on a complete dataset. Note that in the case of complete datasets the automatic method consistently estimates a BMI less than the measured BMI. This is due to an over estimation in the height caused by down pointing toes. Correspondingly, calculated BMI's based on incomplete datasets tend to be higher than the equivalent measured values, for the reasons given. These observations support the assertion that the automatic measurements achieved represent accurate results for the data under examination.




Fig. A.13: Comparing manually measured BMI and the automatic CAD calculation. This illustrates that the CAD system does a good job in matching the manual BMI measurements. This test was used as an initial validation of the segmentation procedure and demonstrates that the CAD system developed produces reliable measurements. See the text for a full discussion.

D. Body Fat Subdivisions

This usefulness can be still further enhanced within the context of a volumetric body fat analysis tool when the individual fat subclassifications of subcutaneous, visceral, intermuscular, and bone marrow are made available to the user. The subdivisions as illustrated in Fig. A.11 were achieved using a manual editing procedure. Another example of this can be seen in Fig. A.14 where the axial slices corresponding to the start and end of the rendered subsection

of the thigh are also shown. In these you can see where the various fat regions intersect the axial planes. Recall that the data has been interpolated in the y direction as the original dataset contains only 32 coronal slices.




Fig. A.14: Thigh section rendering and associated first and last raw axial views depicting subcutaneous, intermuscular, and bone marrow fat deposits in a normal weight 34 year old female subject.

Fig. A.15 illustrates a number of coronal and sagittal slices through the subvolume rendered in Fig. A.11. Again interpolation is needed in order to generate the sagittal slices, and the fat subclassifications have been marked manually in all cases. These views illustrate the detail in which it is possible to examine the fat distribution in a subject, and once the subclassifications have been defined, useful calculations can be performed in order to quantify metrics such as the visceral to total fat ratio.

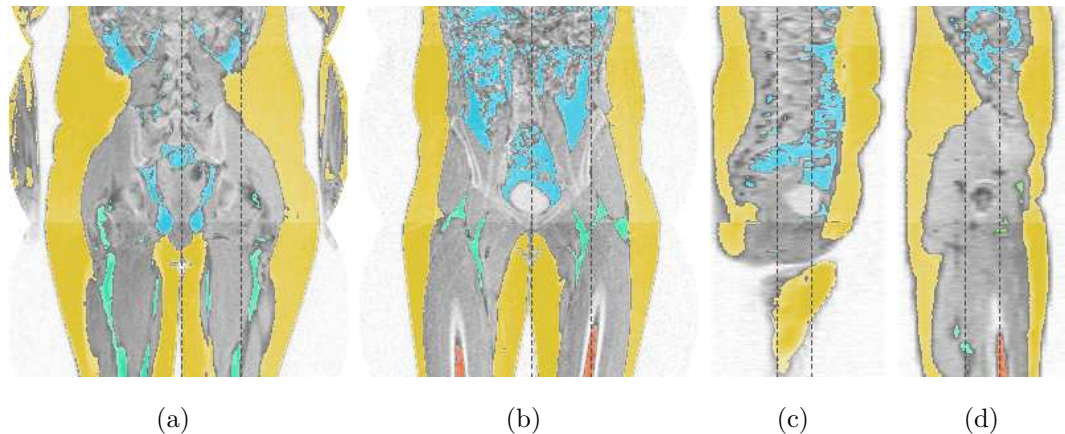


Fig. A.15: Two coronal and two sagittal cross sections through the data volume rendered in Fig. A.11. Dashed lines indicate the points of intersection between the four planes. (a) Posterior coronal section. (b) Anterior coronal section. (c) Medial sagittal section. (d) Lateral sagittal section. The segmented fat regions have again been classified and colour coded as subcutaneous, visceral, intermuscular, and bone marrow deposits.

Our current focus in this project is towards the development of a semiautomated fat subdivision procedure using a model for the expected fat distribution within the body. First we aim to automatically partition the dataset into regions consisting of head, torso, arms, and legs. This will be followed by clustering and labelling phases in order to arrive at a final classification of the fat regions.

IV. CONCLUSION

This appendix has focused on the issues surrounding the preparation, segmentation, and analysis of whole body MRI data for the detailed assessment of body fat level and location. It has described an effective set of preprocessing and segmentation operations designed to overcome the difficulties encountered in the accurate and repeatable isolation of fat tissues in MR data of limited resolution and quality. Results achieved on a database of 42 subjects were presented and discussed, and the future direction for this investigation was indicated. This illustrates the wider applicability and the effectiveness of the preprocessing techniques introduced in the main body of this thesis and thus serves to further emphasise the value and the potential of these techniques in the context of medical image processing and CAD applications.

Appendix B

NeatMRI

We developed the NeatMRI library and environment in order to provide a flexible and easy to use framework in which to conduct the development and testing of our image processing algorithms and procedures. In this appendix we introduce the functionality of the library and illustrate some of the applications that we have developed using the tools that it provides. We illustrate the two routes to interfacing with the library, through Java and C, and provide a quick reference for the full collection of procedures that the library provides.

B.1 Library

B.1.1 Introduction

The library comprises a large and growing collection of tools and procedures for the processing, segmentation, analysis, and visualisation of two and 3-D image data, along with all the low level utility functions necessary in order to manipulate the data in a straightforward and powerful fashion. It exports all these procedures, and provides interfaces to its functionality through both C and Java.

Using the library from C

Include VSGLib.h in the C project and access methods as illustrated below:

```
#include "VSGLib.h"

void main(void) {
    int vol;

    vol = VSGgetVolume(256,256,64);           // Create volume
    //...
    VSGreleaseVolume(vol);                   // Release volume
}
```

Using the library from Java

Include VSGLib.java in the java project and access methods as illustrated below:

```
public class Test {
    public static void main(String[] args)
    {
        int vol = VSGLib.VSGgetVolume(256,256,64); // Create volume
        //...
        VSGLib.VSGreleaseVolume(vol);               // Release volume
    }
}
```

The library interfaces through C and Java are identical except for the requisite differences in the declared types of some formal parameters, as shown in Table B.1. Code fragments in C and Java follow, to illustrate the differences in usage.

C	Java
char *	String
int *	int []
double *	double []

Table B.1: Parameter differences between C and Java.

This fragment of C code loads a PGM image into memory.

```
char *filename = "C:/images/image1.pgm";
int width;
int height;
int *image;

VSGgetPNMSize(filename,&width,&height);
image = (int *)malloc(width*height*sizeof(int));
VSGreadPGM(filename,image);
```

The same code fragment is implemented in Java as follows.

```
String filename = "C:/images/image1.pgm";
int[] width = new int[1];
int[] height = new int[1];
int[] image;

VSGLib.VSGgetPNMSize(filename,width,height);
image = new int[width[0]*height[0]];
VSGLib.VSGreadPGM(filename,image);
```

B.1.2 Quick Reference

This section lists all 148 library functions in the current version, arranged into functional groupings. A detailed description of the operation and use of each routine is provided in the comprehensive ‘VSG Volume Processing Library — Reference Manual’, which has been written to accompany the library.

General routines

VSGabout

Volume creation and deletion routines

VSGgetVolume

VSGcopyVolume

VSGloadVolume

VSGloadSlice

VSGresizeVolume

VSGreorientVolume

VSGcombineVolumes

VSGreleaseVolume

Image reading and writing routines

VSGsavePGMs

VSGsavePPMs

VSGgetPNMSize

VSGreadPGM

VSGreadPPM

VSGwritePGM

VSGwritePPM

AVI movie sequence manipulation routines

VSGjoinAVIs
VSGcreateAVI
VSGopenAVI
VSGcloseAVI
VSGsetAVIFrameData
VSGgetAVIFrameData
VSGsetAVIFrameRate
VSGgetAVIFrameRate
VSGgetAVIFrameWidth
VSGgetAVIFrameHeight
VSGgetAVIFrameCount

Volume manipulation routines

VSGsaveVolume
VSGgetOrthogonalSlice
VSGgetObliqueSlice
VSGgetWidth
VSGgetHeight
VSGgetDepth
VSGgetPixelsPerImage
VSGgetVoxelsPerVolume
VSGgetDataMin
VSGgetDataMax
VSGgetDataSum
VSGgetDataSumSquared
VSGgetVoxel
VSGsetVoxel
VSGgetSlice
VSGsetSlice
VSGgetData
VSGsetData
VSGgetVoxelSize
VSGsetVoxelSize

VSGgetRowOrientation
VSGsetRowOrientation
VSGgetColumnOrientation
VSGsetColumnOrientation
VSGgetGlobalOrigin
VSGsetGlobalOrigin
VSGgetHistogram
VSGsetGreyData
VSGset2DBorder
VSGset3DBorder
VSGinvert
VSGshiftGreyRange
VSGthreshold
VSGadaptiveThreshold
VSGdoubleThreshold

Volume processing routines

VSGcopyData
VSGcopyNonZeroData
VSGcopySubVolume
VSGsubtract
VSGminimum
VSGmaximum

Volume filtering routines

VSGgradient
VSGgrassfireFilter
VSGmeanFilter
VSGmedianFilter
VSGlaplacianFilter
VSGmorphologyErode
VSGmorphologyDilate
VSGmorphologyConditionalDilate
VSGmorphologyOpen

VSGmorphologyClose
VSGmorphologyGradient
VSGmorphologyInternalGradient
VSGmorphologyExternalGradient
VSGmorphologyReconstructByDilation
VSGadaptiveFilter

