



Dublin City University
School of Electronic Engineering

Dissertation submitted in partial fulfilment for the
Degree of Doctor of Philosophy

Multidimensional Image Analysis of Cardiac Function in MRI

Author

Michael Lynch

lynchm@eeng.dcu.ie

Supervisor: Prof. Paul F. Whelan

Dublin, 2006

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: Michael Lynch

ID No.: 98599054

Date: 15th May 2006

To Karen, my father and sisters, and my late mother.

Acknowledgements

I would like to express sincere gratitude to Prof. Paul F. Whelan for his support and assistance during this work. I also wish to thank Dr. Ovidiu Ghita, Dr. Kevin Robinson and all the members of the Vision Systems Group for insightful discussions over the past years.

The cardiac studies used in this thesis were kindly donated by Dr. John Murray and staff, of the Mater Misericordiae Hospital, Dublin, Ireland, Dr. Rob van der Geest, of Leiden University, Holland, with permission from Cory Swingen, University of Minnesota, Minneapolis and Dr. Dudley Pennell and staff of Imperial College Hospital in London, England. The authors would also like to thank Dr. John Murray for his advice and validation on segmentation and assessing the value of this work from a medical point of view.

Contents

Contents	i
List of Figures	iv
List of Tables	viii
List of Symbols and Abbreviations	ix
Publications Arising from this Thesis	xi
Abstract	xv
1 Introduction	1
1.1 Motivation	2
1.2 Aims	2
1.3 Challenges	3
1.4 Contributions	4
1.5 Thesis Overview	5
2 Background	7
2.1 The Heart	7
2.1.1 Morphology	8
2.1.2 Dynamics	9
2.1.3 Measurements	10
2.2 Imaging Modalities	12
2.2.1 X-Ray with Angiocardiology	12
2.2.2 Cardiac Ultrasound	13
2.2.3 SPECT/PET	13
2.2.4 Computer Tomography	14
2.2.5 Magnetic Resonance Imaging	14
2.3 MRI for Cardiac Imaging	14
2.3.1 MRI Physics	16
2.3.2 Protocols	17
2.3.3 ECG Gating	19

2.3.4	Imaging Planes	19
2.3.5	Image Formats	20
2.4	Overview of Image Processing and Analysis	20
2.4.1	Bottom-up Approaches	22
2.4.2	Top-down Approaches	24
2.5	Conclusions	26
3	Advanced Data Filtering	27
3.1	Linear Methods	28
3.1.1	Gaussian Filter	29
3.1.2	Savitzky-Golay Filter	30
3.2	Non-Linear Filters	32
3.2.1	Nonlinear Diffusion Filtering	32
3.2.2	Adaptive Smoothing	34
3.2.3	Anisotropic Gaussian Smoothing	36
3.3	Experiments and Results	37
3.3.1	Performance Characterisation of Non-Linear Filters	40
3.4	Conclusions	45
4	Statistical Partitioning of Data	53
4.1	Data Clustering	54
4.1.1	Automatic Detection of lv cavity	57
4.2	Extension to 3D	58
4.2.1	Automatic Detection of lv cavity using 3D information	59
4.3	Segmentation of epi-cardium border	60
4.3.1	First Approach: Robust-Arc epi-cardium segmentation	61
4.3.2	Second Approach: Model assisted Epi-cardium segmentation	64
4.4	Results	66
4.5	Conclusion	71
5	Boundary-Based Segmentation	75
5.1	Active-Contours	76
5.1.1	Internal Energy	76
5.1.2	External Energy	77
5.1.3	Application of Active Contours	78
5.2	Parametrically Deformable Models	79
5.2.1	Application of Parametrically Deformable Models to Medical Imaging	80
5.3	Active Shape Models	80
5.3.1	Application of ASMs to Medical Imaging	82
5.4	Active Appearance Models	83
5.4.1	Application of AAMs to Medical Imaging	84
5.5	Atlas Based Segmentation	85
5.5.1	Application of Atlas Based Methods in Medical Imaging	85
5.6	Level-set Method	86
5.6.1	Level Set Formulisation	87

5.6.2	Non-gradient based curve propagation	92
5.6.3	Introduction of <i>A-Priori</i> Knowledge	93
5.6.4	Coupling of Level Sets	93
5.6.5	Initialisation using Fast Marching	94
5.6.6	Narrow-band Methods	95
5.7	Initialisation	95
5.8	Coupling Force between Fronts	96
5.9	Improved Stopping term	98
5.10	Introduction of Priors Models	99
5.11	Extension to 4D	102
5.12	Applying level set on 3D+ t data	104
5.12.1	Modelling the temporal movement	105
5.13	Results	109
5.13.1	Testing under different motion approximation	111
5.13.2	Coupled Approach	111
5.14	Conclusions	112
6	Conclusions and Further Developments	117
6.1	Summary	117
6.2	Contributions	119
6.3	Discussion	120
6.4	Further Work	122
6.5	Concluding Remarks	123
A	Expectation Maximisation	127
A.1	EM Algorithm	127
A.1.1	Seed Generation	129
A.2	Results	130
A.3	Applications in Medical Imaging	132
B	Level-set Segmentation for Candidate Polyp extraction in CTC	141
B.1	Convex Surface Extraction	142
B.2	Level-Set Initialisation. Fast-Marching Algorithm	142
B.3	Level-Set Analysis	143
B.4	Classifier	146
B.5	Results	146
C	Mathematical Background	151
C.1	LMS Circle	151
C.2	LMS Ellipsoid	152
C.3	Splines	152
	Bibliography	155

List of Figures

2.1	Anatomy of the heart. From Gray's Anatomy [58].	8
2.2	Diagram of the vertical long-axis, horizontal long-axis and short-axis planes.	10
2.3	The basics of MRI.	17
2.4	Shows two short axis images of the heart. Bright and dark-blood. . . .	19
2.5	Top-down approach to image segmentation.	24
3.1	The four local discontinuity detectors.	35
3.2	Smoothing results.	38
3.3	Smoothing results.	39
3.4	Pixel intensities and gradient intensities after smoothing.	41
3.5	(a) Slice of the heart MRI dataset. Pixel (b) and (c) gradient intensities are plotted for the highlighted edge illustrated in image (a). . . .	43
3.6	(a) Slice of the MRCP dataset. Pixel (b) and gradient intensities (c) are plotted for the highlighted edge illustrated in image (a).	44
3.7	(a) Slice of the brain MRI dataset. Pixel (b) and gradient intensities (c) are plotted for the highlighted edge in image (a).	45
3.8	(a) Slice of the whole body MRI dataset. Pixel (b) and gradient intensities (c) are plotted for the highlighted edge illustrated in image (a).	46
3.9	3D data clustering results - heart dataset.	47
3.10	3D data clustering results - brain dataset.	48
3.11	3D data clustering results - MRCP dataset.	49
3.12	3D data clustering results whole body dataset.	50
4.1	Two iterations in the k -means clustering technique on 2D data.	56
4.2	Shows four images; a gradient-echo images before (a) and after clustering (b), and a spin-echo image before (c) and after clustering (d). . . .	57
4.3	The top three images from left to right show the original short axis image, after smoothing and after clustering.	59

4.4	Illustrating the low grayscale differentiation between the outer wall of the myocardium and other organs in the body, before (top row) and after (bottom row) data partitioning.	60
4.5	From left to right: Original unseen image, calculated edges, robust segments of epi-cardium and the complete segmentation using arcs to complete the epi-cardial boundary	62
4.6	A schematic representation of the two phases involved in the segmentation of the endo- and epi- cardium border.	63
4.7	From left to right: Original unseen image, calculated edges, robust segments of epi-cardium and the complete segmentation using an <i>a priori</i> knowledge database	65
4.8	The rendered images of (a) the end-diastole <i>lv</i> cavity, (b) the end-systole <i>lv</i> cavity, (c) and (d) the diastolic myocardium. These volumes are constructed from the true segmentation of the images excluding fat and papillary muscles.	68
4.9	Figures (a)-(b) shows scatterline plot of manual segmentation against the automatic segmentation and shows Bland-Altman plot for the left ventricle blood pool areas.	69
4.10	(a) illustrates the results using the Robust arc technique and (b) shows the results using the Prior model technique.	69
4.11	Plot shows the error frequency using a point to curve error metric.	70
4.12	The left ventricle contours obtained using our automatic segmentation method in short axis cardiac MR images.	73
5.1	Curve propagating with a force 'F', in the normal direction of the local boundary.	78
5.2	Principal axes. A 2D example where axis 1 and axis 2 are the first two eigenvectors.	82
5.3	Original image (left) following segmentation using AAM (middle) and method described in Mitchell et al. [95] (left).	85
5.4	Level-set representation of the evolution of a circle.	87
5.5	Front propagation using Fast Marching. Adapted from Sethian [134].	94
5.6	Results show the initialisation (marked in white) from a seeded Fast Marching algorithm.	96
5.7	The images above show evolution of the front at four different iterations (a) <i>iteration</i> = 0, (b) <i>iteration</i> = 5, (c) <i>iteration</i> = 10 and (d) <i>iteration</i> = 15.	97
5.8	Segmentation results of the same slice at three separate phases through the hearts cycle, (a) end-diastolic, (b) mid-diastolic and (c) end-systolic.	97
5.9	Graph of the inhibitor function where the values of $d = 6$ and $w = 4$	98
5.10	Improved stopping term on diffused data.	99
5.11	Improved stopping term in textured data.	100
5.12	Images show the probability density functions from a priori hand segmented images.	102
5.13	The images above show the segmentation using our method on the four previously unseen datasets.	102

5.14	Volume, in pixels, of left-ventricle cavity over the cardiac cycle.	105
5.15	Change of a single point on ϕ as the boundary evolves over the cardiac cycle in the short axis view.	106
5.16	Change of a single point on ϕ as the boundary evolves over the cardiac cycle in the long axis view.	106
5.17	Volume, in pixels, of left-ventricle cavity over the cardiac cycle with fitted model using an Adaptive Gaussian Model.	107
5.18	Volume, in pixels, of left-ventricle cavity over the cardiac cycle with fitted Gaussian model.	109
5.19	Results of the 4D segmentation of the left ventricle cavity boundary compared against those obtained from manual segmentation.	110
5.20	Results of the 4D segmentation of the left ventricle cavity boundary compared against those obtained from manual segmentation.	111
5.21	Selected images from a 4D sequence demonstrating a linear volume expansion.	111
5.22	The volume data from the sequence shown in Figure 5.21 with linear function fitted.	112
5.23	Estimation using prior knowledge of the Epi-cardium and Endo-cardium deformation through the cardiac cycle using inverse Gaussian curves.	112
5.24	Results from a coupled 4D segmentation of a cardiac sequence for diastolic, systolic and mid-phase for a basal (top row), mid-slice (middle row), and apical slice (bottom row).	113
6.1	Screen capture for the graphical interface.	124
A.1	Image intensity histogram overlaid with an illustration of the associated Gaussian Mixture model.	128
A.2	Figures show the short axis view of cardiac MRI (<i>a</i>) shows the original image (<i>b</i>) indicates the manually selected areas (<i>c</i>) represents the results after applying the EM using the manually picked initialisation and (<i>d</i>) is the result after applying the automatic seed picking.	131
A.3	Figures show an coronal slice from a brain MRI (<i>a</i>) shows the original image (<i>b</i>) indicates the manually selected areas (<i>c</i>) represents the results after applying the EM using the manually picked initialisation and (<i>d</i>) is the result after applying the automatic seed picking.	133
A.4	Figures show a coronal slice from a section of a full body MRI.	134
A.5	3D space partitioning using EM	136
A.6	Application of EM to Cardiac MRI data.	137
A.7	Application of EM to Brain MRI data.	138
A.8	Application of EM to full body MRI data.	139
B.1	Flow-chart of proposed algorithm	143
B.2	Extracted polyp surface (dotted) using the levelset approach based on curvature.	145
B.3	Images above show the segmentation of the convex polyp candidate. The bottom left image shows the segmentation of a fold.	147

B.4	Extracted polyp surface (dotted) for a small polyp, note the inclusion of healthy colon lumen.	148
B.5	Images above show the polyp candidate renderings of the extracted surface. Figures (a)-(c) show correctly classified polyps, where Figures (d)-(f) show correctly classified folds.	148
B.6	One of the ≤ 5 mm polyps misclassified due to the inclusion of colon wall in the surface extraction.	149
C.1	The two principle axes of a two dimensional data set are plotted and scaled according to the amount of variation that each axis explains.	153

List of Tables

3.1	The RMS of the standard deviation of the homogeneous areas for the original and each filtered image	40
3.2	Shows the edge strength and edge spread on the gradient image after each filtering. While Savitzky-Golay and Gaussian filters spread the edge, the other three maintain and even enhance the edge characteristics.	40
3.3	Performance characterisation results when the algorithms have been applied to an artificially created dataset. F1, F2, F3 denote the standard diffusion, adaptive smoothing and anisotropic diffusion respectively.	42
3.4	The RMS of the standard deviations of the homogeneous areas for the original and filtered MRI datasets used in our experiments.	46
4.1	Mean Percentage Errors \pm 1SD for manual versus automatic	67
4.2	Point to curve Errors between manual and computer segmentation for clustering technique for the endo-cardium boundary segmentation(<i>mm</i>)	70
4.3	Point to curve Errors between manual and automatic segmentation for the epi-cardium boundary(<i>mm</i>) segmentation	70
5.1	Table representing the point to curve error for <i>Method 1</i> using the Gaussian curve with adaptive variance and <i>Method 2</i> using the Expectation-Maximisation of the Gaussian parameters.	110
A.1	Changes in cluster means in the Cardiac data	132
A.2	Changes in cluster means in the brain data	134
A.3	Changes in cluster means in the whole body data	135
B.1	Control parameters used in the level-set segmentation.	147
B.2	Point-to-curve errors between manually segmented data and our method.	147
B.3	Performance Analysis for Polyp Classification. True positive (TP) and False Positive (FP).	148

List of Symbols and Abbreviations

Abbreviation	Description	Definition
CVD	Cardiovascular Disease	page 1
LV	Left Ventricle	page 2
MRI	Magnetic Resonance Imaging	page xv
RF	Radio-Frequency	page 16
CT	Computer Tomography	page 1
DICOM	Digital Imaging and Communications in Medicine	page 20
FWHM	Full Width Half Maximum	page 39
FIR	Finite Impulse Response	page 27
ASM	Active Shape Models	page 24
AAM	Active Appearance Models	page 24
PCA	Principal Component Analysis	page 81
DCT	Discrete Cosine Transform	page 81
SNR	Signal to Noise Ratio	page 22
EM	Expectation-Maximisation	page 127
CTC	Computer Tomography Colonography	page 141

Publications Arising from this Thesis

Journal Publications

Michael Lynch, Ovidiu Ghita and Paul F. Whelan (2006), **Automatic Segmentation of the Left Ventricle Cavity and Myocardium in MRI Data**, Computers in Biology and Medicine 36(4):pp389-407.

Ovidiu Ghita, Kevin Robinson, Michael Lynch and Paul F. Whelan (2005), **MRI diffusion-based filtering: A note on performance characterisation**, Computerized Medical Imaging and Graphics 29(4):pp267–277.

Michael Lynch, Dana Ilea, Kevin Robinson, Ovidiu Ghita, Paul F. Whelan (2006). **Automatic seed initialisation for the Expectation-Maximisation algorithm and its application in 3D medical imaging**. Journal of Medical Engineering and Technology. (Accepted – awaiting publication)

Michael Lynch, Ovidiu Ghita, Paul F. Whelan. **Left-Ventricle myocardium segmentation using a Coupled Level-Set with A-Priori knowledge**. Computerized Medical Imaging and Graphics. (Accepted – awaiting publication)

In Submission

Michael Lynch, Ovidiu Ghita, Paul F. Whelan. **Segmentation of the left ventricle in 3D+t MRI data using an optimised non-rigid temporal model**. Submitted to IEEE Transactions in Medical Imaging, March 2006.

Conference Publications

Michael Lynch, Tarik Chowdhury, Ovidiu Ghita and Paul F. Whelan (2005), **Determining Candidate Polyp Morphology from CT Colonography using a Level-Set Method**, European Medical and Biological Engineering Conference EMBEC 2005, November 2005, Prague, Czech Republic.

Michael Lynch, Ovidiu Ghita and Paul F. Whelan (2005), **Automatic Seed Picking Algorithm for Region-Based Segmentation of Cardiac MRI Images**, European Society of Cardiac Imaging, ESCR 2005, October 2005, Zürich, Switzerland.

Ovidiu Ghita, Paul F. Whelan and Michael Lynch (2005), **Automatic Segmentation of the Left Ventricle Cavity and Myocardium in MRI Data**, Joint Imaging Seminar between Vision Systems Group, Dublin City University and Trauma and Rehab RRG, Imaging Group, Belfast Castle, UK. 18th of May 2005. (Non-peer reviewed).

Dana Ilea, Ovidiu Ghita, Kevin Robinson, Michael Lynch and Paul F. Whelan (2005), **A 3D CAD Tool for Body Fat Identification**, Proceedings of the 17th European Congress of Radiology (ECR), March 2004, Vienna, Austria.

Michael Lynch, Ovidiu Ghita and Paul F. Whelan (2004), **Extraction of Epicardial Contours from Unseen Images Using a Shape Database**, IEEE NSS-MIC 2004 Medical Imaging Conference, October 16-22, 2004, Rome, Italy.

Michael Lynch, Kevin Robinson, Ovidiu Ghita and Paul F. Whelan (2004), **A Performance Characterisation in Advanced Data Smoothing Techniques**, IMVIP 2004 Irish Machine Vision and Image Processing Conference, September 2004, Trinity College, Dublin, Ireland.

Michael Lynch, Ovidiu Ghita and Paul F. Whelan (2004), **Comparison of 2D and 3D clustering on Short Axis Magnetic Resonance Images of the left ventricle**, CARS 2004 Computer Assisted Radiology and Surgery, June 23 - 26, 2004 Chicago, USA.

Dana Ilea, Ovidiu Ghita, Kevin Robinson, Robert Sadleir, Michael Lynch, Darren Brennan and Paul F. Whelan (2004), **Identification of Body Fat Tissues in MRI Data**, OPTIM 2004 Optimization of Electrical and Electronic Equipment, May 2004, Brasov, Romania.

Michael Lynch, Ovidiu Ghita and Paul F. Whelan. **Calculation of Ejection Fraction (EF) from MR Cardio-Images**. Paper in the Irish Machine Vision and Image Processing Conference 2003, Coleraine, Northern Ireland.

Michael Lynch. **Poster in Biomedical Diagnostics Research Seminar**, DCU, Dublin, Ireland, 2003. (Non-peer reviewed).

Abstract

Cardiac morphology is a key indicator of cardiac health. Important metrics that are currently in clinical use are left-ventricle cardiac ejection fraction, cardiac muscle (myocardium) mass, myocardium thickness and myocardium thickening over the cardiac cycle. Advances in imaging technologies have led to an increase in temporal and spatial resolution. Such an increase in data presents a laborious task for medical practitioners to analyse.

In this thesis, measurement of the cardiac left-ventricle function is achieved by developing novel methods for the automatic segmentation of the left-ventricle blood-pool and the left ventricle myocardium boundaries. A preliminary challenge faced in this task is the removal of noise from Magnetic Resonance Imaging (MRI) data, which is addressed by using advanced data filtering procedures. Two mechanisms for left-ventricle segmentation are employed.

Firstly segmentation of the left ventricle blood-pool for the measurement of ejection fraction is undertaken in the signal intensity domain. Utilising the high discrimination between blood and tissue, a novel methodology based on a statistical partitioning method offers success in localising and segmenting the blood pool of the left ventricle. From this initialisation, the estimation of the outer wall (epi-cardium) of the left ventricle can be achieved using gradient information and prior knowledge.

Secondly, a more involved method for extracting the myocardium of the left-ventricle is developed, that can better perform segmentation in higher dimensions. Spatial information is incorporated in the segmentation by employing a gradient-based boundary evolution. A level-set scheme is implemented and a novel formulation for the extraction of the cardiac muscle is introduced. Two surfaces, representing the inner and the outer boundaries of the left-ventricle, are simultaneously evolved using a coupling function and supervised with a probabilistic model of expertly assisted manual segmentations.

Finally, to fully utilise all data presented from a single 4D cardiac ($3D + t$) MRI scan a novel level-set segmentation process is developed that delineates and

tracks the boundaries of left ventricle. By encoding prior knowledge about cardiac temporal evolution in a parametric framework, an expectation-maximisation algorithm tracks the myocardium deformation and iteratively updates the level-set segmentation evolution in a non-rigid sense.

Both methods for the extraction of cardiac functions have been tested on patient data and provide positive qualitative and quantitative experimental results when compared against expertly assisted segmentations.

Chapter 1

Introduction

An estimated 17 million people die annually from Cardiovascular Disease (CVD). In general, CVD claims more lives each year than the next five leading causes of death combined. The World Health Organisation's 2002 report [119], states that 29.3% of deaths in its 191 countries were as a result of CVDs. It is these alarming statistics that has initiated the substantial research into accurate measurements of the heart for the determination of cardiac health through diagnostic imaging. The diagnosis and monitoring of cardiovascular disease, and the planning for appropriate treatment relies on accurate imaging, analysis and visualisation of the heart.

Advances in diagnostic imaging technology, in particular Computer Tomography (CT) and Magnetic Resonance (MR), has enabled greater amounts of information, in both the spatial and temporal dimensions to be generated. This high-resolution volumetric data, as a function of time, can give important physiological information about the heart. The increase in data available has made the hand annotation performed by the physician a very time-consuming task. This has pushed the advancement toward semi-automated and fully-automated approaches to quantify the results obtained from these high resolution scanners. A substantial amount of research is focusing on the accurate measurement of shape, volume and shape dynamics of the heart structure. This thesis develops the methodology for the automatic, quantitative and clinically relevant cardiac analysis in multidimensional data.

1.1 Motivation

Quantitative measurement of the left ventricle of the heart is used as a key indicator of cardiac health. The left ventricle is important as it pumps oxygen-rich blood around the body. The increased volume of data generated by the latest medical scanners presents a vast amount of high resolution volumetric data to be interpreted by the specialist. Interpreting and analyzing this large amount of data represents a tedious and time-consuming task for the cardiologist. Manual or highly supervised tracing of the cardiac boundaries is a widely used method to segment the left ventricle myocardium in current clinical studies. In one such scenario, a skilled operator, using a tracking ball or a mouse, manually traces the region of interest on each slice of the volume [100, 46, 164]. Manual slice editing suffers from many drawbacks. These include the difficulty in achieving reproducible results, operator bias, forcing the operator to view each 2-D slice separately to deduce and measure the shape and volume of 3-D structures, and operator fatigue. Since manual segmentation is labour-intensive, time-consuming and can suffer from inter- and intra-observer variability, the prospect of an automatic and accurate segmentation is highly desirable. Automatic segmentation will therefore enhance comparability between and within cardiac studies and increase accurate evaluation of volumes by allowing acquisition of thinner MRI-slices.

1.2 Aims

The main objective of this thesis is to outline the work carried out for the extraction of volumetric data and shape descriptors from MR images of the heart and to quantify the analysis against a standard of reference. Analysis of the heart function is achieved through segmentation of the left ventricle (LV). From this accurate segmentation prognostic measurements used in the diagnosis of CVDs are obtained, these include the ejection fraction (EF) of the left ventricle cavity, left ventricle mass (LVM) of the myocardium and wall thickness and thickening (WT) of the left ventricle myocardium. Therefore the expected outcome of the work is to assist the cardiologist in the prognosis of CVDs by delineating the true anatomical features present in the image and avoid making assumptions over reading what is present. Cardiac Magnetic Resonance (CMR) is the imaging modality chosen for this study. It is non-invasive, provides high temporal and spatial resolution and high contrast between blood and the myocardium.

This thesis describes the methodology that identifies the boundaries of the left ventricle of the heart with minimum user interaction. The delineation allows for the calculation of key measurements that may show anomalous heart function and therefore may indicate CVD.

1.3 Challenges

There are a number of challenges involved in the delineation of the left ventricle from MR image. Medical images are acquired using the natural and unique response of the bodies tissues to metabolic or nuclear changes. These changes are not ideal and this introduces noise into the image acquisition process in the form of image distortions.

Image distortions can be attributed to many factors, for example there is random image noise, blurring effects due to patient movement and coil intensity fall-off. Added to this, is the heterogeneous properties of the tissues, partial voluming effects between the endocardium and the left ventricle blood pool, particularly at the apex and at end-systole due to the presence of *trabeculae carneae*. In cine-MRI the variation of intensity within a tissue is increased because it may take several cycles of inducing a signal followed by measurement to image the entire sequence. This leads to gray-scale variations between image slices.

In short, there are many challenges that prevent the accurate delineation of the left ventricle myocardium due to the presence of noise in the image, heart dynamics and uneven breath-holds. The developed procedure must remove the ambiguous nature of the images while maintaining the strong anatomical features before an accurate segmentation is achieved.

As previously mentioned, modern scanners create a large amount of data in both temporal and spatial domains. Therefore the developed algorithms should utilise all the information available. Anatomical structures are represented in 3D and therefore the segmentation process of such structures are most accurately extracted using 3D algorithms. Temporal coherence can also be introduced to increase the robustness of the segmentation. Prior knowledge is often used in medical imaging analysis schemes to localise and extract anatomical features. This thesis incorporates prior knowledge in the temporal domain as a generic measure of temporal coherence which is iteratively refined, as opposed to prior models

encoded in the image domain where there may be large variation in anatomical morphology. Hence, one of the largest challenges undertaken in this thesis is to incorporate both spatial and temporal information in a meaningful way to improve the accuracy and robustness of the segmentation.

1.4 Contributions

Based on the challenges outlined in the previous section, the major contributions of this work lie in the segmentation of the left ventricle myocardium in multidimensional MRI data. There are a number of stages that are adopted and these constitute the major contributions to this work.

Firstly, in order to reduce the inherent noise associated with MRI images, a performance characterisation of advanced smoothing filters is performed. The characterisation is performed in both 2D and in 3D.

A novel method for segmentation and localisation of the left ventricle blood pool using an unsupervised clustering technique is presented in Chapter 4. This technique is approached in both a slice by slice and volume image context. After the segmentation of the left ventricle blood pool cavity, an heuristic approach is developed to extract the outer walls of the myocardium in each image slice. This technique is based on gradient information in the image and where such information is lacking, a prior model of previously segmented myocardium images is incorporated into the segmentation. While this approach gives favorable results in good quality data, introduction of temporal information into this framework is cumbersome. Therefore a more involved approach is proposed that can easily incorporate extension to 4D data.

Describing a contour as a particular instance of a higher dimensioned function in the Eulerian space has many advantages. Firstly, errors associated with sampling of the contour are eliminated as the approach is non-marker based. The deformation is numerically stable and has the ability to handle topological changes during the deformation. In Chapter 5 a novel methodology of level-sets is introduced that evolves a coupled surface, representing the inner and outer wall of the left ventricle myocardium. The deformation is guided using a probabilistic model of manual segmentations.

Finally, the Eulerian formulation of the level-set is exploited in a novel fashion to extend the deformation in a supervised way to 4D. A temporal model is constructed for each grid point in Eulerian space using prior knowledge about cardiac motion. This parametric model is then iteratively refined during the segmentation process to capture the myocardium boundaries. This novel approach has many advantages. Firstly, it performs a temporal smoothing of the segmented contours through the cardiac cycle that follows the temporal model from the user defined motion model. Secondly, the model is defined in temporal space and is therefore free from the highly variable anatomical features of the cardiac muscle in image space. The human left ventricle has a harmonic pumping motion which can be modelled for both healthy and unhealthy hearts and is relatively independent of the highly variant cardiac anatomy. Thirdly, initial estimates for the parametric model found through a fast marching algorithm and the parameters are then iteratively updated using an expectation-maximisation algorithm.

Hence, segmentation of the left ventricle in cardiac MRI data is approached in a systematic way, at each step increasing the dimensionality of the problem and incorporating more knowledge and information in more involving methodologies. Initial approaches are based on signal intensity values in 2D and 3D images for the segmentation of the cardiac blood pool followed by a 2D model assisted segmentation of the outer wall of the left ventricle myocardium. In the second phase, a coupled deformation of surfaces is introduced for both the inner and outer boundary which is also guided by models of manually annotated models. In the final stages, temporal information is introduced in a knowledge based way to model the left ventricle motion and ensure smooth temporal transition of segmented surfaces.

1.5 Thesis Overview

This thesis details the progression of ideas for the segmentation of the left ventricle of the heart from multi-dimensional MRI data. Based on the challenges outlined in the previous sections, the thesis details each of the steps.

Chapter 2 gives a background to the problem. This chapter deals mainly with cardiac anatomy, dynamics and clinically relevant measurements associated with diagnosing CVDs. An introduction to some of the most common medical imaging modalities, an in-depth discussion on MRI and the difference

acquisition procedures. And finally a brief overview of image processing and in particular on medical image analysis.

Chapter 3 details the methods employed in image noise suppression. The advantage of non-linear smoothing over linear smoothing is investigated in 2D before a performance characterisation of three non-linear filters applied to MRI data is performed in 3D.

Chapter 4 gives the particulars on how statistical based segmentation algorithms can be used to accurately measure the left ventricle blood pool volumes and consequently the measurement of ejection fraction. Heuristic methods for the segmentation of the outer boundary of the cardiac muscle in 2D are detailed and deficiencies associated with this approach are discussed.

Chapter 5 introduces boundary based methods as an alternative approach in order to circumvent some of the limitations of the statistical based approaches. An overview of previous approaches are detailed. Gradient based level-set segmentation approaches are proposed as an accurate method of segmentation in higher dimensioned data. A novel method for the segmentation of $3D+t$ (4D) is introduced. This method employs a parametric prior model encoded in the temporal domain which is iteratively updated using a expectation-maximisation algorithm.

Chapter 6 concludes the thesis, outlining the novel methods developed, discussing the results obtained and recommending how these approaches may be advanced or can be applied to other temporally variant anatomical structures.

Appendix 1 describes the application of an expectation-maximisation algorithm for partitioning data in medical images.

Appendix 2 details the application of the level-set algorithm to perform accurate segmentation of polyps in CT colonography.

Chapter 2

Background

This chapter gives a brief overview of three distinct areas involved in this project. Firstly, an introduction to the heart, its function, some terminology and the clinical measurements that are to be extracted from cardiac specific images acquired of the thorax. Next, an overview of the imaging modalities used in cardiac imaging, moving to explain why MRI is the chosen modality for this investigation. This is followed by a discussion on the basics of MRI also mentioning the main protocols in common use will ensue. Finally, a background is given on work that has been investigated in the image processing area and in particular in the field of medical imaging and cardiac analysis. In this section a review is given of the most relevant literature published on the subject.

2.1 The Heart

The heart can be thought of as the “pumping station” of the body. Situated between the third and sixth ribs in the center of the thoracic cavity of the body, the heart is a hollow conically shaped muscle about the size of a clenched fist, 12-13cm along its major axis and 7-8cm at its widest point [101, 58]. It rests on the diaphragm between the lower part of the two lungs. Its function is to pump oxygen and nutrient rich blood around the cardiovascular system, where it supplies the oxygen to the cells. The de-oxygenated blood then returns to the heart before being pumped to the lungs to collect more oxygen. The oxygen-rich blood then proceeds back to the heart before it is sent around the cardiovascular network again.

2.1.1 Morphology

The heart is a hollow muscle that is divided internally into four separate chambers. The heart muscle itself is divided into three layers, the *epi-cardium* is the outer protective layer, the middle muscular layer is referred to as the *myocardium* while the inner layer is known as the *endo-cardium*.

The heart is divided down the center with a strong muscle wall called the *interatrial-interventricular septum* into a cylindrical left side and a more crescent shaped right side. The right side of the heart is again divided in two with the upper atrium or auricle separated from the lower ventricle with a one-way valve called the - *Tricuspid* valve. Similarly, the left side is divided into the left-atrium and left-ventricle with the *Bicuspid* or *mitral* valve. In order of size, the left-atrium is the smallest chamber, holding approximately $45ml$ at rest, and having a wall thickness of approximately $3mm$.

This is followed by the right-atrium, with a typical capacity of $63ml$ and wall thickness of $2mm$, the left ventricle with a $100ml$ capacity and a wall thickness as high as $12mm$ and finally the right atrium which can hold up to $130ml$ with a wall thickness of $4mm$. The varying wall thickness is a result of the normal operating pressure of each of the chambers and is explained in the next section. Each of the chambers has an associated major vessel either supplying blood or transporting blood away. The left ventricle has the *aorta*, joined using a one-way *aortic* valve, the left atrium has the *pulmonary vein*, the right ventricle has the *pulmonary artery* which is closed with the *pulmonary semi-lunar* valve while the right atrium is supplied from the *superior* and *inferior venae cavae* and the *coronary sinus*. Disease associated with the valves is mainly caused from *congenital* abnormalities, degeneration or infection and can result in leakage through the valve. In the most common type of valvular disease the *mitral* valve or *aortic* valves are frequently affected. With *mitral dysfunction*, the blood is allowed to




Figure 2.1: Anatomy of the heart. From Gray's Anatomy [58].

regurgitate back to the left *atrium* increasing pressure in that atrium.

Also present in both ventricles are thin columns of muscle along its length, these are referred to as *trabeculae carneae*. The papillary muscles are thin muscles protruding from the walls of both ventricles and are connected to their respective *atrioventricular* valves. Both the *trabeculae carneae* and papillary muscles are more pronounced in the left-ventricle.