Volume segmentation routines

VSGmarkerGuidedWatershed
VSGactiveRegionClassifier
VSGclusteringClassifier

Object map routines

VSGallocateObjectMaps
VSGreleaseObjectMaps
VSGgetObjectMapCount
VSGvolumeToObjectMap
VSGobjectMapToVolume
VSGobjectMapToObjectMap
VSGgetObjectMapMin
VSGgetObjectMapMax
VSGgetObjectMapVoxel
VSGsetObjectMapVoxel
VSGsetObjectMapData
VSGlabelExists
VSGchangeLabel
VSGfloodFill
VSGlabel
VSGsuppressObject
VSGgetObjectBoundingBox

Structuring element routines

VSGgetStructuringElement
VSGbuildStructuringElement
VSGbuildIsotropicStructuringElement
VSGmakeStructuringElement
VSGreleaseStructuringElement
VSGgetStructuringElementSize

Colour lookup table (LUT) routines

VSGgetVolumeLUTPointer
VSGgetRenderLUTPointer
VSGresizeColourLUT
VSGinvertLUT
VSGsetLUTValue
VSGsetLUTRamp
VSGsetLUTChannelRamp
VSGsetLUTRainbow
VSGsetLUTMixedRainbow
VSGgetNumberOfLUTs
VSGgetSizeOfLUTs
VSGgetLUTValue
VSGgetColour

Surface routines

VSGloadSurface
VSGextractSurface
VSGextractReliefSurface
VSGreleaseSurface
VSGsaveSurface

Volume rendering routines

VSGgetVolumeRenderParam
VSGreleaseVolumeRenderParam
VSGsetDefaultVolumeRenderTransform
VSGsetVolumeRenderType
VSGsetVolumeRenderQuality
VSGsetVolumeRenderThreshold
VSGsetVolumeRenderClippingPlane
VSGsetVolumeObjectVisibility
VSGrenderVolume

Surface rendering routines

VSGgetSurfaceRenderParam
VSGreleaseSurfaceRenderParam
VSGsetDefaultSurfaceRenderTransform
VSGsetSurfaceRenderType
VSGrenderSurface

General rendering routines

VSGsetRenderRotation
VSGsetRenderScale
VSGsetRenderTranslation
VSGincRenderRotation
VSGincRenderScale
VSGincRenderTranslation
VSGresizeRenderImage

B.2 Algorithms

In this section we provide concrete implementations for a number of the key algorithms and procedures developed and presented in this thesis. A list of the routines documented in this section, along with brief descriptions and an indication of how many lines of code each routine occupies, is given below in Table B.2.

Name	Description	Size
<code>LinearX</code>	Calculate x given y and two other points, using linear interpolation.	4 lines
<code>LinearY</code>	Calculate y given x and two other points, using linear interpolation.	4 lines
<code>QuadraticMaxMinY</code>	Calculate the y position of the extrema of the quadratic passing through three points.	4 lines
<code>CubicY</code>	Calculate y given x and three other points, using cubic interpolation.	4 lines
<code>FindMinMaxPair</code>	Implementation of Algorithm 2.1 on page 39	32 lines
<code>FindFixedPoints</code>	Implementation of Algorithm 2.2 on page 46	9 lines
<code>ReBinHistogram</code>	Implementation of Algorithm 2.3 on page 48	24 lines
<code>ReSizeBins</code>	Implementation of Algorithm 2.4 on page 49	30 lines
<code>IntensityCorrectImage</code>	Given a histogram distribution allocate the correct number of samples to each bin.	63 lines
<code>Gradient</code>	Calculates a local greyscale gradient vector in 2-D non-isometric data.	16 lines
<code>GaussianSmooth</code>	Performs gradient-weighted adaptive Gaussian smoothing in 2-D non-isometric data.	35 lines
<code>DownhillFilter</code>	Algorithm 4.1 (page 96): The downhill filter for efficient reconstruction by dilation.	40 lines
<code>GrassFire</code>	The grassfire distance transform implemented using the directed filter algorithm.	42 lines

Table B.2: List and brief descriptions of the algorithms and routines for which C code implementations are provided in the rest of this section.

B.2.1 Intensity Non-uniformity Correction

The following C routines implement the key components of the algorithms described in Chapter 2 for histogram based inter-slice intensity non-uniformity correction.

```

double LinearX(double y, double x1, double y1, double x2, double y2)
{
    return((y*(x1-x2)-(x1*y2-x2*y1))/(y1-y2));
}

double LinearY(double x, double x1, double y1, double x2, double y2)
{
    return(x*(y1-y2)/(x1-x2)+(x1*y2-x2*y1)/(x1-x2));
}

double QuadraticMaxMinY(double y1, double y2, double y3)
{
    return((y3-y1)*(y3-y1)/(16.0*y2-8.0*(y3+y1))+y2);
}

double CubicY(double x, double y1, double y2, double y3, double y4)
{
    return(y2+((y3-(y1/3.0+y2/2.0+y4/6.0))+(((y1+y3)/2.0-y2)+((y2-y3)/2.0+(y4-y1)/6.0)*x)*x)*x);
}

void FindMinMaxPair(int *histo, int size, int *A, int *C)
{
    int B;
    int i,j;

    *A=0;
    for (i=1; i<size; i++)
        if (histo[*A] < histo[i])
            *A=i;
    B=*C=0;
    for (i=*A-1; i>0; i--) {
        for (j=i-1; j>=0; j--) {
            if (histo[j]-histo[i]>histo[*C]-histo[B]) {
                B=i;
                *C=j;
            }
        }
    }
    for (i=*A+1; i<size-1; i++) {
        for (j=i+1; j<size; j++) {
            if (histo[j]-histo[i]>histo[*C]-histo[B]) {
                B=i;
                *C=j;
            }
        }
    }
    if (*A>*C) {
        B=*C;
        *C=*A;
        *A=B;
    }
}

void FindFixedPoints(int *histo, int A, int C, double *P)
{
    P[1]=QuadraticMaxMinY(histo[A-1],histo[A],histo[A+1])*2.0/3.0;
    while(histo[A--]>P[1]);
    P[1]=LinearX(P[1],A,histo[A],A+1,histo[A+1]);
    P[2]=QuadraticMaxMinY(histo[C-1],histo[C],histo[C+1])*2.0/3.0;
    while(histo[C++]>P[2]);
    P[2]=LinearX(P[2],C-1,histo[C-1],C,histo[C]);
}

```

```

void ReBinHistogram(int *src, double *dst, int size, double *P, double *PMax)
{
    int i=0;
    double fp,ip;

    for (; i<=PMax[1]; i++) {
        fp=modf(LinearY(i,0.0,0.0,PMax[1],P[1]),&ip);
        if (ip<1)
            dst[i]=CubicY(fp-1.0,src[0],src[1],src[2],src[3]);
        else
            dst[i]=CubicY(fp,src[(int)ip-1],src[(int)ip],src[(int)ip+1],src[(int)ip+2]);
    }
    for (; i<=PMax[2]; i++) {
        fp=modf(LinearY(i,PMax[1],P[1],PMax[2],P[2]),&ip);
        dst[i]=CubicY(fp,src[(int)ip-1],src[(int)ip],src[(int)ip+1],src[(int)ip+2]);
    }
    for (; i<=size-1; i++) {
        fp=modf(LinearY(i,PMax[2],P[2],size-1,size-1),&ip);
        if (ip<size-2)
            dst[i]=CubicY(fp,src[(int)ip-1],src[(int)ip],src[(int)ip+1],src[(int)ip+2]);
        else
            dst[i]=CubicY(ip+fp-(size-3),src[size-4],src[size-3],src[size-2],src[size-1]);
    }
}

int ReSizeBins(double *src, int *dst, int size, int numVoxels)
{
    int i;
    int remainder;
    int sum=0;
    double R=0.0;

    for (i=0; i<size; i++)
        R+=src[i];
    for (i=0; i<size; i++) {
        dst[i]=(int)(src[i]*numVoxels/R+0.5);
        sum+=dst[i];
    }
    size=CompressHistogram(dst,size);
    remainder=numVoxels-sum;
    if (remainder<0) {
        remainder=-remainder;
        size--;
        while (remainder>=dst[size])
            remainder-=dst[size--];
        dst[size++]-=remainder;
    } else if (remainder>0) {
        while (remainder>dst[size-1]) {
            dst[size]=dst[size-1];
            remainder-=dst[size++];
        }
        dst[size++]+=remainder;
    }
    return(size);
}

```

```

double WeightFunction(int pos, int *img, int w, int h)
{
    int i,x,y,dx,dy;
    double top=0.0;
    double bot=0.0;

    x=pos/w;
    y=pos/w;
    for (dy=-2,i=0; dy<=2; dy++) {
        for (dx=-2; dx<=2; dx++,i++) {
            if ((dx!=0 || dy!=0) && x+dx>=0 && y+dy>=0 && x+dx<w && y+dy<h) {
                top+=img[x+dx+(y+dy)*w]*delta[i];
                bot+=delta[i];
            }
        }
    }
    return(top/bot);
}

typedef struct {
    int pos;
    int val;
    double wt;
} tInfo;

int LessThan(const tInfo *a, const tInfo *b)
{
    if (a->val < b->val)
        return(-1);
    else if (a->val > b->val)
        return(1);
    else if (a->wt < b->wt)
        return(-1);
    else if (a->wt > b->wt)
        return(1);
    else
        return(0);
}

void IntensityCorrectImage(int *histo, int *img, int w, int h)
{
    int i;
    int val=0;
    int cnt=0;
    tInfo *ptr;

    ptr=(tInfo *)malloc(w*h*sizeof(tInfo));
    for (i=0; i<w*h; i++) {
        ptr[i].pos=i;
        ptr[i].val=img[i];
        ptr[i].wt=WeightFunction(i,img,w,h);
    }
    qsort(ptr,w*h,sizeof(tInfo),LessThan);
    for (i=0; i<w*h; i++) {
        while (cnt>=histo[val]) {
            cnt=0;
            val++;
        }
        cnt++;
        img[ptr[i].pos]=val;
    }
    free(ptr);
}

```


B.2.2 Gradient-Weighted Smoothing Algorithm

A concise implementation of the gradient-weighted adaptive Gaussian smoothing algorithm in 2-D non-isometric data. Greater efficiency can be achieved by the use of lookup tables, thus avoiding repeated calculations in a number of locations. The implementation given here sacrifices some efficiency in favour of simplicity and clarity in presentation. The implementation directly mirrors the form given in Algorithm 3.1 on page 60 of this thesis.

```
double Gradient(int *data,int w,int h,int px,int py,double dx,double dy,double *gx,double *gy)
{
    double wtx1,wtx2,wty1,wty2;
    double *p1,*p2,*p3;

    wtx1=dx/(3*dx*dx+dy*dy);
    wtx2=(dx*dx+dy*dy)/(3*dx*dx+dx*dy*dy);
    wty1=dy/(3*dy*dy+dx*dx);
    wty2=(dy*dy+dx*dx)/(3*dy*dy+dy*dx*dx);
    p1=data+w*(py-1)+px-1;
    p2=data+w*(py )+px-1;
    p3=data+w*(py+1)+px-1;
    *gx=(*(p1+2)*wtx1 + *(p2+2)*wtx2 *(p3+2)*wtx1)-(*p1*wtx1 + *p2*wtx2 *p3*wtx1);
    *gy=( *p3 *wty1 + *(p3+1)*wty2 *(p3+2)*wty1)-(*p1*wty1 + *(p1+1)*wty2 *(p1+2)*wty1);
    return(sqrt((*gx)*(*gx)+(*gy)*(*gy)));
}

void GaussianSmooth(int *src,int *dst,int w,int h,double dx,double dy,double lambda,double mu)
{
    int px,py; // coordinates of current pixel
    int qx,qy; // offset to neighbour
    double gx,gy,gm; // gradient x, y, and magnitude
    double wtq,wtsum,val;
    double wt1,wt2;
    double v1,v2;

    wt1=1.0/(lambda*lambda*exp(mu)*exp(mu));
    wt2=(exp(2.0*mu)-exp(-2.0*mu))/(lambda*lambda);
    for (py=1; py<h-1; py++)
    {
        for (px=1; px<w-1; px++)
        {
            gm=Gradient(src,w,h,px,py,dx,dy,&gx,&gy);
            wtsum=0.0;
            val=0.0;
            for (qy=-1; qy<=1; qy++)
            {
                for (qx=-1; qx<=1; qx++)
                {
                    v1=sqrt(dx*dx*qx*qx+dy*dy*qy*qy)*gm;
                    v1=v1*v1;
                    v2=dx*qx*gx+dy*qy*gy;
                    v2=v2*v2;
                    wtq=exp(-(v1*wt1 + v2*wt2));
                    wtsum+=wtq;
                    val+=wtq*(*(src+(py+qy)*w+px+qx));
                }
            }
            *(dst+py*w+px)=(int)(val/wtsum+0.5);
        }
    }
}
```

B.2.3 Downhill Filter: Strengthened Precondition

This is an extremely compact and efficient implementation of the downhill filter. It requires that the form of the mask image input be restricted as explained in the text of Chapter 4. Efficiency can be improved still further, in the order of 30 to 40 percent, by unwrapping the two innermost for loops and optimising the resulting code, but even in the compact form presented here it is in the general case faster than the best of the alternatives tested.

```

void DownhillFilter(const int *mask, int *marker, int width, int height)
{
    int ix,iy,ox,oy,offset,maxVal=0;
    int currentQ,currentPixel,pixelsPerImage=width*height;
    int *istart,*iarray;

    for (offset = pixelsPerImage-1; offset >= 0; offset--)
        if (marker[offset] > maxVal)
            maxVal = marker[offset];
    istart = (int*)malloc((maxVal+1)*sizeof(int));
    iarray = (int*)malloc(pixelsPerImage*sizeof(int));
    memset(istart,0xfe,(maxVal+1)*sizeof(int));
    memset(iarray,0xff,pixelsPerImage*sizeof(int));
    for (offset = pixelsPerImage-1; offset >= 0; offset--) {
        if (marker[offset] != 0) {
            iarray[offset] = istart[mask[offset]];
            istart[mask[offset]] = offset;
        }
    }
    for (currentQ = maxVal; currentQ > 0; currentQ--) {
        currentPixel = istart[currentQ];
        while (currentPixel != 0xfefefefe) {
            istart[currentQ] = iarray[currentPixel];
            ix = currentPixel%width;
            iy = currentPixel/width;
            for (oy = iy-1; oy <= iy+1; oy++) {
                for (ox = ix-1; ox <= ix+1; ox++) {
                    if (ox>=0 && oy>=0 && ox<width && oy<height && iarray[offset=ox+oy*width]==0xffffffff) {
                        marker[offset] = mask[offset]>currentQ?currentQ:mask[offset];
                        iarray[offset] = istart[marker[offset]];
                        istart[marker[offset]] = offset;
                    }
                }
            }
            currentPixel = istart[currentQ];
        }
    }
    free(istart);
    free(iarray);
}

```

B.2.4 Grassfire Distance Transform

This is a compact and efficient implementation of the grassfire distance transform, based on the directed filtering algorithmic pattern described in detail in Chapter 4. It provides a fast, single pass implementation of the transform, which outputs a distance map indicating the shortest chessboard distance from each pixel in the image to the nearest background pixel.

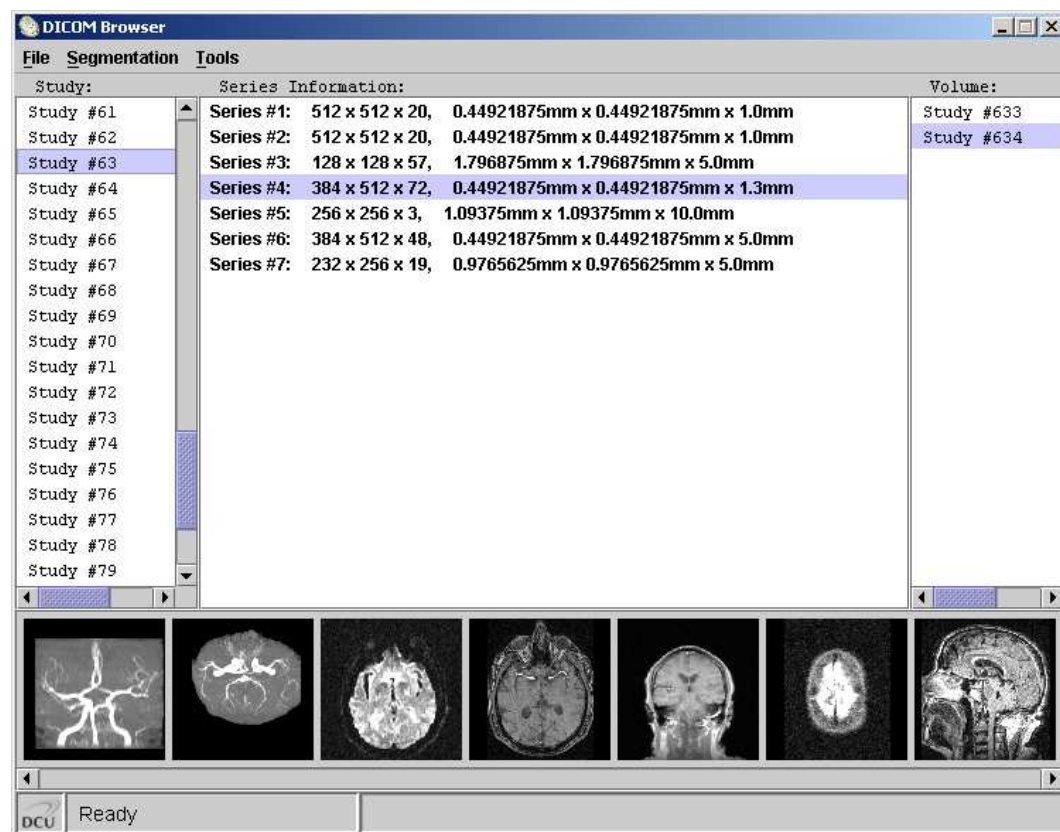
```
// input  : src - figure non-zero, ground zero
// output : dst - chess board distance to nearest ground pixel
void GrassFire(const int *src, int *dst, int width, int height)
{
    int ix,iy,ox,oy,offset;
    int head=-999;
    int tail=-1;
    int pixelsPerImage=width*height;
    int *iarray,*p;

    p = (int*)malloc((pixelsPerImage+1)*sizeof(int));
    iarray = p+1;
    for (offset = pixelsPerImage-1; offset >= 0; offset--) {
        if (src[offset] != 0) {
            dst[offset] = -1;
        } else {
            dst[offset] = 0;
            if (head == -999) {
                head = iarray[-1] = offset;
            } else {
                iarray[head] = offset;
                head = offset;
            }
        }
    }
    if (head != -999) {
        while (head != tail) {
            tail = iarray[tail];
            ix = tail%width;
            iy = tail/width;
            for (oy = iy-1; oy <= iy+1; oy++) {
                for (ox = ix-1; ox <= ix+1; ox++) {
                    if (ox >= 0 && oy >= 0 && ox < width && oy < height &&
                        dst[offset = ox+oy*width] == -1) {
                        dst[offset] = dst[tail]+1;
                        iarray[head] = offset;
                        head = offset;
                    }
                }
            }
        }
    }
    free(p);
}
```

B.3 Applications

We have developed a wide range of data processing, analysis, and visualisation tools within the framework of the NeatMRI library. Here we provide an overview of the kinds of tools that we have constructed in the course of this project, illustrating the versatility and flexibility of the NeatMRI environment.