The heart itself needs to be supplied with oxygen-rich blood and the measure of blood supplied to the heart is called *myocardium viability*. Coronary circulation is achieved through two main arteries, the *right coronary artery* and *left coronary artery*. Both of these arteries stem from the *ascending aorta*. Blood is returned via the *coronary sinus*. Over time, the *coronary arteries* may become clogged from a build-up with *fat*, *cholesterol*, *triglycerides* and *calcium*. This build-up prevents the coronary arteries from functioning properly, and interferes with the delivery of an adequate supply of blood to the heart muscle. Ninety five percent of all coronary artery disease is due to this *atherosclerosis*, the build-up of fatty substances. The insufficient blood supply to the heart is called *ischemia*. Myocarditis is inflammation of the myocardium, the muscular part of the heart. It is generally due to infection (viral or bacterial). It may present with rapid signs of heart failure.

For clinical evaluation of cardiac anatomy and motion, a standard left ventricle representation has been developed [24] whereby the cardiac muscle is divided into 17 segments, Figure 2.2. These 17 segments creates a distribution of 35%, 35% and 30% for the basal, mid cavity and apical thirds of the heart, which is close to the observed autopsy data.

2.1.2 Dynamics

The heart has two distinct phases, *diastole* and *systole*. The diastole, or filling cycle, occurs when the muscle relaxes and the left and right ventricles fill with blood from the respective atria, this can take 480 *ms* of the complete 750 *ms* cycle. During this cycle the tricuspid and mitral valves are open while the aortic and semi-lunar pulmonary valves are closed. When the end-diastole volume (EDV) has being reached the heart sends an electronic pulse for the systole cycle to start. The systole phase is much shorter where the muscle contracts and closes the tricuspid and mitral valves while opening the aortic and pulmonary valves.




Figure 2.2: Diagram of the vertical long-axis, horizontal long-axis and short-axis planes showing the name, location and anatomical landmarks for the selection of basal, mid-cavity and apical short axis slices for the 17 segment system.

Approximately half of the ventricles capacity is emptied during the systolic phase, the remainder is called the cardiac reserve volume. The cardiac cycle is timed using the hearts own intrinsic nervous system and can survive *in-vitro* for several hours. The main switch in the nervous system is called the *Sinus Node*, this triggers the *AV Node* which in turn connects to the *Bundle of His* to conduct the triggering pulse through the septum of the heart.

The ventricular working pressures are much higher than atria pressures. The left and right ventricles have an approximate working pressure of 140 mmHg and 40 mmHg respectively, this gives rise to the left ventricle muscle being three times thicker than that of the right ventricle.

2.1.3 Measurements

The volumetric data acquired with time can produce a number of important measurements that can indicate disease of the heart [48]. While these descriptors are well documented in research literature [43] they are not always clinically assessed in everyday practice.

In practice, clinical measurements still rely on global volumetric measure such as left-ventricle end-diastolic volume (EDV), end-systolic volume (ESV) and mass (LVM). These are then used in conjunction with other measurements to calculate the stroke volume (SV), cardiac output (CO) and ejection fraction (EF). The inclusion of papillary muscles and trabeculations is still undecided and is usually dependent on the center performing the scan. Recent research [138] has shown that the difference between subtracting and not subtracting the papillary muscles and trabeculations has little clinical relevance when calculating the left ventricular volumes and ejection fractions. The extraction of the epi-cardium boundary aids the accurate measurement of wall thickening (WT) over the cardiac cycle, this can indicate areas with reduced contractibility.

- **End-diastolic volume (EDV)** and **End-systolic volume (ESV)** is the amount of blood contained in the left ventricle at its maximum and minimum respective capacities, measured in *ml*.
- **Left ventricle mass (LVM)** is an important indicator for left ventricle hypertrophy (LVH). LVH is an enlargement of the muscle fibers of the left ventricle, mainly around the interventricular septum. LVH is a late complication of congestive heart disease and cardiac arrhythmias. The LVM is measured to be the volume (cm^3) enclosed by the epi-cardium boundary minus the left ventricle cavity and multiplied by the density of muscle tissue ($1.05g/cm^3$).

$$LVM = 1.05 \times (V_{epi} - V_{endo}) \quad (2.1)$$

- **Stroke Volume (SV)** is the volume (cm^3) of blood ejected from the left ventricle between the end-diastole and the end-systole. This value can then be normalised to body surface area and called the stroke-volume index (SVI).

$$SV = V_{endo}(t_D) - V_{endo}(t_S) \quad (2.2)$$

where V_{endo} is the volume of the left ventricle cavity, $V_{endo}(t_D) = \max_t[V_{endo}(t)]$ at end-diastole and $V_{endo}(t_S) = \min_t[V_{endo}(t)]$ is the end-systole.

- **Ejection Fraction (EF)** is the percentage of blood ejected from the left ventricle with each heart beat, and can be represented by the equation:

$$EF(\%) = \frac{V_{endo}(t_D) - V_{endo}(t_S)}{V_{endo}(t_D)} \times 100 \quad (2.3)$$

- **Cardiac output (CO)** is the amount of oxygenated blood supplied to the body (ml/min). This can be normalised to the body surface area and called the cardiac index (CI). The calculation is shown in Eq. 2.4 where HR is the heart rate.

$$CO = (V_{endo}(t_D) - V_{endo}(t_S)) \times HR \quad (2.4)$$

- **Wall thickening (WT)** is the measurement of the myocardium thickness over time (mm_t). This can give a direct indication to the myocardial viability and therefore can forecast *ischemic* disease. The wall thickness can be computed using the centerline method, along lines that are perpendicular to a curve that is equidistant from both the endo- and epi-cardial boundaries.

2.2 Imaging Modalities

In this section the reader is given a brief introduction into the imaging modalities that are commonly used for cardiac analysis. A brief description of each method is given along with their advantages and disadvantages. This is followed by a brief discussion on the suitability of MRI in cardiac analysis, a more in-depth background describing some of the physics involved and the different protocols in clinical use.

2.2.1 X-Ray with Angiocardiology

X-ray angiocardiology (XRA) is a projection image of the left ventricle usually in the oblique view after a contrast agent has been injected into the ventricle via a pigtail catheter. In XRA the contrast agent is not uniformly spread throughout

the left ventricle because of the dilution with blood at the mitral valve. It may not reach to the apex of the heart and there is also a limitation on the amount of contrast agent used due to the risk to the patient. Surrounding structures such as ribs can be removed from the image using Digital Subtraction Angiography (DSA). DSA involves a temporal subtraction where the image obtained without a contrast agent is subtracted from the contrast image. Complications associated with cardiac angiography are *cardiac arrhythmias* (irregular heartbeat) and *embolism* (by dislodging plaque from the artery wall while treading the catheter). XRA can be used for the calculation of the ejection fraction using geometric approximations [43] but cannot be used for the calculation volumes or delineating the epi-cardial boundary.

2.2.2 Cardiac Ultrasound

Cardiac ultrasound is a tomographic imaging system, it is relatively cheap, non-invasive and can image on arbitrary planes [24]. It gives low contrast when compared to MR and X-ray, cannot image through gaseous mediums and has a low signal-to-noise ratio due to frequency attenuation in the tissue. The signal-to-noise ratio is further reduced in cases where the patient presents obesity. 3D ultrasound [88, 125] has been introduced and can quantify ventricular volumes and myocardium mass without the need for geometric models. Ultrasonographic heart images suffer from the need for acoustic windows, operator subjectiveness and are often characterised by weak echoes, echo dropouts and high levels of speckle noise causing erroneous detection of the LV boundaries.

2.2.3 SPECT/PET

Single-photon Emission Computed Tomography (SPECT) and Positron Emission Tomography (PET) scans are part of the non-invasive nuclear imaging techniques. SPECT was introduced in the 1970's and is used to detect subtle metabolic changes in the organ under investigation. PET was introduced shortly after SPECT because of its increased temporal resolution. Both methods work by the injection of radionuclides (radioactive isotopes) into the organ under investigation. These radionuclide tracers are absorbed at different rates by the healthy and dysfunctional muscle. While these methods are good for the measurement of metabolic changes, the resolution does not lend to the delineation of anatomical structures in the image.

2.2.4 Computer Tomography

A traditional X-Ray Computer Tomography (CT) scan is an X-Ray procedure which combines many X-Ray images with the aid of a computer to generate cross-sectional views of the body. CT is increasingly used in cardiac analysis. It provides increasing resolution in data with the introduction of the 64 slice CT and is particularly useful for evaluating coronary atherosclerosis. With conventional CT, cardiac motion causes blurring and artifacts in the image but advances such as Electronic Beam Computer Tomography (EBCT), Ultrafast[®] and Dual-Source CT have increased the acquisition time sufficiently to capture the beating heart. These approaches can be gated to capture information at a precise phase in the hearts cycle. However CT suffers from low contrast between soft tissues like blood and myocardium and the patient is exposed to ionising radiation.

2.2.5 Magnetic Resonance Imaging

Magnetic Resonance Image (MRI) was first introduced in medical imaging in 1971. Since its introduction cardiac magnetic resonance (CMR) has progressively improved with increased spatial and temporal resolution. CMR is considered by some authors [43, 128, 130] to be the standard of reference for evaluating the cardiac function. MR has proved to be more accurate than echo-cardiology and cardiac angiography in the calculation of the ejection fraction and also has shown superior results in endo-cardium border segmentation [128]. MRI boasts a wide topographical field of view and high contrast between soft tissues without the need for a contrast agent. It is non-invasive with high spatial resolution and can be gated using an *electrocardiogram* (ECG) at different phases during the hearts pulse [158, 102]. However, it can suffer from noise and grey scale variation between adjacent slices. More details are discussed in the next section.

2.3 MRI for Cardiac Imaging

Cardiac Magnetic Resonance (CMR) has very clear advantages over the other imaging modalities, discussed in the previous section. It has proved to be more accurate [43] for the evaluation of cardiac function measurements due mainly to its independence from any geometric assumptions about the ventricle shape and can be used without the need for a contrast agent. Cine-MR has being introduced to capture a collection of images over one or several phases of the cardiac cycle. MRI tagging has been introduced to obtain heart twist through the cardiac cycle.

The use of MRI has many benefits over other types of acquisition.

- Images of soft-tissue structures such as the heart and major vessels are clearer and more detailed when compared to other imaging methods.
- The detail of MRI makes it an invaluable tool in early detection and evaluation of coronary disease.
- Even without the use of contrast material, MRI often shows sufficient detail of the heart to be valuable in diagnosis and treatment planning.
- When it is used, MRI contrast material is less likely to produce an allergic reaction than the iodine-based materials used for conventional X-Rays and CT scanning and does not contain the radioisotopes used in nuclear medicine exams.
- MRI enables the detection of abnormalities that might be obscured by bone tissue with other imaging methods.
- MRI provides a fast, noninvasive and often less expensive alternative to other techniques of cardiac diagnosis.
- MRI can help physicians evaluate the function, as well as the structure, of the heart muscles and valves.
- MRI does not require exposure to radiation or the introduction of radioisotopes in the body.
- MRI has the advantage that images can be obtained from arbitrary planes.

The use of MRI also has the following associated risks or limitations.

- An undetected metal implant may be affected by the strong magnetic field.
- MRI is generally avoided in the first 12 weeks of pregnancy. Doctors usually use other methods of imaging such as ultrasound on pregnant women, unless there is a strong medical reason to conduct an MRI exam.

In this section an overview of the basic physics of MRI is given to the reader, the imaging planes used in a conventional heart examination are shown and finally the different protocols that have been introduced with their advantages and disadvantages.

2.3.1 MRI Physics

MRI applies a Radio Frequency (RF) excitation pulse to the protons that sit in a static magnetic field. When the protons return to a state of equilibrium they emit a quantified energy as an RF signal. This signal is then collected and analysed. On the scan this corresponds to high intensity meaning high signal collected by that group of protons.

MR uses the magnetisation effects of hydrogen to create the intensity map, or image. Hydrogen is an abundant atom in almost all biological systems. As can be seen in figure 2.3 these atoms do not naturally align in any particular direction but instead spin around their own axes in arbitrary orientations and therefore the magnetic effect is negligible. If however, a strong static magnetic field B_0 is applied to these atoms they align themselves either in the parallel or anti-parallel direction to the direction of the field (in most cases they align parallel). The atoms do not strictly align parallel to the magnetic field but at a small angle θ , or flip-angle, and precess around the magnetic field at a frequency f , or the Larmor frequency. If an external frequency B_1 is pulsed at the Larmor frequency perpendicular to B_0 the atoms absorb the energy and tend to precess away from B_0 and toward B_1 momentarily. When the pulse has finished the atom returns to the static magnetic field and releases the energy as a small RF signal. This signal is collected and used to produce the image. TE is the echo time, the time at which the signal echoes are obtained after the excitation pulse. TR is the repetition time, the time in which the excitation pulse is repeated to obtain the image lines.

In order to locate the position of the signal spatially, a third magnetic field called a gradient magnetic field that varies the magnetic field strength with respect to its position is applied. The most common type of reconstruction used to create the image is a two-dimensional Fourier transform. Measurements are taken at important relaxation times $T1$ and $T2$. $T1$, or spin-lattice relaxation time, is the settling time for the atoms to return to equilibrium after being disturbed by the RF pulse while $T2$, also called the spin-spin relaxation time, is the decay of the RF signal after it has being created, both these measures are tissue dependent which gives the MRI its ability to distinguish between different tissues in the body. For example water has a longer $T1$ time when compared to fat because it does not give up its energy as quickly as fat, similarly water has a longer $T2$ time when compared with fat. Using these and other properties a host

of different imaging protocols have been devised to optimize image quality.




Figure 2.3: The basics of MRI. Figure (a) shows random hydrogen atoms, (b) shows the aligned atoms after the introduction of a static magnetic field B_0 , (c) shows results after applied RF pulse B_1 and (d) plots the T_1 and T_2 decay graphs.

Image derived from U.S. Department of Health and Human Services, Food and Drug Administration, Center for Devices and Radiological Health, Magnetic Resonance Working Group.

2.3.2 Protocols

Echo planar imaging (EPI) is a fast imaging technique, introduced in the late 1970s that reads multiple lines of the image with just one excitation pulse. This method greatly increased the speed of MRI meaning images could be acquired in fractions of a second compared to minutes with early MRI.

Gradient Echo

Gradient Echo images are also called bright-blood images due to the high signal intensity of the blood. Gradient echo images are acquired using either T_1 and T_2 weighting or a combination of both. The RF excitation pulse is applied once and the signal is obtained after a short TE , usually between 1-10 ms . Due to the low TE time, the blood does not have the opportunity to flow away from the imaging plane, explaining the high intensity in the blood but this can cause heterogeneity

within the blood-pool especially pronounced along the endo-cardium and the *mitral valve*. TR s are also short, $< 50\text{ ms}$, which allows for rapid acquisition cine-MR.

Spin Echo

Spin-Echo, or dark-blood sequences, apply two RF pulses, usually at both 90° and 180° . This second pulse, applied at $TE/2$, reorients the atoms. It is the resulting echo signal that is used to construct the image. The TR for spin echo sequences is much higher than that of gradient echo. Spin Echo is therefore not used for the generation of cine-MR sequences because of this increased acquisition time. Spin-echo does however provide higher contrast-to-noise ratio (CNR) between the blood and the myocardium. Fast spin-echo sequences, also called turbo spin-echo, Rapid Acquisition and Relaxation Enhancement (RARE), increase the temporal resolution but at the cost of soft tissue contrast. Typical images taken from both spin-echo and gradient-echo images can be seen in figure 2.4.

Balanced Sequences

Steady-state free precession (SSFP) method has been recently developed where the contrast of the tissues depends more on the $T1$ and $T2$ contrast and less on the flow dynamics. It involves rapid excitation using the RF pulse, never allowing the MR signal to completely decay. This means that the images can have the high tissue contrast of $T1$ and the high blood tissue contrast of $T2$ -weighted acquisition. There are a whole family of SSFPs which include Balanced Fast Field Echo (bFFE), Balanced Turbo Field Echo (bTFE), Fast Imaging with Steady Precession (TrueFISP), Completely Balanced Steady State (CBASS) and Balanced SARGE (BASG).

Recently, methods such as Sensitivity Encoding (SENSE) have been introduced to speed up imaging and therefore increase the resolution. This method is based on the use of multiple RF coils and receivers. Other advances include Prospective Acquisition and Correction (PACE) which allows free breathing during the exam by detecting the diaphragm and correcting for its movement. MRI tagging has been a well documented method of evaluating the twist and torque of the myocardium during the heart-beat by non-invasively placing a grid, known as tagging, on an image with changing radio frequencies.




Figure 2.4: Shows two short axis images of the heart. (a) gradient-echo image; $TE= 1.6\text{ms}$, $TR= 3.2\text{ms}$, flip angle = 60° and (b) spin-echo image, $T1$ -weighted approximate times of $TE=10\text{-}20\text{ms}$ and $TR=300\text{-}600\text{ms}$. Image (b) used courtesy of the Auckland MRI Research Group (<http://www.scmr.org/education/atlas/intro/>).