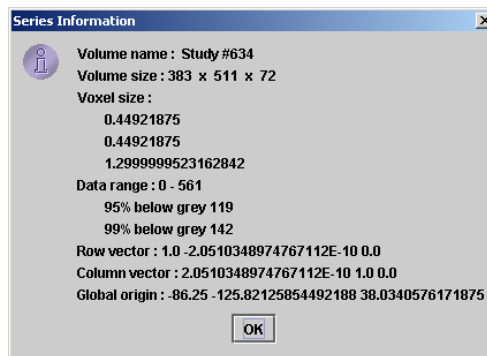

At the top level of all of our applications we provide a database browser. Below is an example of such a browser, displaying a study from our general MRI database. This particular study consists of seven individual series, and represents an examination of the subjects head.




The menus at the top of this browser provide access to all the subsidiary tools integrated into the application, as described below. The list to the left contains all the studies registered within the current database. A study consists of one or more series. The list of series that comprise the current study, highlighted in the study list, are displayed in the centre panel, and a thumbnail image for each is also shown along the bottom of the browser. Double clicking a series or its thumbnail loads that series into memory. Currently loaded series are displayed in the rightmost panel.

The file menu provides the following set of functionality:

- New DICOM studies can be scanned, reconstructed, and added to the current database.
- Multi-section volumes can be reconstructed using the DICOM header information.
- The database can be purged, deleting all studies so as to leave an empty database.
- Volumes can be loaded in four different ways:
 - 1 – Load exactly as saved in size and orientation
 - 2 – Load interpolated to yield cubic voxels ($1 \times 1 \times 1 \text{mm}^3$)
 - 3 – Load interpolated and resliced into an axial orientation
 - 4 – Load a volume not registered in the current database
- The selected volume in the loaded volume list, or all volumes, can be unloaded from memory, freeing all associated space.
- Volume details can be displayed in a dialog box, as below.




- An intensity histogram for the current volume can be calculated and displayed in a popup window.
- The ‘About’ dialog box can be displayed.



- The application can be exited.

The segmentation menu provides the following functions:



- The four steps in body fat analysis can be applied one at a time:
 - 1 – Isolate and characterise the whole body.
 - 2 – Apply the basic segmentation procedure.
 - 3 – Apply the adaptive procedure.
 - 4 – Store the results as an html page.





- The four steps in body fat analysis can be applied in one go.
- The four steps can be applied to every volume in the database.
- Any one of a number of segmentation procedures can be applied to the current volume.

The tools menu provides the following set of functionality:

- The current volume can be browsed in an axial, coronal, or sagittal orientation.
- The current volume can be examined using a flexible multi-view orthogonal sections viewer



- The current volume can be rendered and manipulated in a set of rendering tools, using a number of possible volume and surface rendering algorithms. Below we see a maximum intensity projection (MIP) volume rendering and a Gouraud shaded surface rendering of the same dataset showing the segmented biliary tree. The surface was constructed using a variant of the marching cubes algorithm.



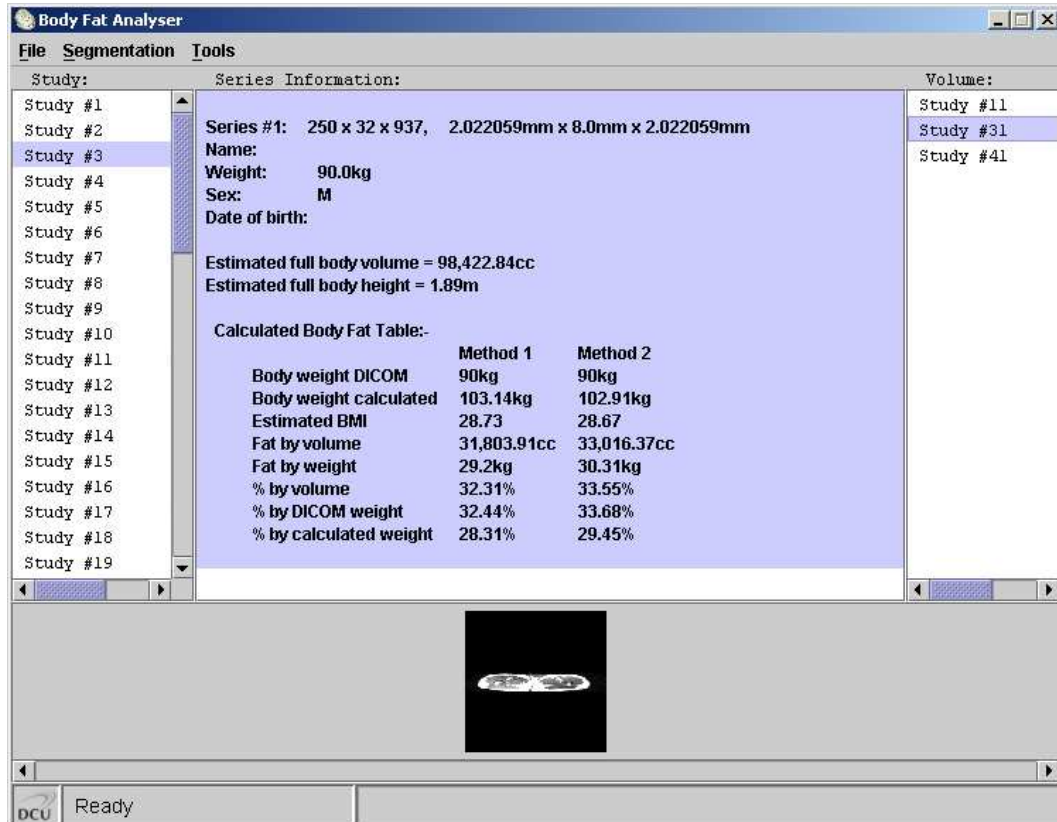
- Multiple volumes from the same study can be merged using the recorded DICOM coordinate system information to achieve registration.
- An adaptive smoothing filter based on gradient-weighted Gaussian smoothing can be applied to the current volume to reduce noise in the data.

B.3.1 Body fat tools

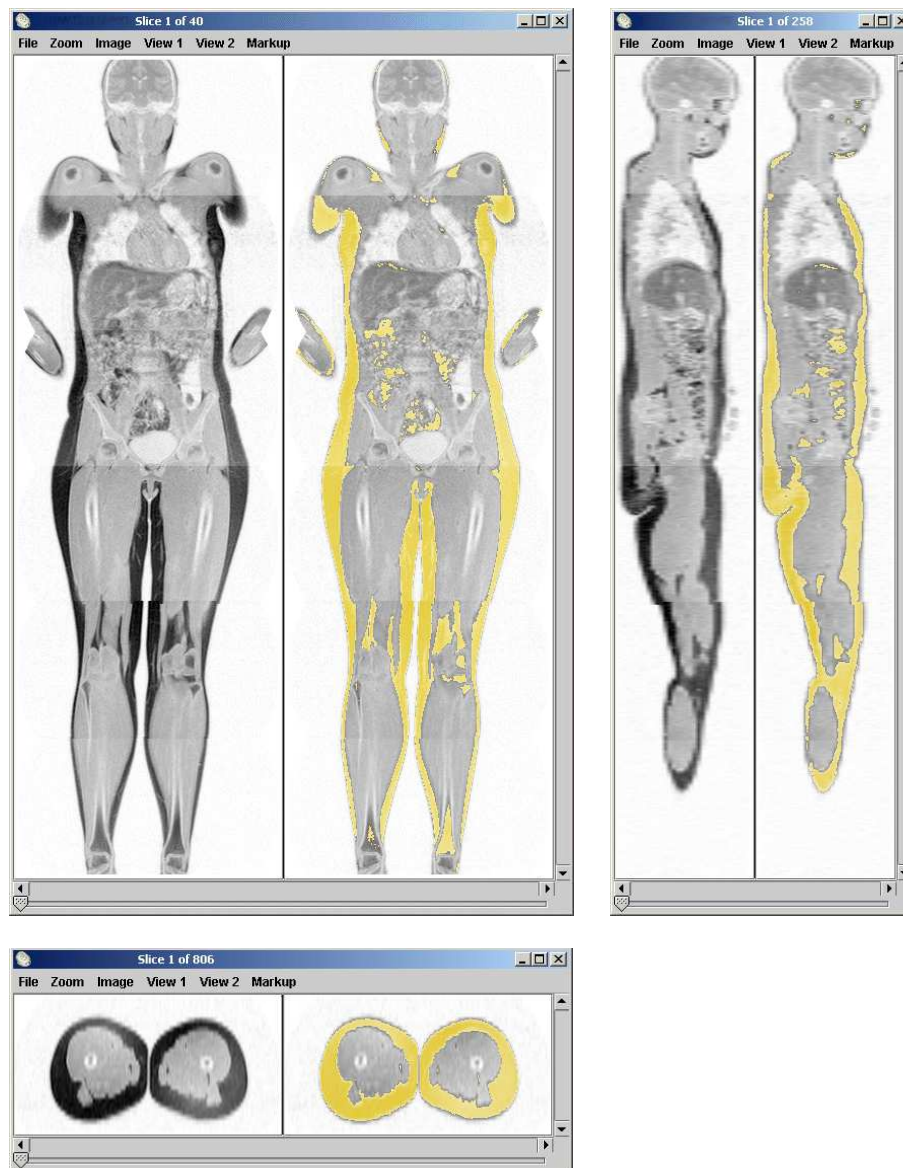
The prospective study conducted into the measurement and localisation of body fat in whole body MRI lead to the development of a number of notable enhancements to the basic set of tools described so far.




Firstly the basic browser interface was extended so as to provide full reporting of segmentation and analysis results. Since in whole body MRI each study consists of only a single series the space in the central panel is used to display each set of results as they are calculated, as illustrated below.



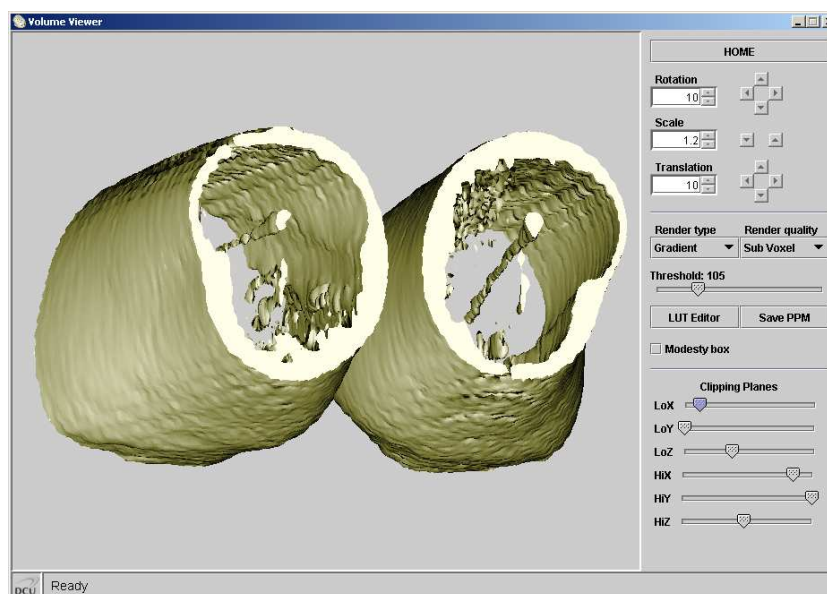
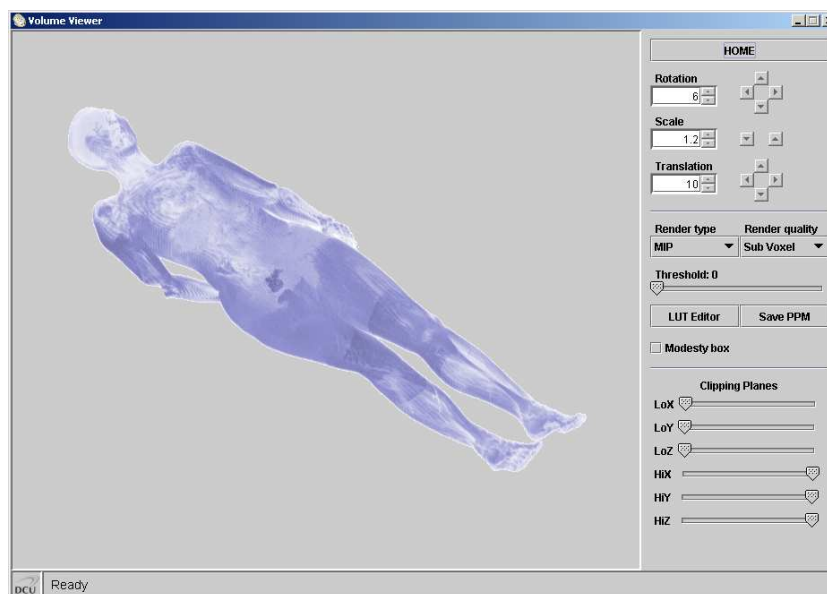
The orthogonal section viewer tools were enhanced to display multiple views of the orthogonal slice data, providing useful options for zoom, highlighting, contrast enhancement, and data markup, as well as image saving.



A versatile LUT editor was also developed in order to allow flexible control over the colour lookup table used in generating the displayed images in these and other visualisation tools such as the surface and volume rendering applications. An extensible library of useful standard LUTs is provided to help maximise the ease of use of the LUT editor.



A more comprehensive volume rendering interface was also implemented allowing for full control over the rendering procedure. Below are shown two views of this new interface. The first is a MIP of a whole body dataset, while the second shows a greyscale gradient rendering of a dataset subregion.



The home button returns the rotation, scale, and translation parameters to predefined values allowing quick return to a default view. Tradeoff between render quality and speed is possible from the 'Render quality' dropdown menu, which provides four settings from course to fine. The six 'Clipping Planes' sliders allow for dataspace clipping to be performed such that region of interest can be specified and examined in detail, eliminating interference from non-relevant regions of the data.

Bibliography

- Acton, S. (1998), ‘Multigrid anisotropic diffusion’, *IEEE Transactions on Image Processing* **7**(3), 280–291.
- Agrafiotis, D., Jones, M., Nikolov, S., Halliwell, M., Bull, D. and Canagarajah, N. (2001), Virtual liver biopsy: Image processing and 3d visualization, *in* S. Cantu, ed., ‘Proc. International Conference on Image Processing’, Thessaloniki, Greece, pp. 331–334.
- Angulo, J. and Serra, J. (2003), ‘Automatic analysis of dna microarray images using mathematical morphology’, *Bioinformatics* **19**(5), 553–562.
- Araujo, A., Guimaraes, S. and Cerqueira, G. (2001), A new approach for old movie restoration, *in* ‘Proc. SPIE-High-speed Imaging and Sequence Analysis’, pp. 67–77.
- Ashburner, J. and Friston, K. (2000), ‘Voxel-based morphometry - the methods’, *NeuroImage* **11**, 805–821.
- Barkhausen, J., Quick, H., Lauenstein, T., Goyen, M., Ruehm, S., Laub, G., Debatin, J. and Ladd, M. (2001), ‘Whole-body mr imaging in 30 seconds with real-time true fisp and a continuously rolling table platform: Feasibility study’, *Radiology* **220**(1), 252–256.
- Beucher, S. and Meyer, F. (1993), The morphological approach to segmentation: The watershed transformation, *in* E. Dougherty, ed., ‘Mathematical Morphology in Image Processing’, Marcel Dekker Inc, pp. 433–481.
- Bigand, A., Bouwmans, T. and Dubus, J. (2001), ‘Extraction of line segments from fuzzy images’, *Pattern Recognition Letters* **22**, 1405–1418.
- Black, M., Sapiro, G., Marimont, D. and Heeger, D. (1998), ‘Robust anisotropic diffusion’, *IEEE Transactions on Image Processing* **7**(3), 421–432.
- Boissonnat, J. (1984), ‘Geometric structures for three-dimensional shape representation’, *ACM Transactions on Graphics* **3**(4), 266–286.
- Boissonnat, J. (1988), ‘Shape reconstruction from planar cross sections’, *Computer Vision, Graphics, and Image Processing* **44**, 1–29.

- Boraschi, P., Braccini, G., Gigoni, R., Geloni, M. and Perri, G. (1999a), 'Mr cholangiopancreatography: Value of axial and coronal fast spin-echo fat-suppressed t2-weighted sequences', *European Journal of Radiology* **32**, 171–181.
- Boraschi, P., Neri, E., Braccini, G., Gigoni, R., Carmella, D., Perri, G. and Bartolozzi, C. (1999b), 'Choledocolithiasis: Diagnostic accuracy of mr cholangiopancreatography. three-year experience', *Magnetic Resonance Imaging* **17**(9), 1245–1253.
- Brennan, D., Whelan, P., Robinson, K., Ghita, O., Sadleir, R., O'Brien, J. and Eustace, S. (2005), 'Rapid automated measurement of body fat distribution from whole body mri', *American Journal of Roentgenology* **185**(2), 418–423.
- Bromiley, P., Thacker, N. and Courtney, P. (2002), 'Non-parametric image subtraction using grey level scattergrams', *Image and Vision Computing* **20**, 609–617.
- Bueno, G., Fisher, M., Burnham, K., Mills, J. and Haas, O. (2001), Automatic segmentation of clinical structures for rtp: Evaluation of a morphological approach, in 'Proc. Medical Image Understanding and Analysis', Birmingham, UK, pp. 73–76.
- Canero, C. and Radeva, P. (2003), 'Vesselness enhancement diffusion', *Pattern Recognition Letters* **24**, 3141–3151.
- Canny, J. (1986), 'A computational approach to edge detection', *IEEE Transactions on Pattern Analysis and Machine Intelligence* **8**, 679–698.
- Cesari, S., Liessi, G., Balestreri, L. and Morassut, S. (2000), 'Raysum reconstruction algorithm in mr cholangiopancreatography', *Magnetic Resonance Imaging* **18**, 217–219.
- Chen, Y. and Wang, M. (2004), 'Mr cholangiography 3d biliary tree automatic reconstruction system', *Computerized Medical Imaging and Graphics* **28**, 13–20.
- Chen, Z. and Molloy, S. (2002), 'Vascular tree object segmentation by deskeletonization of valley courses', *Computerized Medical Imaging and Graphics* **26**, 419–428.
- Crank, J. (1975), *The Mathematics of Diffusion*, Oxford.
- Dalal, P., Howlett, D., Sallomi, D., Marchbank, N., Watson, G., Marr, A., Dunk, A. and Smith, A. (2004), 'Does intravenous glucagon improve common bile duct visualisation during magnetic resonance cholangiopancreatography? results in 42 patients', *European Journal of Radiology* **49**, 258–261.
- Dauguet, J., Mangin, J.-F., Delzescaux, T. and Frouin, V. (2004), Robust inter-slice intensity normalization using histogram scale-space analysis, in 'Proc. Medical Image Computing and Computer Assisted Intervention', Saint-Malo, France.