2.3.3 ECG Gating

An electrocardiogram (ECG \ EKG) is a recording of the hearts electrical pulses as it stimulates the myocardium. In imaging, ECGs are used to establish the hearts R-wave which is a high peak wave, in a normal patient and depending on acquisition, coming between the Q and S wave and indicates the start of the myocardium contraction. This is used to trigger the imaging at the correct time in the hearts phase. ECG gating suffers in MRI from a phenomenon called the *magneto-hydrodynamics* effect where the signal gets distorted when the patient enters the static magnetic field. However, this can be eliminated with Vector-cardiogram (VCG) which uses multiple ECG-channels to accurately detect the R-wave.

2.3.4 Imaging Planes

MRI has the advantage that images can be obtained in arbitrary planes. This is useful to obtain the best orientation for the images to be viewed, as the orientation of the heart changes from patient to patient. Traditional views in cardiac imaging are sagittal, which divides the body into left and right, orthogonal where the images are taken from the head to feet direction and long axis where the images are oriented to show the best view of the four chambers of the heart (see

figure 2.4). In practice the orientation for the evaluation of the left ventricle is the short-axis view as it gives the best view of the left ventricle chamber for volume calculations. The short-axis is the plane perpendicular to center line running from the mitral valve to the apex of the heart.

2.3.5 Image Formats

All images used in this work were encoded in the DICOM (Digital Imaging and Communications in Medicine) format, taken along the short axis plane traversing the left ventricle cavity from the base to the apex of the heart as shown in Figure 2.2.

2.4 Overview of related Image Processing and Analysis Techniques

Image processing first evolved in the late 50s and early 60s where simple algorithms were implemented in hardware. Many of these implementations were derived from signal processing ideas. It wasn't until the mid to late 1960's and early 1970s that digital image processing became an active area for research. Applications such as the NASA 1964 project aimed to remove imperfections from lunar images returned on the Ranger 7 expedition. It was at the early stages of image processing that ideas into medical image analysis were first investigated and many of the first projects in image processing were funded by the National Institute of Health (NIH) as well as the National Science Foundation (NSF) in the US. One of the earliest publications on medical image analysis by Strauss et al. [153] where nuclear images of the heart were obtained using a scintiphotographic method and the computer semi-automatically outlined a region of interest for the quantitative measurement of the ejection fraction.

Image processing is inextricable tied to the advancement of the computer and it was in the past and still is the increase in computational power that drives the level of complexity entailed in image processing techniques. As the discipline of image processing grew, more sophisticated algorithms were developed to achieve more complex tasks. Today, the major problems where image processing are in the areas of:

- Photography

- Satellite Imaging
- Face Detection
- Medical Imaging
- Natural Scene Analysis

The field of image processing includes acquisition where the main challenges are the reduction of distortion and develop sensors that aim to improve the signal-to-noise ratio (SNR). Image storage has always stretched the boundaries of computer memory capacities and therefore image compression in both still and video data has also attracted researchers. Post processing of images include geometric transformations of the object or coordinate system, colour corrections for image enhancements, distortion corrections to rectify camera inaccuracies, noise suppression and filtering to correct sensor inaccuracies, edge detection to define boundaries between objects in the image, segmentation of an image into distinct regions and pattern recognition for the localisation and classification of objects from a scene.

Many of these operations that are common in image processing and image analysis can also be implemented in medical image analysis, but with subtle differences. For instance, problems such as illumination difficulties are replaced by more acquisition specific limitations such as coil intensity fall off in MRI. Many image processing and computer vision tasks deal with the extraction of 3D data from stereo images but in medical image analysis, very often with modern scanners, the data can easily be reconstructed into 3D and therefore accurate shape recovery and tracking in 3D is the major issue. Pattern recognition is also implemented in medical images using prior knowledge of anatomical shape or structure.

The main issues that drive research in medical image analysis are:

- Image segmentation
- Image matching / registration
- Motion tracking

The in plane resolution of modern scanners are in the domain of $0.5\text{-}2.5\text{mm}$ for CT and MRI scanners, therefore, medical image analysis is performed at macroscopic or organ level as opposed to microscopic or atomic level.

Image segmentation deals with the accurate dividing up of an image or a volume into smaller relevant collection of pixels or voxels. In the case of medical image analysis these smaller subgroups generally represent anatomical features such as tissue, blood or bone. It is the methods by which these divisions can be made that is the subject of segmentation. Segmentation is a deceptively difficult problem to solve and many approaches require much user intervention such as live-wire techniques [164, 46]. Image segmentation has received a significant amount of attention in the past number of decades. With the exponential growth of computational power and memory, more complex algorithms can be applied to larger amounts of data. There are a number of proposed techniques which can be broadly classified in bottom-up approaches and top-down approaches.

2.4.1 Bottom-up Approaches

Bottom-up approaches perform the separation normally based on no prior knowledge and divisions are made based on the intensity or gray level values. The most basic form of bottom-up or intensity based segmentation is thresholding. Thresholding is a binary classification problem where all elements in an image with gray level values higher than a user defined number are classified as one object and all elements with gray level value below are classified as a second object, adaptive methods to find the threshold values have also been evaluated [175, 57]. Other methods for selecting thresholds include histogram analysis and global and local thresholding. Thresholding methods are susceptible to noise in low contrast images and are therefore normally combined with some morphological operators such as opening and closing to remove outliers. Other bottom up approaches search for divisions of objects within the image called edge detectors. This division is characterised by a difference in local grayscale values. This differential operator can give information regarding the strength of the division returned by the gradient and also the direction returned by the orientation. Common edge detector operators include Canny and Sobel. Similar to thresholding, these methods are limited in images with low Signal-to-Noise Ratio (SNR). In these circumstances, methods such as edge linking [55] and non-maximum suppression [20] may be employed.

More advanced methods involve partitioning the image into a greater number of final classes, how best to classify the objects into the appropriate classes and how to determine the appropriate number of classes in a specific image. Statis-

tical partitioning of images into higher than two classes is a very active area of research.

Clustering methods have evolved and try to minimise the variance of pixels within clusters while maximising the variance between clusters. Inclusion into a certain cluster may be based on gray level value or a number of other metrics. Cluster membership may be a hard classification, as is the case in k -means clustering, or a soft membership classification, as is the case with fuzzy c -means clustering or Expectation-Maximisation classification [40, 14]. In the first case, each element is assigned to a particular class but on the other hand, in a soft classification, membership to a cluster is given as a probabilistic measurement. More advanced clustering methods use multiple scales [136] to alleviate over segmentation whereby the object to be extracted is divided into multiple regions.

Delibasis et al. [38] implemented a number of standard bottom-up techniques for evaluating the segmentation of the left ventricle cavity from cine MR sequences in a small number of normal and abnormal patients. These included an adaptive region growing technique from a seeded position, where the new voxels are added to the object of interest if its value is close to the mean of all the voxels contained in the object. A k -means algorithm, which partitions voxels in feature space into a predefined number of classes [65] using a distance metric of each voxel feature from the class feature average. A fuzzy C -means algorithm [118], similar to the k -means with the introduction of a fuzzy function which defines the probability of membership to each class. A neural network based Self Organizing Maps (SOMs) based on Kohonens [75] work. Delibasis et al [38] proved that k -means gave the most robust results on average over the normal and abnormal data when compared to manual segmentations.

A more in-depth discussion on statistical partitioning of data is continued in Chapter 3 but these methods may suffer in noisy images where there is a significant variation in gray scale values. In medical segmentation, its is often the task to extract a closed structure, however these partitioning algorithms based on intensity values do not take spatial relationships into consideration. This is why many researchers have investigated the value of approaching the problem from a top-down angle.

2.4.2 Top-down Approaches

Top-down approaches apply some information about the desired results and then tries to perform some sort of fitting and deformation to achieve the final segmentation and aims to closer resemble a cognitive approach to segmentation. Template matching is an example of top-down segmentation in which a predefined shape is fitted to the data by means of scaling, rotating and translating (see Figure 2.5). This method performs a search of the image using a predefined template and tries to fit the template to gradients in the image which minimises the error and maximises the overlay. Of course, in this case, the template is a rigid structure and can only be used for localisation of the object and only in cases where there the template does not differ greatly from the final object to be located.




Figure 2.5: Top-down approach to image segmentation. (a) Shows the prior model to be fitted to the data in (b) giving the resulting image shown in (c) [56].

One significant advancement on this idea Active Shape Models (ASMs) was proposed by Cootes et al. [34], (see also [168, 48]) whereby the template consisted of numerous shapes which were encoded into a shape model. Also encoded into this model were the principal modes of variability and this was used in the deformation process to minimise the template to object error. This is a very powerful idea in medical imaging and the extension of this method to include other parameters in the model, such as Active Appearance Models (AAMs) which integrates texture into the model [151, 150, 152, 78, 77, 17]. All model-based approaches are limited by the number and variation of the prior templates used in the model building process.

Active contours or Snakes which were first proposed by Kass et al. [68] are an extension of this top-down approach where a closed contour or surface is located in

the image and is then deformed until the final segmentation is achieved. Normally this deformation is constituted from two separate parts, the first is the intrinsic properties of the contour in order to maintain its shape through rigidity or elasticity and the second part of the deformation energy. This form of segmentation has been employed in medical image analysis, where the anatomical feature in question can be encapsulated within a closed contour [25, 26, 121, 4, 52, 67, 66]. Segmentation is then achieved by evolving this closed contour using intrinsic properties such as curvature and external properties obtained from the image. Combinations of snakes and statistical shape models have also been developed [60] whereby snake evolution is additionally guided using a predefined model of what the final shape should approximate. Non-parametric snakes were introduced in order to address some of the limitations of traditional snakes and have proved successful in medical image analysis [86, 110, 6, 2, 163]. These techniques are discussed in more detail in Chapter 5.

While these approaches have been shown to perform robust segmentation, even in noisy images, accuracy of the segmentation is bounded by the initial shape. This is particularly the case in medical imaging, where anatomical features present a significant variation between patients none more so than in the presence of disease.

There are many algorithms which try to employ a combination of bottom-up and top-down approaches to segmentation to capture the advantages from both approaches [16]. Prior knowledge about a particular segmentation task can be incorporated as low level information such as expected intensity values, gradient strength of orientation or incorporated at a higher level such as texture variation over an object and object shape.

With the increasing temporal resolution available in modern scanners, the tracking of clinical structures over time may hold particular clinical significance. This area has been investigated in the myocardium of the heart more than in any other biological structure (a excellent reviews of applying image processing techniques to left ventricle segmentation can be found in [156, 49, 44, 167]). Deformation tracking of the cardiac muscle over the temporal cycle has been investigated in many studies in order to measure the regional function of the left ventricle (LV) in an effort to isolate the location, severity and extent of *ischemic* myocardium [137]. Myocardium twist and torque can be measured with using *tagged*-MRI.

Single breath hold images means registration of the images is not as significant a factor as in multiple breath hold images. Registration methods [84] deal with the registration of cardiac images from multiple modalities. Registration techniques were first performed in brain images for the registration of higher resolution images acquired using MRI or CT to images of lower resolution such as Magneto-Encephalo-Graphy (MEG) or Electro-Encephalo-Graphy (EEG). Registration in cardiac images is more complicated due to the non-rigid and mixed motions of the cardiac muscle and thorax structures. Much attention is focussed on registration of the modalities MRI and PET [85, 139], MRI and SPECT [62] or CT and PET [179, 19].

2.5 Conclusions

This chapter introduces the key areas associated with this thesis. Firstly, an overview of the heart is given with particular emphasis on anatomical morphology and cardiac dynamics. This is followed by some of the most common CVDs and the clinically acquired measurements used in their diagnosis.

In the second part of this chapter, an overview of image acquisition is provided. MRI is the chosen modality for this study, based on the outlined advantages over other modalities. This is followed with a fundamental background in MRI physics and common protocols.

Finally, in order to extract the clinically relevant features from the data presented from the image acquisition, image processing is proposed and introduced. The remainder of the chapter is devoted to the exploration of how medical image analysis has evolved by classifying the approaches into two rudimentary methodologies (see review [44]).

Chapter 3

Advanced Data Filtering

Image smoothing is a procedure employed in image processing to reduce or simplify the data present in an image in order to make image understanding more attainable. In a practical sense, this can be achieved by the removal of noise or redundant signal intensities from the image in order to obtain a more appropriate model of the underlying structures within the image.

The motivation behind smoothing images is therefore two-fold, firstly it removes unwanted noise from the image to facilitate further processing and secondly to eliminate features irrelevant to the given problem to reduce the complexity for subsequent processing. Specifically in MRI, increased magnet strength may resolve somewhat the associated low SNR, but advances to 3T magnets are limited by the higher RF power disposition in the body [8]. Nayak et al. [105] showed in 2004 how 3T imaging improved SNR and CNR on cine sequences but note signal fall-off due to decreased RF penetration.

There are two main types of smoothing, linear and non-linear. Both of these types have been extensively studied in literature [116, 140, 159]. When filtering images, it is mostly an advantageous property of the smoothing filter to smooth areas of homogeneity while preserving areas of interest in the image such as edges. This is typically achieved by means of a convolution of a number of pixels or voxels with a smoothing kernel, this is also called Finite Impulse Response (FIR) filtering. Linear filters convolve an image patch with a smoothing kernel that is independent of the data in the image. Standard linear smoothing techniques based on local averaging or Gaussian weighted spatial operators reduce the level of noise but this is achieved at the expense of poor feature preservation.

Consequently, the filtered data appears blurry as step intensity discontinuities such as edges are attenuated. Non-linear filters smooth the image but try to maintain edges by smoothing less. Among these, the median filter is the simplest non-linear operator to remove impulse-like noise [142, 116]. More complex non-linear techniques include statistical approaches based on non-parametric estimation [140, 160]. However, while these methods do alleviate somewhat the shortcomings associated with linear techniques, they still perform only modestly when the data is affected by long tailed noise distributions. To complement these filtering approaches, a number of adaptive techniques have been proposed [140, 53, 33, 124, 28]. These methods try to achieve the best trade-offs between smoothing efficiency, feature preservation and the generation of artefacts. Koenderink [73] expressed the blurring operation of smoothing as heat conductance or diffusion. Diffusion-based filtering was originally developed by Perona and Malik [115] in order to implement an optimal feature preserving smoothing strategy. Many implementations follow their original approach where the main aim was to improve numerical stability [172]. This was advanced by Weikert [171] where he developed a new smoothing algorithm by permitting diffusion along the direction of edges. Gerig et al [53] extended this work to 3D and evaluated its usefulness when applied to medical 2D and 3D datasets.

In this chapter, a performance characterisation is evaluated on some advanced smoothing filters both in 2D and 3D. The performance of a filter is evaluated as a means of simplifying the image before segmentation. Therefore, advantageous characteristics are defined as their ability to flatten the signal intensity values within a structure while maintaining a strong separation of signal intensity values between structures. Firstly, five filters are introduced and assessed, two linear filters (Gaussian and Savitzky-Golay) and three non-linear filters (Diffusion-based, Adaptive and Anisotropic) are evaluated to detail the advantage of non-linear filters over linear filters. Finally, a rigorous performance characterisation is performed on the three non-linear filters using homogeneity within regions and edge strength as the indicators of performance.

3.1 Linear Methods

Traditional linear filters such as mean, average and Gaussian attempt to remove noise by replacing pixels by an average or weighted average of its spatial neighbours [116]. While this reduces the amount of noise present in the image, it also has the disadvantage of removing or blurring the edges. The Savitzky-Golay [127]

linear filter smooths the image but tries to preserve higher moments, like edges, in the image. It achieves this by selecting coefficients that are the least squares approximation of a higher degree polynomial.

Firstly, let us look at the basic linear causal smoothing filter given in equation 3.1. This is the 2D case where each pixel in the smoothed image g at position (x, y) is calculated to be the average or weighted average of the original image f 's spatial neighbours. The convolution matrix C is of size $N \times N$ where $N = 2n + 1$ and the sum of its elements is normalised to unity.

$$g_{x,y} = \sum_{j=-n}^n \sum_{i=-n}^n C_{i,j} f_{x+i,y+j} \quad (3.1)$$

This type of filtering introduces a blurring effect to the image which is undesirable for most image processing applications. The basic filter illustrated in Equation 3.1 is linear and is independent of the data being processed. Some common causal filters are mean, Gaussian and Savitzky-Golay.

3.1.1 Gaussian Filter

The Gaussian smoothing technique is very straightforward and is similar to the average filter. The Gaussian filter differs from the average filter in that it involves the convolution of the original image with a Gaussian mask where the standard deviation and the size of the smoothing kernel selects the scale and size of the blurring operation. The resulting image $S_{x,y}$ is defined as,

$$S_{x,y} = I_{x,y} \circ \text{Gauss}(x, y, \sigma) \quad (3.2)$$

where $I_{x,y}$ is the original image, $\text{Gauss}(x, y, \sigma)$ represents the Gaussian kernel with scale parameter σ and \circ implements the 2D convolution operation.