- Davies, E. (1990), *Machine Vision: Theory, Algorithms and Practicalities*, Academic Press.
- de Smet, P. and de Vleeschauwer, D. (1997), Motion-based segmentation using a thresholded merging strategy on watershed segments, *in* 'Proc. IEEE International Conference on Image Processing', Washington, DC, USA.
- Delibasis, K., Matsopoulos, G., Mouravliansky, N. and Nikita, K. (2001), 'A novel and efficient implementation of the marching cubes algorithm', *Computerized Medical Imaging and Graphics* **25**, 343–352.
- Deschamps, T. and Cohen, L. (2001), 'Fast extraction of minimal paths in 3d images and applications to virtual endoscopy', *Medical Image Analysis* **5**(4), 281–299.
- Dominguez, G., Bischof, H. and Beichel, R. (2003), Fast 3d mean shift filter for ct images, *in* 'Proc. Scandinavian Conference on Image Analysis', Goteborg, Sweden, pp. 438–445.
- Du, Y., Parker, D., Davis, W. and Cao, G. (1994), 'Reduction of partial-volume artifacts with zero-filled interpolation in three-dimensional mr angiography', *Journal of Magnetic Resonance Imaging* **4**, 733–741.
- Dubno, B., Debatin, J., Luboldt, W., Schmidt, M., Hany, T. and Bauerfeind, P. (1998), 'Virtual mr cholangioscopy', *American Journal of Roentgenology* **171**(6), 1547–1550.
- Duda, R. and Hart, P. (1973), *Pattern Classification and Scene Analysis*, John Wiley & Sons.
- Dulyakarn, P., Rangsanseri, Y. and Thitimajshima, P. (1999), Histogram transformation based threshold selection for image segmentation, *in* 'Proc. Asian Conference on Remote Sensing', Hong Kong, China.
- Farid, M. and Murtagh, F. (2001), Vessel tracking in digital retinopathy using wavelet correlation and percolation clustering, *in* 'Proc. Irish Machine Vision and Image Processing', Dublin, Ireland, pp. 37–42.
- Faugeras, O., Hebert, M., Mussi, P. and Boissonnat, J. (1984), 'Polyhedral approximation of 3-d objects without holes', *Computer Vision, Graphics, and Image Processing* **25**, 169–183.
- Felkel, P., Bruckschwaiger, M. and Wegenkittl, R. (2001), 'Implementation and complexity of the watershed-from-markers algorithm computed as a minimal cost forest', *Computer Graphics Forum (Conference Issue)* **20**(3), C26–C35.
- Fetita, C., Preteux, F., Beigelman-Aubry, C. and Grenier, P. (2004), 'Pulmonary airways: 3-d reconstruction from multislice ct and clinical investigation', *IEEE Transactions on Medical Imaging* **23**(11), 1353–1364.
- Flasque, N., Desvignes, M., Constans, J. and Revenu, M. (2001), 'Acquisition, segmentation and tracking of the cerebral vascular tree on 3d magnetic resonance angiography images', *Medical Image Analysis* **5**(3), 173–183.

- Foley, J., van Dam, A., Feiner, S., Hughes, J. and Phillips, R. (1993), *Introduction to Computer Graphics*, Addison-Wesley.
- Frangi, A., Niessen, W., Hoogeveen, R., van Walsum, T. and Viergever, M. (1999), 'Model-based quantitation of 3-d magnetic resonance angiographic images', *IEEE Transactions on Medical Imaging* **18**(10), 946–956.
- Frangi, A., Niessen, W. and Viergever, M. (2001), 'Three-dimensional modeling for functional analysis of cardiac images: A review', *IEEE Transactions on Medical Imaging* **20**(1), 2–25.
- Frei, W. and Chen, C. (1977), 'Fast boundary detection: A generalization and a new algorithm', *IEEE Transactions on Computers* **C-26**(10), 988–998.
- Gerig, G., Kubler, O., Kikinis, R. and Jolesz, F. (1992), 'Nonlinear anisotropic filtering of mri data', *IEEE Transactions on Medical Imaging* **11**(2), 221–232.
- Ghita, O., Robinson, K., Lynch, M. and Whelan, P. (2005a), 'Mri diffusion-based filtering: A note on performance characterisation', *Computerized Medical Imaging and Graphics* **29**(4), 267–277.
- Ghita, O., Whelan, P., Robinson, K. and Ilea, D. (2005b), A 3d cad tool for body fat identification, in 'Joint Imaging Seminar', Belfast, UK.
- Gobbi, D. and Peters, T. (2003), 'Generalized 3d nonlinear transformations for medical imaging: an object oriented implementation in vtk', *Computerized Medical Imaging and Graphics* **27**, 255–265.
- Gonzalez, R. and Woods, R. (1992), *Digital Image Processing*, Addison-Wesley.
- Greenspan, H., Anderson, C. and Akber, S. (2000), 'Image enhancement by nonlinear extrapolation in frequency space', *IEEE Transactions on Image Processing* **9**(6), 1035–1048.
- Grevera, G. and Udupa, J. (1998), 'An objective comparison of 3-d image interpolation methods', *IEEE Transactions on Medical Imaging* **17**(4), 642–652.
- Griswold, M., Jakob, P., Chen, Q., Goldfarb, J., Manning, W., Edelman, R. and Sodickson, D. (1999), 'Resolution enhancement in single-shot imaging using simultaneous acquisition of spatial harmonics (smash)', *Magnetic Resonance in Medicine* **41**(6), 1236–1245.
- Guibaud, L., Bret, P., Reinhold, C., Atri, M. and Barkun, A. (1995), 'Bile duct obstruction and choledocholithiasis: Diagnosis with mr cholangiography', *Radiology* **197**(1), 109–115.
- Haignon, P., Bellemare, M., Acosta, O., Goksu, C., Kulik, C., Rioual, K. and Lucas, A. (2004), 'Depth-map-based scene analysis for active navigation in virtual angiography', *IEEE Transactions on Medical Imaging* **23**(11), 1380–1390.

- Hajnal, J., Saeed, N., Soar, E., Oatridge, A., Young, I. and Bydder, G. (1995), 'A registration and interpolation procedure for subvoxel matching of serially acquired mr images', *Journal of Computer Assisted Tomography* **19**(2), 289–296.
- Haralick, R. and Shapiro, L. (1992), *Computer and Robot Vision*, Addison-Wesley.
- Hoe, L. V., Vanbeckevoort, D. and Steenberg, W. V. (1998), *Atlas of Cross-Sectional and Projective Mr Cholangiopancreatography*, Springer-Verlag.
- Hundt, W., Petsch, R., Scheidler, J. and Reiser, M. (2002), 'Clinical evaluation of further developed mrcp sequences in comparison with standard mrcp sequences', *European Radiology* **12**, 1768–1777.
- Ilea, D., Ghita, O., Robinson, K., Lynch, M. and Whelan, P. (2005), A 3d cad tool for body fat identification, in 'Proc. European Congress of Radiology', Vienna, Austria.
- Ilea, D., Ghita, O., Robinson, K., Sadleir, R., Lynch, M., Brennan, D. and Whelan, P. (2004), Identification of body fat tissues in mri data, in 'Proc. Optimization of Electrical and Electronic Equipment', Brasov, Romania, pp. 227–232.
- Jayant, N. and Noll, P. (1984), *Digital Coding of Waveforms : Principles and Applications to Speech and Video*, Prentice-Hall.
- Ji, L., Piper, J. and Tang, J.-Y. (1989), 'Erosion and dilation of binary images by arbitrary structuring elements using interval coding', *Pattern Recognition Letters* **9**(3), 201–209.
- Johnson, C. and Dachman, A. (2000), 'Ct colonography: The next colon screening examination?', *Radiology* **216**(2), 331–341.
- Jung, C. and Scharcanski, J. (2004), 'Wavelet transform approach to adaptive image denoising and enhancement', *Journal of Electronic Imaging* **13**(2), 278–285.
- Kanitsar, A., Wegenkittl, R., Felkel, P., Fleischmann, D., Sandner, D. and Gröller, E. (2001), Computed tomography angiography: A case study of peripheral vessel investigation, in 'Proc. IEEE Visualization', San Diego, California, USA.
- Karron, D. (1992), The "spiderweb" algorithm for surface construction in noisy volume data, in 'Proc. Visualization in Biomedical Computing', Chapel Hill, North Carolina, USA, pp. 462–476.
- Ko, C., Sun, Y., Mao, C. and Lin, X. (1995), 'Three-dimensional reconstruction of biliary tract from biplane projections', *Computer Methods and Programs in Biomedicine* **47**, 21–33.

- Koito, K., Namieno, T., Hirokawa, N., Ichimura, T., Syonai, T., Yama, N., Mukaiya, M., Hirata, K., Sakata, K. and Hareyama, M. (2001), 'Virtual ct cholangioscopy: Comparison with fiberoptic cholangioscopy', *Endoscopy* **33**(8), 676–681.
- Kondo, H., Kanematsu, M., Shiratori, Y., Itoh, K., Murakami, T., Hori, M., Yasuda, I., Matsuo, M., Nakamura, H., Hoshi, H. and Moriwaki, H. (2001), 'Mr cholangioscopy with volume rendering: Receiver operating characteristic curve analysis in patients with choledocholithiasis', *American Journal of Roentgenology* **176**(5), 1183–1189.
- Kondo, S., Isayama, H., Akahane, M., Toda, N., Sasahira, N., Nakai, Y., Yamamoto, N., Hirano, K., Komatsu, Y., Tada, M., Yoshida, H., Kawabe, T., Ohtomo, K. and Omata, M. (2005), 'Detection of common bile duct stones: Comparison between endoscopic ultrasonography, magnetic resonance cholangiography, and helical-computed-tomographic cholangiography', *European Journal of Radiology* **54**, 271–275.
- Krissian, K. (2002), 'Flux-based anisotropic diffusion applied to enhancement of 3-d angiogram', *IEEE Transactions on Medical Imaging* **21**(11), 1440–1442.
- Lacroute, P. and Levoy, M. (1994), Fast volume rendering using a shear-warp factorization of the viewing transformation, in 'Computer Graphics Proceedings, Annual Conference Series (SIGGRAPH)', Orlando, USA, pp. 451–458.
- Lai, S. and Fang, M. (2003), 'A dual image approach for bias field correction in magnetic resonance imaging', *Magnetic Resonance Imaging* **21**, 121–125.
- Lapeer, R., Tan, A. and Aldridge, R. (2002), A combined approach to 3d medical image segmentation using marker-based watersheds and active contours: The active watershed method, in 'Proc. Medical Image Understanding and Analysis', Portsmouth, UK.
- Larena, J., Astigarraga, E., Saralegui, I., Merino, A., Capelastegui, A. and Calvo, M. (1998), 'Magnetic resonance cholangiopancreatography in the evaluation of pancreatic duct pathology', *British Journal of Radiology* **71**, 1100–1104.
- Lee, T.-Y. and Lin, C.-H. (2001), 'Growing-cube isosurface extraction algorithm for medical volume data', *Computerized Medical Imaging and Graphics* **25**, 405–415.
- Lin, X., Sun, Y., Hu, J., Ko, C., Chen, C. and Wang, T. (1995), 'Three-dimensional reconstruction of the biliary tract from two-dimensional biliary images', *Endoscopy* **27**, 400–403.
- Lorensen, W. and Cline, H. (1987), 'Marching cubes: A high resolution 3d surface construction algorithm', *Computer Graphics* **21**(4), 163–169.

- Lorigo, L., Faugeras, O., Grimson, W., Keriven, R., Kikinis, R., Nabavi, A. and Westin, C. (2001), 'Curves: Curve evolution for vessel segmentation', *Medical Image Analysis* **5**(3), 195–206.
- Lynch, M., Robinson, K., Ghita, O. and Whelan, P. (2004), A performance characterisation in advanced data smoothing techniques, *in* 'Proc. Irish Machine Vision and Image Processing Conference', Dublin, Ireland, pp. 123–128.
- Mangin, J.-F., Coulon, O. and Frouin, V. (1998), Robust brain segmentation using histogram scale-space analysis and mathematical morphology, *in* 'Proc. Medical Image Computing and Computer Assisted Intervention', Boston, USA.
- Mariani, A. (2001), 'Is secretin magnetic resonance cholangio-pancreatography an effective guide for a diagnostic and/or therapeutic flow-chart in acute recurrent pancreatitis?', *Journal of the Pancreas* **2**(6), 414–421.
- Metzler, V., Thies, C. and Lehmann, T. (2001), Segmentation of medical images by feature tracing in a selfdual morphological scale-space, *in* 'Proc. SPIE-Medical Imaging', pp. 139–150.
- Mohamed, C. and Auda, G. (2002), An algorithm for enhancing retinal vessel detection filters, *in* 'Proc. Medical Image Understanding and Analysis', Portsmouth, UK.
- Neri, E., Boraschi, P., Braccini, G., Caramella, D., Perri, G. and Bartolozzi, C. (1999a), 'Mr virtual endoscopy of the pancreaticobiliary tract', *Magnetic Resonance Imaging* **17**(1), 59–67.
- Neri, E., Boraschi, P., Braccini, G., Carmella, D., Gigoni, R., Perri, G., Lencioni, R. and Bartolozzi, C. (1999b), 'Mr virtual endoscopy of the pancreaticobiliary tract: A feasible technique?', *Abdominal Imaging* **24**, 289–291.
- Neri, E., Boraschi, P., Caramella, D., Braccini, G., Gigoni, R., Cosottini, M., Lodovigi, S. and Bartolozzi, C. (2000), 'Real-time volume rendering of mrcp: Clinical applications', *Magnetic Resonance Materials in Physics, Biology, and Medicine* **10**, 35–42.
- Newman, T., Tang, N., Dong, C. and Choyke, P. (2002), 'Slice-adaptive histogram for improvement of anatomical structure extraction in volume data', *Pattern Recognition Letters* **23**(1-3), 25–38.
- Nguyen, H., Worring, M. and van den Boomgaard, R. (2003), 'Watersnakes: Energy-driven watershed segmentation', *IEEE Transactions on Pattern Analysis and Machine Intelligence* **25**(3), 330–342.
- Oto, A. (2002), 'Virtual endoscopy', *European Journal of Radiology* **42**, 231–239.
- Otsu, N. (1979), 'A threshold selection method for gray-level histograms', *IEEE Transactions on Systems, Man, and Cybernetics* **9**(1), 62–66.

- Papanikolaou, N., Karantanas, A., Heracleous, E., Costa, J. and Gourtsoyianis, N. (1999), ‘Magnetic resonance cholangiopancreatography: Comparison between respiratory-triggered turbo spin echo and breath hold single-shot turbo spin echo sequences’, *Magnetic Resonance Imaging* **17**(9), 1255–1260.
- Park, H. and Yoo, J. (2001), ‘Structuring element decomposition for efficient implementation of morphological filters’, *IEE Proceedings: Vision, Image and Signal Processing* **148**(1), 31–35.
- Parker, D., Chapman, B., Roberts, J., Alexander, A. and Tsuruda, J. (2000), ‘Enhanced image detail using continuity in the mip z-buffer: Application to magnetic resonance angiography’, *Journal of Magnetic Resonance Imaging* **11**(4), 378–388.
- Pavone, P. and Passariello, R. (1997), *Mr Cholangiopancreatography: Techniques, Results, and Clinical Indications*, Springer-Verlag.
- Perona, P. and Malik, J. (1990), ‘Scale-space and edge detection using anisotropic diffusion’, *IEEE Transactions on Pattern Analysis and Machine Intelligence* **12**(7), 629–639.
- Pirsiavash, H., Soleymani, M. and Hossein-Zadeh, G.-A. (2005), An iterative approach for reconstruction of arbitrary sparsely sampled magnetic resonance images, in ‘Proc. 18th IEEE International Symposium on Computer-Based Medical Systems’, Dublin, Ireland, pp. 27–32.
- Prassopoulos, P., Papanikolaou, N., Maris, T., Gogas, C., Mouzas, J. and Gourtsoyiannis, N. (2002), ‘Development of contrast-enhanced virtual mr cholangioscopy: A feasibility study’, *European Radiology* **12**, 1438–1441.
- Prassopoulos, P., Raptopoulos, V., Chuttani, R., McKee, J., McNicholas, M. and Sheiman, R. (1998), ‘Development of virtual ct cholangiopancreatography’, *Radiology* **209**(2), 570–574.
- Preim, B., Selle, D., Spindler, W., Oldhafer, K. and Peitgen, H. (2000), Interaction techniques and vessel analysis for preoperative planning in liver surgery, in ‘Proc. International Conference on Medical Image Computing and Computer Assisted Intervention’, Pittsburg, USA.
- Pruessmann, K., Weiger, M., Scheidegger, M. and Boesiger, P. (1999), ‘Sense: Sensitivity encoding for fast mri’, *Magnetic Resonance in Medicine* **42**(5), 952–962.
- Rajon, D. and Bolch, W. (2003), ‘Marching cube algorithm: Review and trilinear interpolation adaptation for image-based dosimetric models’, *Computerized Medical Imaging and Graphics* **27**, 411–435.
- Reinhold, C. and Bret, P. (1996), ‘Current status of mr cholangiopancreatography’, *American Journal of Roentgenology* **166**(6), 1285–1295.
- Robinson, K. (2003a), Morphological reconstruction: A downhill filter, Technical Report RM-03-03, Vision Systems Laboratory, Dublin City University, Ireland.