This form of smoothing has the advantage of giving more influence to the pixels or in close neighbourhood to the element being replaced, with exponentially less influence the further away the pixels are from the center of the kernel. In 2D the Gaussian mask is constructed using the following equation,

$$\text{Gauss}(x, y, \sigma) = \frac{1}{\sqrt{2\pi\sigma^2}} e^{-\frac{x^2+y^2}{2\sigma^2}} \quad (3.3)$$

where (x, y) are the 2D position of the element and σ represents the diffusive properties of the filter.

The standard blurring operation involving Gaussian filtering attempts to remove the noise from the image. From Equation 3.3 it is obvious that the smoothing becomes more pronounced for higher values of the scale parameter but at the same time we can notice a significant attenuation of the optical signal associated with image boundaries. This result is highly undesirable for many applications including image segmentation and edge tracking where a precise identification of the object boundary is required.

3.1.2 Savitzky-Golay Filter

The Savitzky-Golay [127] smoothing filter was introduced for smoothing one-dimensional tabulated data and for computing the numerical derivatives. The smoothed points are found by replacing each data point with the value of its fitted polynomial. These filters preserve edges far better than a moving average filter but this is achieved at the expense of not removing as much noise. The process of the Savitzky-Golay is to find the coefficients of the polynomial which are linear with respect to the data values. Therefore the problem is reduced to finding the coefficients for fictitious data and applying this linear filter over the complete data.

Savitzky-Golay can be used for smoothing image data by extending the filter to two dimensions with a two dimensional polynomial. The size of the smoothing window is given as $N \times N$ where N is an odd number, and the order of the polynomial to fit is k , where $N > k + 1$. The general smoothing causal filter equation is given as:

$$g_{x,y} = \sum_{j=-n}^n \sum_{i=-n}^n C_{i,j} f_{x+i,y+j} \quad (3.4)$$

n is equal to $\frac{N-1}{2}$. $g_{x,y}$ is the resulting smoothed data, C is the convolution matrix and f_{xy} is the original image data.

$$f(x_i, y_i) = a_{00} + a_{10}x_i + a_{01}y_i + a_{20}x_i^2 + a_{11}x_iy_i + a_{02}y_i^2 + \dots + a_{0k}y_i^k \quad (3.5)$$

We then want to fit a polynomial of type in equation (3.5) to the data. Solving the least squares we can find the polynomial coefficients. We start with the general equation:

$$A \cdot a = f$$

where a is the vector of polynomial coefficients

$$a = (a_{00} \quad a_{01} \quad a_{10} \quad \dots \quad a_{0k})^T$$

We can then compute the coefficient matrix as follows.

$$(A^T \cdot A) \cdot a = (A^T \cdot f)$$

$$a = (A^T \cdot A)^{-1} \cdot (A^T \cdot f)$$

Due to the least-squares fitting being linear to the values of the data, the coefficients can be computed independent of data. To achieve this we can replace f with a unit vector thus, the coefficient matrix becomes $C = (A^T A)^{-1} A^T$. C can then be reassembled back into a traditional looking filter of size $N \times N$.

The resulting coefficient matrix from a polynomial of order 3 and with a matrix window size of 5 (i.e. n_L and n_R is 2). In order to smooth the image the first coefficient is used, higher order coefficients are used to calculate derivatives. Here are the values for the first coefficients using a 5×5 windowing and orders of 3 and 4 respectfully:

$$C_{k=3} = \begin{pmatrix} -0.0742 & 0.01142 & 0.04001 & 0.01142 & -0.0742 \\ 0.01142 & 0.09714 & 0.12571 & 0.09714 & 0.01142 \\ 0.03999 & 0.12571 & 0.15428 & 0.12571 & 0.03999 \\ 0.01142 & 0.09714 & 0.12571 & 0.09714 & 0.01142 \\ -0.0742 & 0.01142 & 0.04001 & 0.01142 & -0.0742 \end{pmatrix} \quad (3.9)$$

$$C_{k=5} = \begin{pmatrix} 0.04163 & -0.0808 & 0.07836 & -0.0808 & 0.04163 \\ -0.0808 & -0.0196 & 0.20082 & -0.0196 & -0.0808 \\ 0.07836 & 0.20082 & 0.44163 & 0.20082 & 0.07836 \\ -0.0808 & -0.0196 & 0.20082 & -0.0196 & -0.0808 \\ 0.04163 & -0.0808 & 0.07836 & -0.0808 & 0.04163 \end{pmatrix} \quad (3.10)$$

The advantage Savitzky-Golay filters have over moving average and other FIR filters is its ability to preserve higher moments in the data and thus reduce smoothing on peak heights. It can be seen in Equations 3.9 and 3.10 that the higher the order of the polynomial the higher moments are preserved, this leads to less smoothing over data peaks and line widths. In more homogeneous areas the smoothing approaches an average filter over the smoothing kernel.

3.2 Non-Linear Filters

Nonlinear filters, the most common being the median filter, modifies the value of the pixel by some nonlinear function of the pixel value and its spatial neighbours. Nonlinear filters aim to maintain the edges but the filtering may result in a loss of resolution by suppressing fine details. Three non-linear filter are investigated. Firstly a non-linear diffusion based filter based on gradient information, secondly an adaptive filter which uses both gradient and variance within a neighbourhood as a measure of inhomogeneity and finally an anisotropic filter which changes the shape and strength of the smoothing kernel based on gradient strength and orientation.

A more useful way to think of smoothing is as a type of diffusion of intensities within an image, first expressed by Koenderink [73]. Diffusion occurs according to Fick's Law, given in equation 3.11[115], where ΔI is the Laplacian of the intensity value, $c(x, y, t) = \text{constant}$ represents the conductance coefficient and I_t is the recovered value at iteration t .

$$I_t = c\Delta I \quad (3.11)$$

When this equation is implemented it acts as a linear filter, similar to a Gaussian, but it becomes more effective when the non-linear terms are introduced into the diffusion equation. A review of nonlinear diffusion is compiled in [171].

3.2.1 Nonlinear Diffusion Filtering

To alleviate the problems associated with the standard Gaussian smoothing technique, Perona and Malik [115] proposed an elegant smoothing scheme based on

non-linear diffusion*. In their formulation the blurring would be performed within homogeneous image regions with no interaction between adjacent or neighbouring regions that share a common border. The non-linear diffusion procedure can be written in terms of the derivative of the flux function:

$$\phi(\nabla I) = \nabla I \cdot D(\|\nabla I\|) \quad (3.12)$$

where ϕ is the flux function, I is the image and D is the diffusion function. Equation 3.12 can be implemented in an iterative manner and the expression required to implement the non-linear diffusion is illustrated in Equation 3.13.

$$I_{x,y}^{t+1} = I_{x,y}^t + \lambda \sum_{R=1}^4 [D(\nabla_R I) \nabla_R I]^t \quad (3.13)$$

where I^t represents the image at iteration t , R defines the 4-connected neighbourhood, D is the diffusion function, ∇ is the gradient operator that has been implemented as the 4 connected nearest-neighbour differences and λ is a parameter that takes a values in the range $0 < \lambda < 0.25$.

$$\begin{aligned} \nabla_1 I_{x,y} &= I_{x-1,y} - I_{x,y} \\ \nabla_2 I_{x,y} &= I_{x+1,y} - I_{x,y} \\ \nabla_3 I_{x,y} &= I_{x,y-1} - I_{x,y} \\ \nabla_4 I_{x,y} &= I_{x,y+1} - I_{x,y} \end{aligned} \quad (3.14)$$

The diffusion function $D(x)$ should be bounded between 0 and 1 and should have the peak value when the input x is set to zero. This would translate with no smoothing around the region boundary where the gradient has high values. In practice, a large number of functions can be implemented to satisfy this requirement and in the implementation detailed in this thesis the exponential function proposed by Perona and Malik [115] is used.

$$D(\|\nabla I\|) = e^{-\left(\frac{\|\nabla I\|}{k}\right)^2} \quad (3.15)$$

*Perona and Malik discuss in their paper the topic of scale-space. This has not been investigated as it is beyond the scope of this thesis and a single scale space proved to be sufficient for the applications detailed in this document.

where k is the diffusion parameter. The parameter k selects the smoothness level and the smoothing effect is more noticeable for high values of k .

3.2.2 Adaptive Smoothing

The algorithm for adaptive smoothing implemented in this evaluation is adapted from Chen [28]. The technique measures two types of discontinuities in the image, local and spatial. Local variable discontinuities can detect local intensity changes but is susceptible to errors where there is a lot of noise, so in addition to the local discontinuities the contextual information is also utilised given the attributes of neighboring pixels. From both these measures a less ambiguous smoothing solution is found. In short, the local discontinuities indicate the detailed local structures while the contextual discontinuities show the important features.

Local Variable Discontinuities

In order to measure the local discontinuities, four detectors are set up as shown:

$$E_{H_{xy}} = |I_{x+1,y} - I_{x-1,y}|, \quad (3.16)$$

$$E_{V_{xy}} = |I_{x,y+1} - I_{x,y-1}|, \quad (3.17)$$

$$E_{D_{xy}} = |I_{x+1,y+1} - I_{x-1,y-1}|, \quad (3.18)$$

$$E_{C_{xy}} = |I_{x+1,y-1} - I_{x-1,y+1}|, \quad (3.19)$$

$I_{x,y}$ is the intensity of the pixel at the position (x,y) . We can then define a local discontinuity measure E_{xy} as:

$$E_{xy} = \frac{E_{H_{xy}} + E_{V_{xy}} + E_{D_{xy}} + E_{C_{xy}}}{4} \quad (3.20)$$

These pixel positions are illustrated below in Figure 3.1.

Contextual Discontinuities

In order to measure the contextual discontinuities, a spatial variance is employed. Firstly, a square kernel is set up around the pixel of interest, $N_{xy}(R)$. The mean




Figure 3.1: The four local discontinuity detectors.

intensity value of all the members of this kernel is calculated for each pixel as follows:

$$\mu_{xy}(R) = \frac{\sum_{(i,j) \in N_{xy}(R)} I_{i,j}}{|N_{xy}(R)|} \quad (3.21)$$

From the mean the spatial variance is then calculated to be:

$$\sigma_{xy}^2(R) = \frac{\sum_{(i,j) \in N_{xy}(R)} (I_{i,j} - \mu_{xy}(R))^2}{|N_{xy}(R)|} \quad (3.22)$$

This value of sigma is then normalised to $\tilde{\sigma}_{xy}^2$ between the minimum and maximum variance in the entire image. A transformation is then added into $\tilde{\sigma}_{xy}^2$ to alleviate the influence of noise and trivial features. It is given a threshold value of $\theta_\sigma = (0 \leq \theta_\sigma \leq 1)$ to limit the degree of contextual discontinuities.

Overall Adaptive Algorithm

Finally, the actual smoothing algorithm runs through the entire image updating each pixels intensity value I_{xy}^t , where t is the iteration value.

$$I_{xy}^{t+1} = I_{xy}^t + \eta_{xy} \frac{\sum_{(i,j) \in N_{xy}(1)/\{(x,y)\}} \eta_{ij} \gamma_{ij}^t (I_{i,j}^t - I_{xy}^t)}{\sum_{(i,j) \in N_{xy}(1)/\{(x,y)\}} \eta_{ij} \gamma_{ij}^t} \quad (3.23)$$

where,

$$\eta_{ij} = \exp(-\alpha \Phi(\tilde{\sigma}_{xy}^2(R), \theta_\sigma)), \quad (3.24)$$

$$\gamma_{ij}^t = \exp(-E_{ij}^t/S) \quad (3.25)$$

The variables S and α determine to what extent the local and contextual discontinuities should be preserved during smoothing. If there are many contextual discontinuities in the image then the value of η_{ij} will have a large influence on the updated intensity value. On the other hand, if there are a lot of local

discontinuities then both γ_{ij} and η_{ij} will have the overriding effect, as η_{ij} is used for gain control of the adaptation.

3.2.3 Anisotropic Gaussian Smoothing

An anisotropic filter based on the familiar Gaussian model is implemented in order to provide edge enhancing, directional smoothing. The goal is to develop a versatile smoothing filter based on a straightforward and highly adaptable form. The approach reduces to a convolution with a scaled and shaped Gaussian mask, where the determination of the mask weights becomes the key step governing the performance of the filter. By calculating the local grayscale gradient vector and favouring smoothing along the edge over smoothing across it can achieve an effective boundary preserving filtering approach, where regions are homogenised while edges are retained.

The weight $wt(\vec{pq}, \nabla u)$ at each location in the mask is a function of the local gradient vector at the centre of the mask and the distance of the current neighbour from that centre. There are a large number of possibilities for the formulation of the mask weight calculation, based on the desired form for the non-linear and anisotropic components of the filter. The weight for some neighbour q is calculated as a function of the gradient of point p , at the mask origin, and the distance from the origin to the neighbour q . The relationship used in our approach is given in Equation 3.26, where \vec{pq} is the vector from the mask centre point p to some neighbour q , ∇u is the gradient vector at p , λ is the scale parameter, controlling smoothing strength, and μ is the shape parameter, controlling anisotropy. When μ equals zero the anisotropic term $(\frac{\vec{pq} \cdot \nabla u}{\lambda})^2(2\mu + \mu^2)$ has a negligible effect and the filter reduces to the non-linear, isotropic form, where smoothing decreases close to strong edges but is applied equally in all directions, at any given location in the image.

$$wt(\vec{pq}, \nabla u) = e^{-((\frac{\|\vec{pq}\| \|\nabla u\|}{\lambda})^2 + (\frac{\vec{pq} \cdot \nabla u}{\lambda})^2(2\mu + \mu^2))} \quad (3.26)$$

The images in Figures 3.2 and 3.3 illustrate the operation of the anisotropic filter. As the smoothing strength and the number of iterations is increased more noise and small features are eliminated, but even in extreme cases the most important edges in the image are well preserved in both location and strength.

3.3 Experiments and Results

The aim of each filter evaluated in the first study is to measure the linear and non-linear filters ability to smooth areas of homogeneity while preserving areas of interest such as edges. Smoothing of homogeneous areas is measured using the standard deviation while the preservation of edges is measured using the strength and spread of the edge in the filtered images. To show the advantage of using non-linear filters, both the linear are tested on two 2D images, see figures 3.2(a) and 3.3(a). The first image of a laboratory having a high SNR (signal-noise-ratio) and high CNR (contrast-to-noise-ratio) with a high density of edges. The second medical image has a much lower SNR and CNR. Parameters were chosen to give the optimal results on visual inspection. Visual results are presented in Figures 3.2 and 3.3.

To be statistically relevant [42] the standard deviation should be calculated over a large region but on the other hand the results will be affected by small non-uniformities such as intensity gradients or structural image variations [53]. This requirement is quite difficult to be accomplished if we want to develop an automatic performance characterisation scheme where user intervention is not required. One solution has been advanced by Canny [20] when he decided to select the threshold parameters for an edge detector based on analysis of the cumulative histogram of the gradients. However due to the nature of MR datasets this criteria to identify the gradients generated by noise proved to be inefficient. Thus, in this implementation an alternative strategy based on observation has been developed. In this sense, we computed the standard deviation for all data points in the original dataset in a 7×7 neighbourhood. These values were sorted with respect to their magnitude and from these values the 25% of the highest values were eliminated, as they are likely to belong to edges and 25% of the lowest values are also eliminated as they are calculated from areas that have no significant texture (such as image regions defined by air). This strategy was applied to select the seed points that belong to image regions defined by a low SNR. Then, the standard deviation for each of the filtered datasets is measured at the same data point locations (also in a 7×7 neighbourhood). The results are presented in Table 3.3.