- Robinson, K. (2003*b*), Segmentation and analysis of the biliary tree in mrcp, *in* 'Biomedical Diagnostics Research Seminar', Dublin, Ireland.
- Robinson, K. (2004), Non-linear noise suppression strategies in medical image processing, *in* 'RINCE Research Seminar Series', Dublin, Ireland.
- Robinson, K., Ghita, O. and Whelan, P. (2004), Volumetric reconstruction: Matching and merging in dicom data, *in* 'Proc. Computer Assisted Radiology and Surgery', Chicago, USA, p. 1236.
- Robinson, K., Ghita, O. and Whelan, P. (2005*a*), Intensity non-uniformity correction in multi-section whole body mri, *in* 'Proc. SPIE OPTO-Ireland', Dublin, Ireland.
- Robinson, K. and Whelan, P. (2004*a*), 'Efficient morphological reconstruction: A downhill filter', *Pattern Recognition Letters* **25**(15), 1759–1767.
- Robinson, K. and Whelan, P. (2004*b*), Narrow branch preservation in morphological reconstruction, *in* 'Proc. Irish Machine Vision and Image Processing Conference', Dublin, Ireland, pp. 42–49.
- Robinson, K. and Whelan, P. (2005), Analysis of the pancreato-biliary system from mrcp, *in* 'Proc. 18th IEEE International Symposium on Computer-Based Medical Systems', Dublin, Ireland, pp. 253–258.
- Robinson, K., Whelan, P., Ghita, O. and Brennan, D. (2005*b*), Measurement and localisation of body fat in whole body mri, *in* 'IEI Biomedical Engineering Research Award', Dublin, Ireland.
- Robinson, K., Whelan, P. and Stack, J. (2002*a*), Non-linear diffusion filtering for the enhancement of mrcp data, *in* 'Proc. IPEM Volumetric Analysis of MR Images', London, UK, pp. 16–17.
- Robinson, K., Whelan, P. and Stack, J. (2002*b*), Segmentation of the biliary tree in mrcp data, *in* 'Proc. SPIE OPTO-Ireland', Galway, Ireland, pp. 192–200.
- Sai, J. and Ariyama, J. (2000), *MRCP: Early Diagnosis of Pancreatobiliary Diseases*, Springer-Verlag.
- Salembier, P., Brigger, P., Casas, J. and Pardas, M. (1996), 'Morphological operators for image and video compression', *IEEE Transactions on Image Processing* **5**(6), 881–898.
- Salembier, P. and Serra, J. (1995), 'Flat zones filtering, connected operators, and filters by reconstruction', *IEEE Transactions on Image Processing* **4**(8), 1153–1160.
- Schnorr, C. (1999), Variational methods for adaptive smoothing and segmentation, *in* B. Jahne, H. Haussecker and P. Geissler, eds, 'Handbook of Computer Vision and Applications, Volume 2, Signal Processing and Pattern Recognition', Academic Press, pp. 459–482.

- Selle, D., Preim, B., Schenk, A. and Peitgen, H. (2002), ‘Analysis of vasculature for liver surgical planning’, *IEEE Transactions on Medical Imaging* **21**(11), 1344–1357.
- Serra, J. (1982), *Image Analysis and Mathematical Morphology*, Academic Press.
- Sijbers, J., den Dekker, A., van der Linden, A., Verhoye, M. and van Dyck, D. (1999), ‘Adaptive anisotropic noise filtering for magnitude mr data’, *Magnetic Resonance Imaging* **17**(10), 1533–1539.
- Sijbers, J., Scheunders, P., Verhoye, M., van der Linden, A., van Dyck, D. and Raman, E. (1997), ‘Watershed-based segmentation of 3d mr data for volume quantization’, *Magnetic Resonance Imaging* **15**(6), 679–688.
- Sivakumar, K., Patel, M., Kehtarnavaz, N., Balagurunathan, Y. and Dougherty, E. (2000), ‘A constant-time algorithm for erosions/dilations with applications to morphological texture feature computation’, *Real Time Imaging* **6**, 223–239.
- Smith, S. (1999), *The Scientist and Engineer’s Guide to Digital Signal Processing*, California Technical Publishing.
- Sobel, I. (1990), An isotropic 3×3 image gradient operator, in H. Freeman, ed., ‘Machine Vision for Three-Dimensional Scenes’, Academic Press, pp. 376–379.
- Soille, P. (1999), *Morphological Image Analysis: Principles and Applications*, Springer-Verlag.
- Soille, P. (2004), ‘Morphological carving’, *Pattern Recognition Letters* **25**, 543–550.
- Summers, R. (2000), Morphometric methods for virtual endoscopy, in I. Bankman, ed., ‘Handbook of Medical Imaging, Processing and Analysis’, Academic Press, pp. 747–755.
- Suri, J., Wu, D., Gao, J., Singh, S. and Laxminarayan, S. (2002), Comparison of state-of-the-art diffusion imaging techniques for smoothing medical/non-medical image data, in ‘Proc. International Conference on Pattern Recognition’, Quebec.
- Takehara, Y. (1999), ‘Fast mr imaging for evaluating the pancreaticobiliary system’, *European Journal of Radiology* **29**, 211–232.
- Tang, Y., Yamashita, Y., Abe, Y., Namimoto, T., Tsuchigame, T. and Takahashi, M. (2001), ‘Congenital anomalies of the pancreaticobiliary tract: Findings on mr cholangiopancreatography (mrCP) using half-fourier-acquisition single-shot turbo spin-echo sequence (haste)’, *Computerized Medical Imaging and Graphics* **25**, 423–431.
- ter Haar Romeny, B., ed. (1994), *Geometry-Driven Diffusion in Computer Vision*, Kluwer Academic Publishers.

- Thacker, N. and Jackson, A. (2001), ‘Mathematical segmentation of grey matter, white matter, and cerebral spinal fluid from mr image pairs’, *British Journal of Radiology* **74**, 234–242.
- Thacker, N., Jackson, A., Moriarty, D. and Vokurka, E. (1999), ‘Improved quality of re-sliced mr images using re-normalized sinc interpolation’, *Journal of Magnetic Resonance Imaging* **10**, 582–588.
- Ulichney, R. (1987), *Digital Halftoning*, MIT Press, Cambridge, MA.
- van der Heijden, F. (1995), ‘Edge and line feature extraction based on covariance models’, *IEEE Transactions on Pattern Analysis and Machine Intelligence* **17**(1), 16–33.
- van Droogenbroeck, M. and Talbot, H. (1996), ‘Fast computation of morphological operations with arbitrary structuring elements’, *Pattern Recognition Letters* **17**, 1451–1460.
- van Vliet, L. and Verwer, J. (1988), ‘A contour processing method for fast binary neighbourhood operations’, *Pattern Recognition Letters* **7**(1), 27–36.
- Vincent, L. (1993), ‘Morphological grayscale reconstruction in image analysis: Applications and efficient algorithms’, *IEEE Transactions on Image Processing* **2**(2), 176–201.
- Vincent, L. and Soille, P. (1991), ‘Watersheds in digital space: An efficient algorithm based on immersion simulations’, *IEEE Transactions on Pattern Analysis and Machine Intelligence* **13**(6), 583–598.
- Vokurka, E., Thacker, N. and Jackson, A. (1999), ‘A fast model independent method for automatic correction of intensity nonuniformity in mri data’, *Journal of Magnetic Resonance Imaging* **10**, 550–562.
- Vokurka, E., Watson, N., Watson, Y., Thacker, N. and Jackson, A. (2001), ‘Improved high resolution mr imaging for surface coils using automated intensity non-uniformity correction: Feasibility study in the orbit’, *Journal of Magnetic Resonance Imaging* **14**, 540–546.
- Wang, L. and Bhalerao, A. (2002), Model based detection of branching structures, in ‘Proc. Medical Image Understanding and Analysis’, Portsmouth, UK.
- Webb, S. (1988), *The Physics of Medical Imaging*, Adam Hilger.
- Weickert, J. (1999), Nonlinear diffusion filtering, in B. Jahne, H. Haussecker and P. Geissler, eds, ‘Handbook of Computer Vision and Applications, Volume 2, Signal Processing and Pattern Recognition’, Academic Press, pp. 423–450.
- Weickert, J., ter Haar Romeny, B. and Viergever, M. (1998), ‘Efficient and reliable schemes for nonlinear diffusion filtering’, *IEEE Transactions on Image Processing* **7**(3), 398–410.

- Whelan, P., Ghita, O. and Robinson, K. (2004), Computer assisted diagnosis (cad) for the rapid automated measurement of body fat tissue from whole body mri, *in* '12th Seminar on Theoretical Foundations of Computer Vision (Seminar No 04251: Imaging Beyond the Pin-hole Camera)', Dagstuhl, Germany.
- Whelan, P. and Molloy, D. (2000), *Machine Vision Algorithms in Java: Techniques and Implementation*, Springer-Verlag.
- Whelan, P., Sadleir, R., Robinson, K., Fenlon, H., Stack, J. and MacMathuna, P. (2002), Computer aided diagnosis of medical image data in gastrointestinal procedures, *in* '7th Annual Conference of the Healthcare Informatics Society of Ireland', Dublin, Ireland.
- Williams, D., Vitellas, K. and Sheafor, D. (2001), 'Biliary cystadenocarcinoma: Seven year follow-up and the role of mri and mrcp', *Magnetic Resonance Imaging* **19**, 1203–1208.
- Woodcock, J. and Loomes, M. (1988), *Software Engineering Mathematics*, Pitman.