For the laboratory image, Adaptive smoothing gives the best results followed by the two other non-linear filters. Both linear Savitzky-Golay and Gaussian filters have the highest deviation after smoothing. In the medical image there

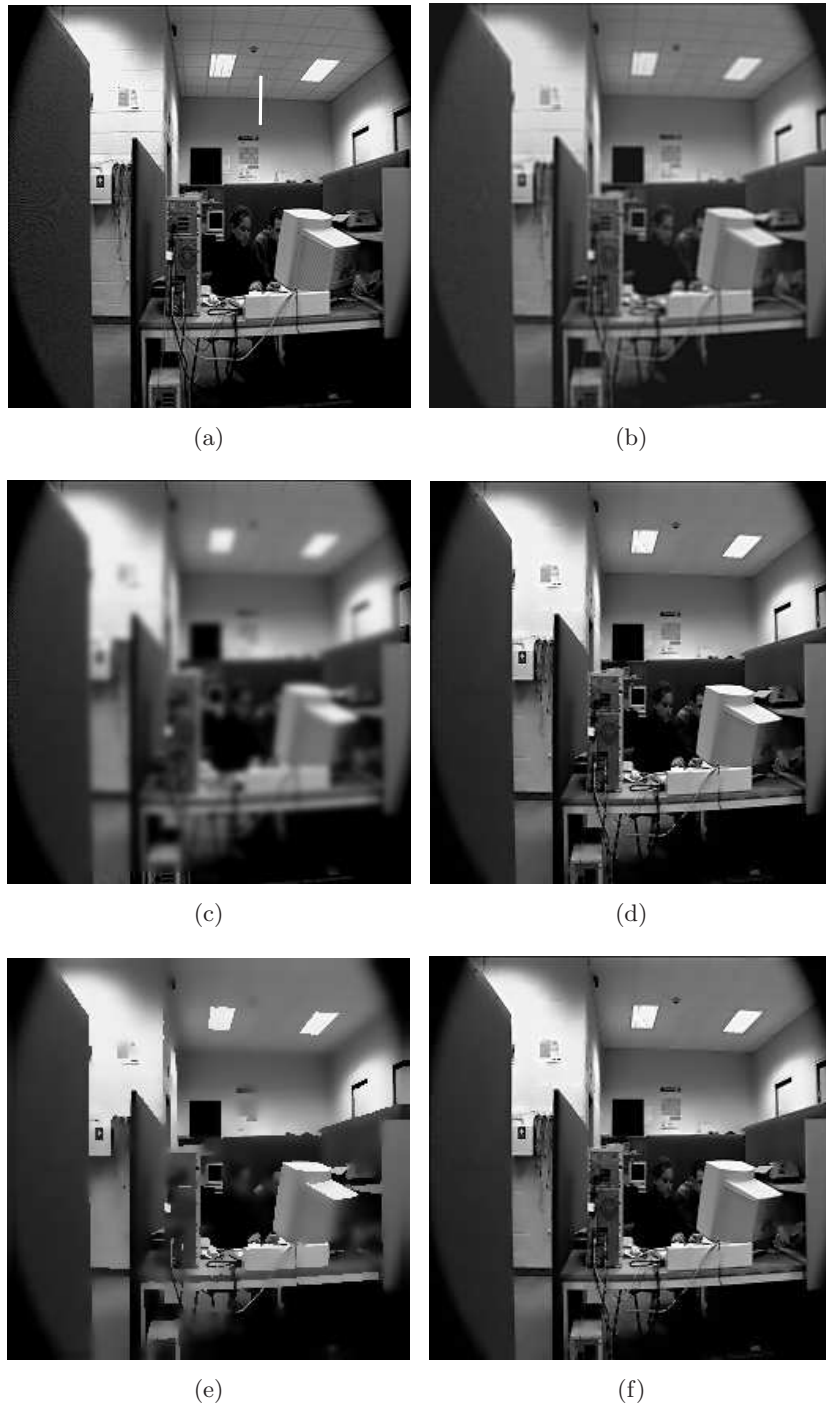


Figure 3.2: Smoothing results. Original image (a) is shown after the application of the Savitzky-Golay (b), Gaussian (c), Diffusion (d), Adaptive (e) and Anisotropic (f) filters.




Figure 3.3: Smoothing results. Original image (a) is shown after the application of the Savitzky-Golay (b), Gaussian (c), Diffusion (d), Adaptive (e) and Anisotropic (f) filters.

are more significant differences with the anisotropic and adaptive giving the best results while the gaussian suffers in the low SNR image.

The strength, shift and spread of the edge is evaluated on each of the images. Histogram plots across two edges, see the white lines across edge features in figures 3.2 and 3.3. In Figure 3.4, the histogram plots show both the image pixels and the gradient for the lab image and medical image. For the lab image the results are similar for all filters with more significant differences between filters in the medical image. Two measurements are taken from these histograms which indicate edge strength and spread where edge spread is taken as the Full Width Half Maximum (FWHM) of the gradient plot. These results are compiled in Table 3.2.

	<i>Laboratory Image</i>	<i>MR Image</i>
Original	57.4	277.65
Savitzky-Golay	40.804	61.232
Gaussian	40.966	102.08
Diffusion	27.658	69.633
Adaptive	24.241	42.99
Anisotropic	31.905	35.05

Table 3.1: The RMS of the standard deviation of the homogeneous areas for the original and each filtered image

	<i>Laboratory Image</i>		<i>MR Image</i>	
	Edge height	Edge width	Edge height	Edge width
Original	31	2.26	219	2.04
Savitzky-Golay	23	2.5	158	2.48
Gaussian	15	4.4	196	2.16
Diffusion	25	2.17	214	2.00
Adaptive	26	2.13	211	2.00
Anisotropic	30	2.17	219	1.99

Table 3.2: Shows the edge strength and edge spread on the gradient image after each filtering. While Savitzky-Golay and Gaussian filters spread the edge, the other three maintain and even enhance the edge characteristics.

From all the experiments detailed above, it is clear that the non-linear filters outperform the linear filters using the criteria specified at the beginning of the test. The next step is to perform a more rigorous characterisation of the non-linear filter described above in medical images. The following experiments have been performed in 3D using the extension of the 2D to 3D of the non-linear algorithms described previously.

3.3.1 Performance Characterisation of Non-Linear Filters

The first set of experiments is conducted on a synthetic dataset that is defined by a homogeneous cubic object with a known grayscale value surrounded by background pixels. To test smoothing algorithms on this artificial dataset is advantageous as the ground truth data is known and the smoothing results are easy to evaluate. The efficiency of the algorithms when the artificial dataset was corrupted with various types of 3D image noise is tested, including Gaussian, Poisson




Figure 3.4: Pixel intensities and gradient intensities along white lines from images figure 3.2 and figure 3.3. (i) and (iii) show the pixel intensities and (ii) and (iv) show the gradient values from the lab image and the medical image respectively. (a) is the original image, (b) image after Savitzky-Golay, (c) Gaussian, (d) Adaptive, (e) Nonlinear Diffusion and (f) Anisotropic and Gaussian.

Noise Type	S. Dev noise	S. Dev F1	S. Dev F2	S. Dev F3	Grayscale Expected	Grayscale F1	Grayscale F2	Grayscale F3
G-15	13.72	1.91	1.62	2.06	127	127	128	128
G-30	31.93	7.64	3.03	5.57	127	128	129	133
P-15	13.02	1.07	0.76	1.74	127	139	138	138
P-30	26.97	9.6	7.62	3.69	127	141	141	142
W-15	4.63	1.5	0.21	0.69	127	126	127	127
W-30	8.56	1.71	0.6	1.14	127	125	126	127

Table 3.3: Performance characterisation results when the algorithms have been applied to an artificially created dataset. F1, F2, F3 denote the standard diffusion, adaptive smoothing and anisotropic diffusion respectively.

and additive uniformly distributed white noise [42]. Similar to the previous experiments, as quantitative values the local uniformity sampled by the $7 \times 7 \times 7$ standard deviation is evaluated at the location situated at the centre of the cube and the alteration of the grayscale value at the same position when compared with the expected known value. Some experimental results are depicted in Table 3.3.

In Table 3.3 the symbols G-15 and G-30 indicate that the synthetic dataset has been corrupted with Gaussian noise (standard deviation 15 and 30 grayscale values). Similarly P-15 and P-30 denote the fact that the test dataset has been corrupted with Poisson noise (distribution 15 and 30 grayscale values) and W-15 and W-30 indicate that the dataset has been corrupted with uniformly distributed white noise (mean deviation 15 and 30 grayscale values). In order to evaluate globally the noise removal efficiency on real datasets we need to define quantitative measures that indicate the overall performance of the smoothing algorithms that are evaluated. In this regard, we propose to evaluate jointly two quantitative measurements: the smoothness factor that assesses the global uniformity and the edge preservation factor that indicates to what extent the strong edge features are retained and enhanced. To this end, the standard deviation as a measure to evaluate the image local homogeneity was employed. As before, the standard deviation is measured in a $7 \times 7 \times 7$ window over the entire original image. These values were sorted with respect to their magnitude and 25% of the highest values were eliminated as belonging to edges in the image and 25% of the lower values as having no significant texture. The standard deviation for each of the filtered images is then taken at the same pixel locations. To evaluate a quantitative estimation we calculate the RMS value of the standard deviations from the original

and smoothed datasets resulting after the application of the smoothing strategies described in previous sections (for details refer to Table 3.3).

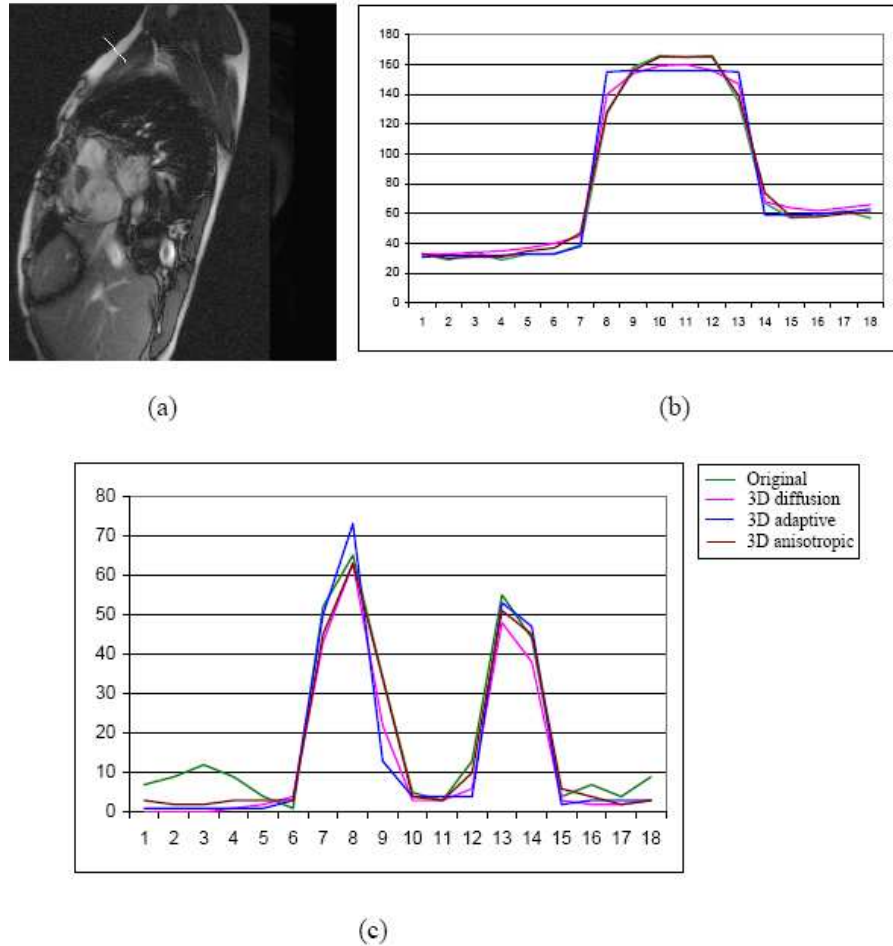


Figure 3.5: (a) Slice of the heart MRI dataset. Pixel (b) and (c) gradient intensities are plotted for the highlighted edge illustrated in image (a).

The edge strength is evaluated by plotting the intensity and gradient data at selected locations where edges are located, before and after the application of the smoothing operations. Some graphical results are depicted in Figures 3.5 to 3.8. The experimental data illustrated in Figures 3.5 to 3.8 indicate that the 3D adaptive smoothing and 3D anisotropic smoothing algorithms perform better than the standard diffusion. The 3D adaptive smoothing algorithm returned better results than the 3D anisotropic when applied to heart, brain and whole body datasets. The 3D anisotropic algorithm performed better when applied to

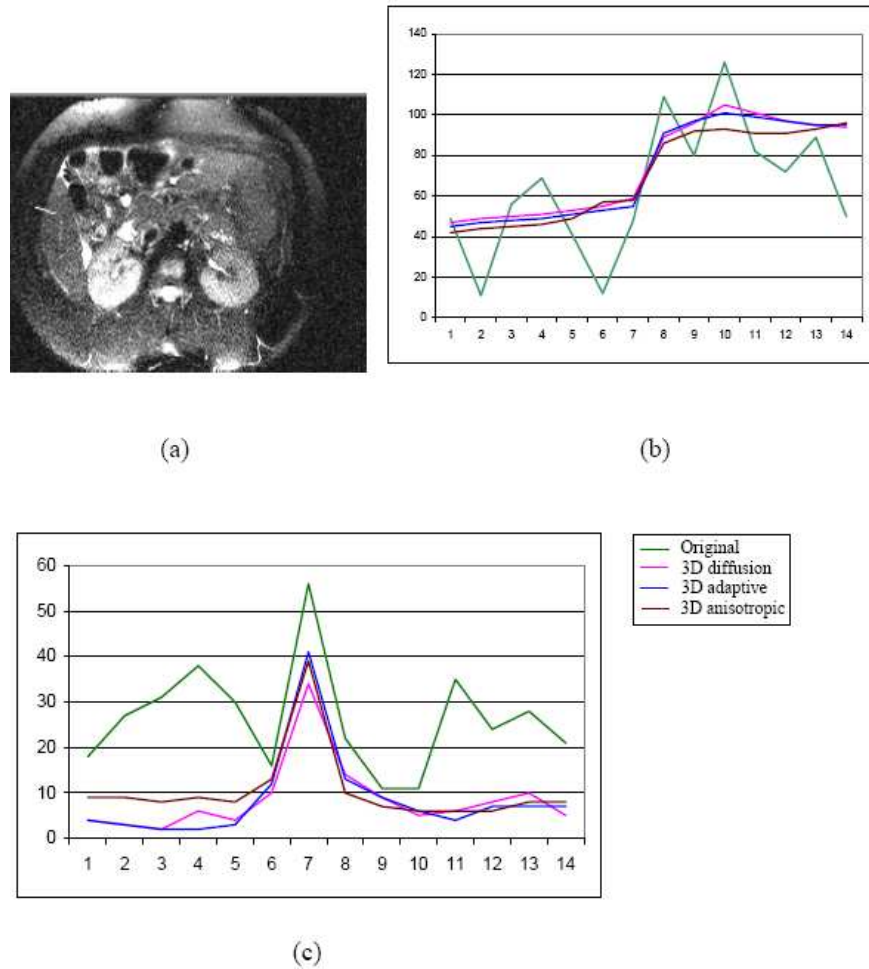


Figure 3.6: (a) Slice of the MRCP dataset. Pixel (b) and gradient intensities (c) are plotted for the highlighted edge illustrated in image (a).

Magnetic Resonance Cholangiopancreatography (MRCP) dataset.

The graphs illustrated in Figures 3.5 and 3.8 demonstrate the edge enhancement around image data defined by step-like edges. It can be noticed that the edge localisation is significantly improved. The effect of edge strengthening is even more pronounced for weaker edges in an MRI brain sequence (see Figure 3.7) or in image areas affected by a high level of noise, as is the case of the MRCP dataset illustrated in Figure 3.6. The performance of the non-linear smoothing algorithms described in this section is remarkable in discriminating a true edge from image noise (see Figure 3.6c). Also notice the improved performance of the adaptive 3D smoothing algorithm as compared with the performance of the

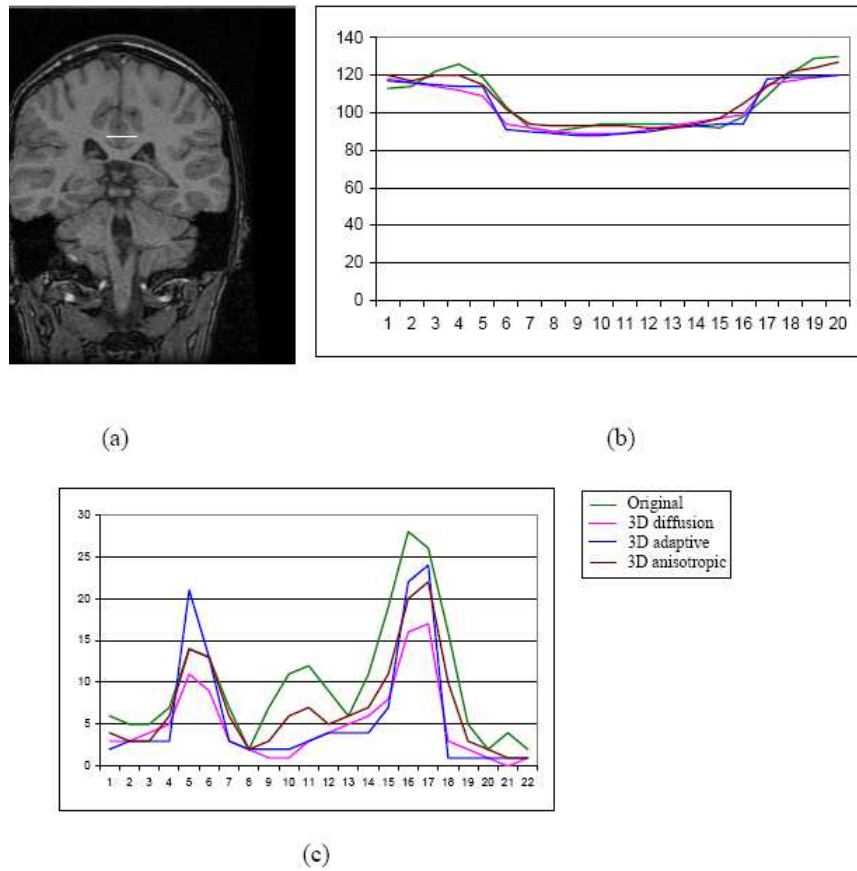


Figure 3.7: (a) Slice of the brain MRI dataset. Pixel (b) and gradient intensities (c) are plotted for the highlighted edge in image (a).

standard diffusion and the 3D anisotropic diffusion algorithms. In order to emphasise the effectiveness of the smoothing strategies described in this chapter the segmentation resulting after the application of a 3D clustering algorithm [42] to the original and smoothed data is presented. Samples of the segmentation results are depicted in Figures 3.9 to 3.12.

3.4 Conclusions

In this chapter, the performance in smoothing for a number of linear and non-linear filters was evaluated. In the first part, experiments were performed in order to show the advantage of non-linear filters over linear filters. In the second part, three diffusion-based smoothing schemes were implemented and their appli-

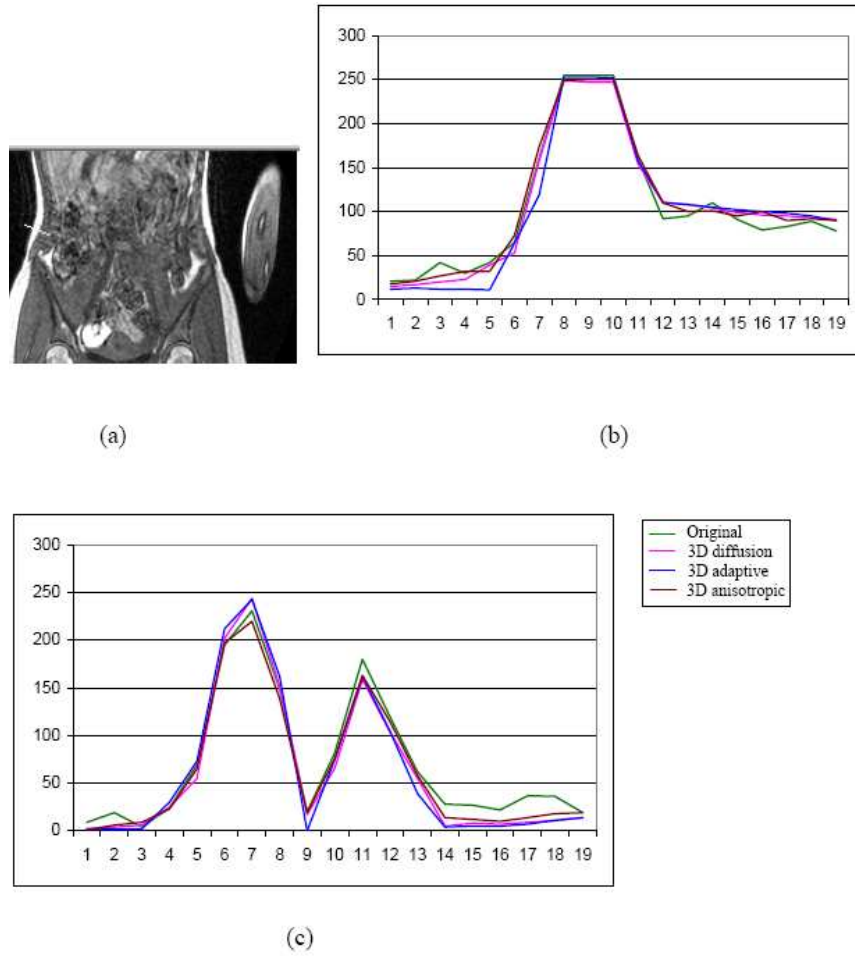


Figure 3.8: (a) Slice of the whole body MRI dataset. Pixel (b) and gradient intensities (c) are plotted for the highlighted edge illustrated in image (a).

	Heart	Brain	Whole body	MRCP
Original data	4.95	9.21	20.46	18.8
3D diffusion	1.88	6.28	14.47	10.96
3D adaptive	1.73	6.16	14.05	10.83
3D Anisotropic	2.08	6.48	16	9.77

Table 3.4: The RMS of the standard deviations of the homogeneous areas for the original and filtered MRI datasets used in our experiments.




Figure 3.9: 3D data clustering results - heart dataset. (First row) Original dataset (slice 9), and corresponding image resulted after clustering. (Second row) 3D diffusion smoothed data (slice 9) and corresponding image resulted after clustering. (Third row) 3D adaptive smoothed data (slice 9) and corresponding image resulted after clustering. (Forth row) 3D anisotropic smoothed data (slice 9) and corresponding image resulting after clustering.




Figure 3.10: 3D data clustering results - brain dataset. (First row) Original dataset (slice 4), and corresponding image resulted after clustering. (Second row) 3D diffusion smoothed data (slice 4) and corresponding image resulted after clustering. (Third row) 3D adaptive smoothed data (slice 4) and corresponding image resulted after clustering. (Forth row) 3D anisotropic smoothed data (slice 4) and corresponding image resulted after clustering.




Figure 3.11: 3D data clustering results - MRCP dataset. (First row) Original dataset (slice 10), and corresponding image resulted after clustering. (Second row) 3D diffusion smoothed data (slice 10) and corresponding image resulted after clustering. (Third row) 3D adaptive smoothed data (slice 10) and corresponding image resulted after clustering. (Forth row) 3D anisotropic smoothed data (slice 10) and corresponding image resulted after clustering.




Figure 3.12: 3D data clustering results whole body dataset. (First row) Original dataset (slice 6), and corresponding image resulted after clustering. (Second row) 3D diffusion smoothed data (slice 6) and corresponding image resulted after clustering. (Third row) 3D adaptive smoothed data (slice 6) and corresponding image resulted after clustering. (Forth row) 3D anisotropic smoothed data (slice 6) and corresponding image resulted after clustering.

cation to medical 3D data was described. The main interest was focused on MRI acquisition modalities as MRI datasets are characteristically defined by a low signal to noise ratio (SNR). Hence, the aim was to demonstrate that far superior results are achieved if the MRI data is initially filtered in order to reduce the level of image noise and improve the SNR. In this regard, a detailed performance characterisation was performed for each smoothing operators evaluated on both synthetic and real data (including heart, brain, whole body and MRCP image sequences). We conclude that in our experiments the non-linear diffusion-based smoothing technique provided the most efficient approach to noise reduction, and more importantly this advantage is not achieved at the expense of feature preservation in our experimentation. Computational time was higher for the non-linear iterative approaches, but as computational expense is not a limiting factor in our application this parameter was not included in the characterisation. The experimental data presented and discussed in this chapter highlights the ability of the diffusion-based smoothing schemes to distinguish the high gradient image features from the MRI image acquisition noise.

Publications associated with this chapter

Journal Publication

Ovidiu Ghita, Kevin Robinson, Michael Lynch and Paul F. Whelan (2005), **MRI diffusion-based filtering: A note on performance characterisation**, Computerized Medical Imaging and Graphics.

Conference Publication

Michael Lynch, Kevin Robinson, Ovidiu Ghita and Paul F. Whelan (2004), **A Performance Characterisation in Advanced Data Smoothing Techniques**, IMVIP 2004 Irish Machine Vision and Image Processing Conference, September 2004, Trinity College, Dublin, Ireland.

Chapter 4

Statistical Partitioning of Data for LV Localisation and Extraction

The advanced filtering techniques employed in the last chapter alleviates much of the work needed in the classification process. Preprocessing the data has removed much of the inherent noise associated with MRI therefore the process of segmenting the data into the relevant anatomical features can be achieved using data partitioning technique. To this end, it is the aim of this chapter to use cluster analysis to successfully segment the left ventricle blood pool. The left ventricle blood pool can then be automatically located using shape characteristics before a more heuristic method is employed to segment the outer boundary of the left ventricle muscle.

Data clustering remains a very active topic in image processing. The application of robust techniques for object identification in images are extensive, none more so than in the rapidly advancing field of medical imaging [30, 117]. Region-based methods [117] are used to segment the image, this is generally achieved without using *a priori* information. The most basic form of region-based segmentation is thresholding. Thresholding techniques create a binary image of pixels above and below a user defined threshold value. Thresholding does not take into account the structure or connectivity of the points that it segments and the threshold value is seldom automatically determined. Segmentation results can sometimes be filled with holes or ragged edges, which in a crude way can be

eliminated with a combination of morphological operators [63, 141]. In medical imaging, thresholding is not widely used without some advanced preprocessing steps due to its sensitivity to noise. More complex statistical methods, like clustering, join pixels of similar intensities to create a segmentation of structures in the image.

All statistical based classification methods [61, 40, 64, 42, 65, 114, 113] aim to optimise the results based on an initialisation. This initialisation is commonly chosen randomly, and as a consequence results are not reproducible, do not take advantage of inherent patterns in the data or may be initialized on outliers. Methods for automatic initialisation of clusters have been proposed in literature [3, 97, 71]. Al-Daoud and Roberts [3] proposed two methods, the first picks points randomly in evenly spaced cells across the entire histogram of the data and reduces the number until the required seeds are found. The second method tries to optimize the sum of squares of the distances from the cluster centers. Mitra *et al.* [97] describe a *rough-set* initialisation provided by graph-search methods. Khan and Ahmad [71] assume a normal distribution over the data attributes and divide the normal distribution curve into equal percentile cells. The seeds are chosen as the midpoints of the interval of each of these partitions. In Appendix A, a novel method developed by the VSG for the initialisation of cluster centers is given where the cluster centers are automatically detected using histogram analysis and applied to medical images.

In order to extract clinical measurements from the smoothed data, a novel method is proposed whereby data is first clustered in order to segment highly differentiated features, i.e. the blood and myocardium. A localisation of the left ventricle is detailed. Using this preliminary step, a new method for the extraction of the epi-cardium boundary is developed which is based on a knowledge driven search of gradient information. Where appropriate gradient information is lacking prior knowledge is used to augment the segmentation solution.

4.1 Data Clustering

Clustering is a well documented image segmentation technique which classifies pixels into groups or clusters using a distance criteria to join data values to each cluster. The most basic form of clustering is Hierarchical clustering, off which there are two types – agglomerative and divisive. Agglomerative clustering in-

volves creating n clusters from the data $X_i = \{x_1, x_2, x_3 \dots x_n\}$ where n is the amount of elements and $X \in \mathfrak{R}^m$. The process then iteratively combines this clusters in a branching formation until there is just one cluster containing all n elements. The clusters are joined using a distance criteria, which can be measured in different ways; single-linking, complete linking, unweighted average pair and weighted average pair. Divisive clustering works in the opposite way by creating one cluster with n elements and then dividing the clusters until n clusters remain. Successful analysis of both these joining methods comes from knowing at which iteration in the process will return the optimal amount of clusters to create a meaningful segmentation.

The k -means, or c -means, clustering method is a well established as a partitioning method [61, 136]. Delibasis et al. [38] proved how the k -means algorithm performed more robustly in a comparative study with an adaptive region growing, fuzzy C -means clustering and Kohonen self-organising maps for the segmentation of the left ventricle blood pool from cardiac MRI images. This comparison was performed on both normal and abnormal cases and results were evaluated against a manual delineation of the left ventricle cavity.

Unlike the Hierarchical methods, the k -means algorithm requires a user defined set of clusters. The process then exchanges the elements between clusters with two aims; to minimise the variation within each cluster and to maximise the variation between clusters. The algorithm has four main steps to find the image clusters, this is also illustrated in figure 4.1. The process terminates when no more elements are exchanged between clusters and it can be shown that the method is always convergent. The process is the minimization of the following equation.

$$E = \min \sum_j (\mathbf{x}_j - \mathbf{m}_{c(\mathbf{x}_j)})^2 \quad (4.1)$$

where j is the number of data points index and $\mathbf{m}_{c(\mathbf{x}_j)}$ is the class centroid closest to the data point.

In this thesis, the smoothed MRI images are then clustered using an improved version of the k -means algorithm proposed by Duda and Hart [42, 61]. An adaptive form of clustering is developed whereby the initial user defined number of clusters is iteratively reduced until a more appropriate number of clusters is achieved. This is based on thresholding the inter and intra cluster variability. Firstly, the image is clustered using an initial guess of 15-20 independent cluster centres which is sufficient to capture all the relevant features. The pixels are

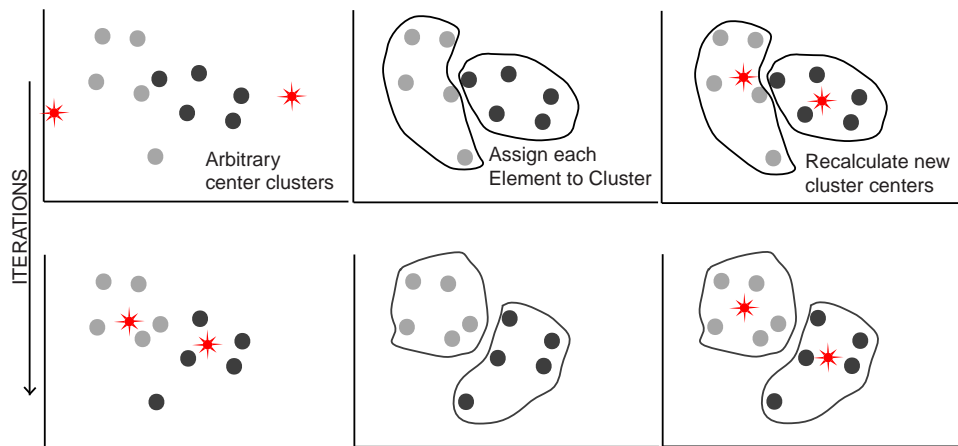


Figure 4.1: Two iterations in the k -means clustering technique on 2D data. The objects change with each iteration to join the cluster whose centre is closest.

clustered together using the following strategy. This algorithm has four steps to find the image clusters:

- (i) Initialise the position of the means $m_1 \rightarrow m_k$.
- (ii) Assign each of the k -items to the cluster whose mean is nearest.
- (iii) Recalculate the mean for the cluster gaining the new item and the mean for the cluster losing the same item. Recalculation is made using the intra cluster variance.
- (iv) Loop through steps (ii) and (iii) until there are no movements of items.

Initialisation of cluster centres can have a significant effect on the results of the classifier, therefore random initialisation is avoided. Alternatively, seeds may be placed at specific regions or equidistantly in the image space or in grayscale space. A better solution to maximise the use of input data in initialising the cluster centres is choosing them based on histogram analysis of the data. This approach is detailed in Appendix A.

In the second phase of the algorithm, each of the k clusters are sorted and compared. The number of clusters is then optimised by merging clusters with similar attributes. This is repeated until there are no more clusters to be merged. The stopping criterion for this joining process is defined using a threshold on the intra cluster variability and is chosen experimentally. Given the high differentiation in intensity signal between the blood pool and the myocardium, experimental

results display robust separation of the blood pool from the myocardium. As can be seen in Figure 4.2 the generality of the method as it is applied to two separate protocols, spin-echo and gradient echo with satisfactory results.




Figure 4.2: Shows four images; a gradient-echo images before (a) and after clustering (b), and a spin-echo image before (c) and after clustering (d).

4.1.1 Automatic Detection of *lv* cavity

The image has now been segmented into separate clustered regions. The next step is to automatically detect which of these clusters represents the *lv* cavity on the first slice. The *lv* cavity is located using shape descriptors only and not using the gray scale values which allows for application of this method in various MRI imaging protocols. The images are short axis, therefore we assume that the *lv* cavity approximates a circular shape and that the *lv* feature is present in successive slices. Approximation to a circle is calculated as the error of the

fitted areas of a least squares approximation to a circle. The approximation is obtained by minimising the error of the areas of the fitted circle and the areas of the associated circles at each data point (see mathematical background Appendix C.1). It is also assumed that the lv is not located on the periphery of the image.

The volume of the left ventricle is then extracted using two criteria:

- (i) Overlapping area of the regions contained in successive slices.
- (ii) Gray scale value of the regions under investigation.

The regions cannot be connected using just gray scale values alone due to the variation in the intensity values through the volume caused, to some extent, by coil intensity falloff. The lv regions are then connected in 3D and the volumes are then rendered for visualisation purposes (see Figure 4.8). The ejection fraction is calculated using the systolic and diastolic volumes. The ejection fraction is defined as “the proportion, or fraction, of blood pumped out of your heart with each beat” [104] and can be calculated using the equation:

$$EF = \frac{V_{endo}(t_D) - V_{endo}(t_S)}{V_{endo}(t_D)} \quad (4.2)$$

where V_{endo} is the volume of the inner walls of the heart, $V_{endo}(t_D) = \max_t[V_{endo}(t)]$ is the end-diastolic volume and $V_{endo}(t_S) = \min_t[V_{endo}(t)]$ is the end-systolic volume.

The corresponding region is found by maximising the result of a cost function where the overlapping and the mean gray-scale value of the areas under investigation are used as parameters.

This works well on basal and mid-cavity slices, the blood pool is large and relatively homogeneous. The apical region is more challenging due to the increase in *trabeculae carneae* and papillary muscles, the low volumes of blood present, partial voluming along the z axis and blurring due to movement of the diaphragm. The extension of this segmentation algorithm to 3D is appropriate as the higher level of knowledge leads to improved segmentation results plus it eliminates the need to match relevant clusters through the volume using overlapping criterion.

4.2 Extension to 3D

In order to improve the robustness of the segmentation technique it is favorable to extend the clustering to the third dimension. The extension means that the blood pool is clustered as a whole and therefore it is more robust in areas where

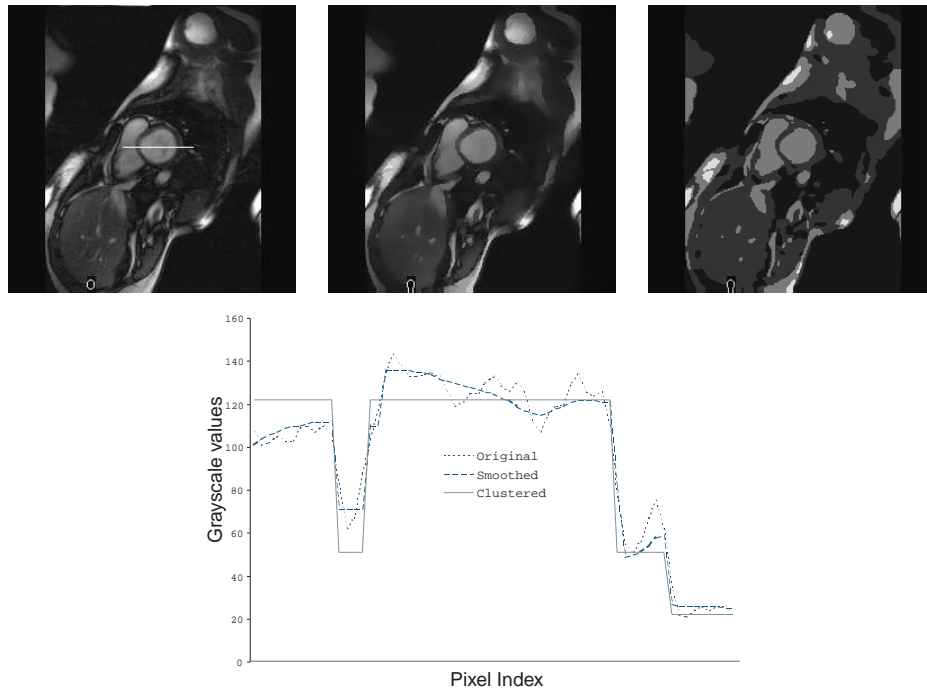


Figure 4.3: The top three images from left to right show the original short axis image, after smoothing and after clustering. The graph plots the intensity values for the white line running through the original image.

artifacts such as the papillary muscles are present. This is particularly the case around the apical regions of the left ventricle cavity.

The end-systole and end-diastole volumes are smoothed in 3D, as in the previous chapter. Once smoothed they are then clustered using the 3D k -means technique using the volume data. The left-ventricle can be manually picked or automatically using the volumetric shape properties of the cavity, as developed in the following section.

4.2.1 Automatic Detection of lv cavity using 3D information

In order to locate the left-ventricle in the image a number of shape descriptors were used. The images are short axis so therefore we use the anatomical knowledge that the *lv* cavity approximates a circular shape and that the *lv* feature is continuous in successive slices. In the 2D scenario, approximation to a circle is calculated as the error between the shape and the least squares approximation to it's circle. Also, a smooth interpolation of the curves is achieved using a spline

fitting.

In the 3D case, the left ventricle cavity is located using its shape description. In this case it is known that on the short axis the left ventricle approximates an ellipsoid in shape, although it is flat at one end, perpendicular to its major axis. The approximation to an ellipsoid parameters (radii and centres) is calculated using the first three principal axes of the PCA of the boundary data points. The error is then calculated between the shape and the fitted ellipsoid using the summation of the normalised point deviations with respect to the calculated ellipsoid radii (see mathematical background in Appendix C.2).

4.3 Segmentation of epi-cardium border

Once the left ventricle blood pool has been successfully segmented, the outside of the myocardium or epi-cardium boundary presents a more challenging problem. Parts of the outer wall of the left ventricle displays low gradient information and low differentiation between neighbouring tissues, as in Figure 4.4.




Figure 4.4: Illustrating the low grayscale differentiation between the outer wall of the myocardium and other organs in the body, before (top row) and after (bottom row) data partitioning.

This is especially true in areas close to the lungs and liver. Therefore clustering techniques are not applicable because the differentiation between tissues is so low and edge detection will only have limited success when used without

supervision or a more involved approach. Noble et al. [107] attempt to change the coordinate system to polar coordinates followed by a constrained snake segmentation to capture the epi-cardium boundary. In order to address these issues, a novel heuristic approach is developed which uses all the available information in a supervised way and where information does not exist or is not found, the segmentation is augmented using prior information of the epi-cardium boundary shape.

Calculation of the wall-thickness and wall-thickening is dependent on the accurate segmentation of the epi-cardium boundary. The main problem associated with the segmentation is the low contrast-to-signal ratio along the epi-cardial boundary in particular on the inferior and inferolateral side where the muscle becomes indistinguishable from the lung. To this end two novel approaches are explained and have been evaluated, both a robust approximation for the epi-cardium thickness to determine strong features of the epi-cardium present in the image. Where strong information is lacking, the algorithms aim to approximate the epi-cardium boundary using in the first case an arc, centered at the center of gravity of the blood pool and connecting two known segments of the epi-cardium boundary. In the second approach, where no information is present, the algorithm uses information obtained from a probabilistic model consisting of manually segmented images to complete the epi-cardium boundary.

4.3.1 First Approach: Robust-Arc epi-cardium segmentation

The robust arc approximation technique works on the 2D slice taken from the previously segmented blood pool volume. Firstly the centre of gravity of the left ventricle blood pool is located. The least squares approximation for the radius of the endo-cardium border is calculated on each slice. The original image is re-clustered again around a smaller region of interest with a smaller predefined number of clusters in order to find the right ventricle blood pool. The right ventricle blood pool is found to be the largest cluster close to the left ventricle cavity with similar intensity attributes to the left ventricle blood pool. The interventricular septum between the two ventricles is measured and this measurement gives an approximate thickness for the myocardium around the left ventricle.

A Canny edge-detection [20] is performed on the original image slice. A 1D radial search is carried out from the centre of gravity on the gradient image and




Figure 4.5: From left to right: Original unseen image, calculated edges, robust segments of epi-cardium and the complete segmentation using arcs to complete the epi-cardial boundary

image edge points are connected together into edge segments using an Euclidean distance criteria. Spurious segments are eliminated by length, orientation away from the endo-cardium border and using the approximation for the myocardium from the septum.

In between these segments are parts of the epi-cardium border that do not have any gradient. Therefore there is no other information in the image to help find the correct path between these segments. In this case the end points of the robust segments are joined with an arc, pivoted around the center of gravity of the endo-cardium. Results can be seen in figure 4.5.

The procedure for segmenting the epi-cardium can be followed in the diagram illustrated in Figure 4.6, Stage II. The position of the *lv* cavity is already known in each slice as explained in the previous section. In order to determine the epi-cardium border a region of interest is defined around the *lv* cavity. Two copies of this region of interest are taken. The first image *Image1* is used to find a value for the approximate radius of the myocardium and the second image *Image2* is used to find real borders around the myocardium. The two are combined to find the true value of the epi-cardium around the *lv*.

Image1 is again clustered using a predefined low number of clusters around the region of interest. A low number of clusters is chosen because of the scarcity of important features around the *lv* cavity. Anatomically, the closest blood pocket to the *lv* cavity is the right ventricle cavity, it is also assumed that the thickness of the myocardium will not change drastically over the entire circumference. The thickness of the wall, or septum, between the two blood pockets can give a reliable estimate for the thickness of the rest of the myocardium.

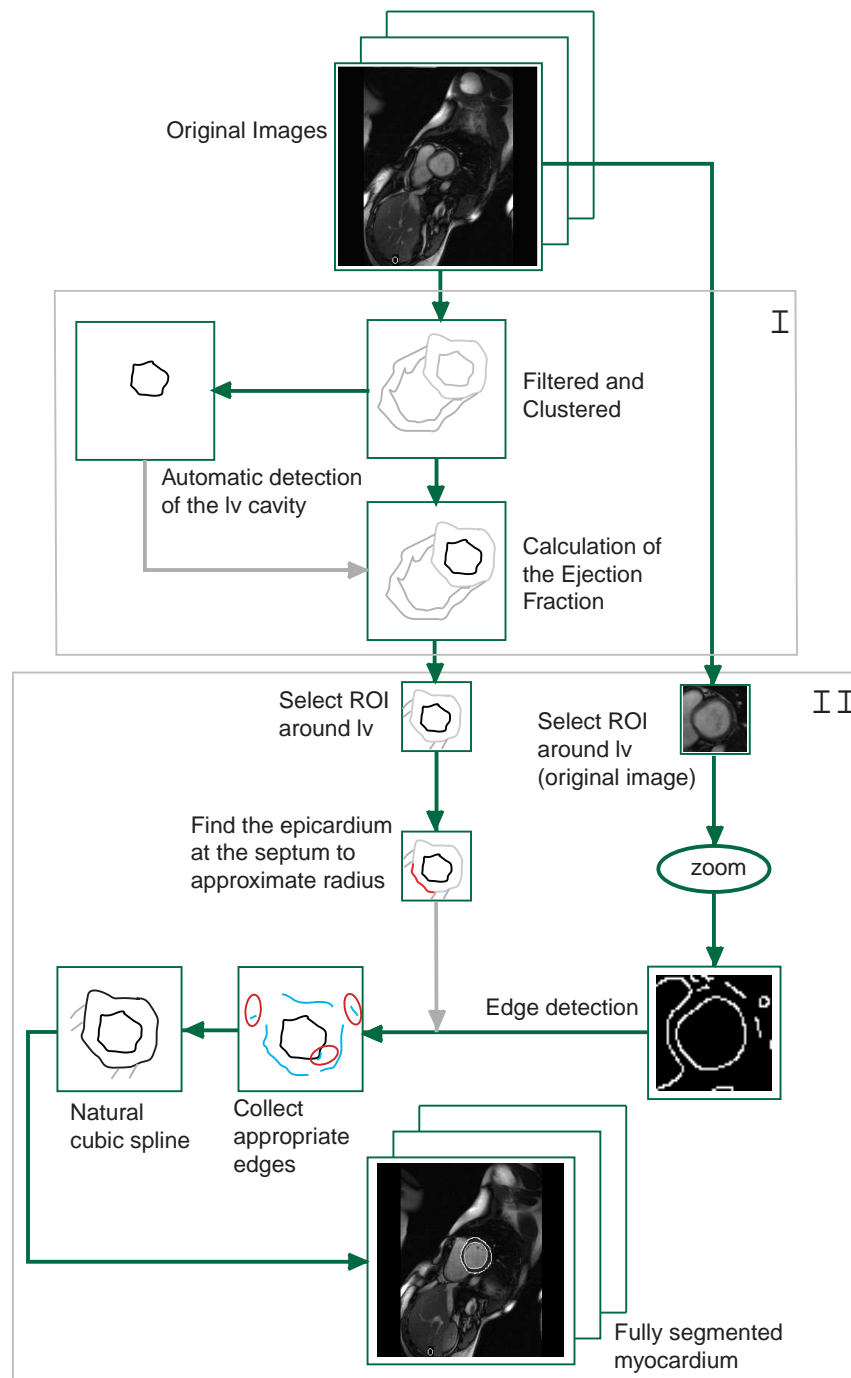


Figure 4.6: A schematic representation of the two phases involved in the segmentation of the endo- and epi- cardium border. *Stage I* shows the preprocessing and segmentation processes, the automatic detection of the *lv* cavity and the connection of the cavity through the volume. *Stage II* shows the method for segmenting the epi-cardium border in each image

Image2 is zoomed using an area averaging technique around the area of interest. The image is then segmented using a thresholded edge-based algorithm [20]. The zooming operation is applied to increase the edge separation. The largest connected segments within certain bounds of the estimated thickness found from *Image1* are taken as potential border segments. There is an angular restraint placed on the transition of these segments around the epi-cardium to eliminate stepping into the endo-cardium border or stepping out to other organs.

A closed natural cubic spline is fitted around the points on the epi-cardium [144, 12], for the formulation see section C.3. The spline is used to close the epi-cardium contour by connecting all the points on the curve in a smooth way. Splines are piece-wise polynomials with the pieces smoothly joined together. The joining points of the polynomial pieces are called control points which do not have to be evenly spaced. Each segment of a spline is a polynomial of degree n , for this implementation n was chosen to be $n = 3$. More details on the mathematical formulation of the natural cubic spline can be found in Appendix C.3.

4.3.2 Second Approach: Model assisted Epi-cardium segmentation

In order to incorporate more realistic approximations for missing data, a new method is developed which uses a probabilistic model of previously segmented heart images. Once each slice is taken from the volume the centre of gravity of the left ventricle blood pool is located. The least squares approximation for the radius of the endo-cardium border is calculated. By re-clustering the original image again around a smaller region of interest with a predefined number of clusters in order to find the right ventricle blood pool. The right ventricle blood pool is found to be the largest cluster close to the left ventricle cavity with similar intensity attributes to the left ventricle blood pool. The myocardium wall (septum) between the two ventricles is measured and this measurement gives an approximate thickness for the myocardium around the left ventricle.

An edge-detection is performed on the original image slice. A 1D radial search is carried out from the centre of gravity on the gradient image and image edge points are connected together into edge segments using a Euclidean distance criterion. Spurious segments are eliminated by length, by orientation away from the endo-cardium border and using the approximation for the myocardium from the septum.

A database of contour points is created based on manual segmentations of the endo and epi cardium boundaries. This database contained 180 2D contours with the associated radii calculated using the least squares approximation based on minimising the error of the areas (detailed in AppendixC.1). Where epi-cardial boundary is not defined by the edge information, the boundary is then completed from a generic database of hand-segmented shapes. The database is searched using the ratio of epi-cardium and endo-cardium radii. The searching uses the two end-points of the robustly located segment from the gradient image. Prior to searching, each contour is scaled with respect to radii parameters extracted from the model. Each scaled contour in the database is searched to minimise the Euclidean distance from these endpoints to their nearest corresponding points on the contour. The contour that minimises this error is chosen. The appropriate section is extracted from the contour and joined to the edge defined boundary using a natural closed spline (see figure 4.7).




Figure 4.7: From left to right: Original unseen image, calculated edges, robust segments of epi-cardium and the complete segmentation using an *a priori* knowledge database

In figure 4.7(b) the segment points obtained from gradient image figure 4.7(a) are illustrated. In between these segments are parts of the epi-cardium border that do not have any gradient. Therefore there is no other information in the image to help find the correct path between these segments. In this case *a priori* knowledge about the shape of the epi-cardium border, obtained from previously hand-segmented can be used to join the segments. In this way we introduce a form of supervision, and by inferring previously drawn contours we hope to maintain continuity of the shape. Because the contours contain the original segments while the manually drawn contours are only inferred where there is no information to be rendered from the image, it is believed that this approach generates more appropriate results than the previous technique, when the model provides a

good approximation to the object being segmented. When complete models are inferred onto the image there is a danger that details may be lost.

4.4 Results

In order to assess the performance of the automatic segmentation, results were compared against those obtained by manually segmenting volume image sequences for the endo- and epi-cardium borders. The manual segmentation was assisted by an experienced cardiologist*. Each volume includes 5-12 images containing the *lv*, transversing the length of the cavity and includes the papillary muscles. The automatic segmentation results can be seen in figure 4.12. The method shows good visual results for bright blood images 4.12(a)-(f) and dark blood images 4.12(g)-(i). The errors are calculated on volumes, endo and epi contours areas, myocardium thickness and finally point correspondence.

Table 4.1 shows the signed average and root mean square error of the ejection fraction from eight volumes from the sequence. The ejection fractions were worked out using pairs of volumes, not necessarily the end-systole and end-diastole and compared with the ejection fraction calculated from the manually segmented volumes. We can see in Table 4.1 low errors between the manual and automatic results.

The errors for the manually segmented endo-cardium area and the automatically traced area are given in Table 4.1. The signed average and root mean square errors are shown. Errors around the apex have a significant effect because the errors are described in proportion to the overall area calculated from the manual segmentation. Linear regression analysis was also performed in Figure 4.9(a) and high correlation value of $r = 0.98$ is obtained. Reproducibility is assessed using the Bland-Altman plot, Figure 4.9(b) [15]. From the Bland-Altman plot we can see that there is a tendency to underestimate the areas of the endo-cardium boundary, this is due to the inclusion of some endo-endocardium fat in the manual segmentation and perhaps due slightly to partial voluming effects. Also evident from the graphs is the accurate performance of this procedure in both systolic and diastolic phases, represented by the lack of skew in the plots as the areas increase. Note that the graphs are relatively zoomed to show the detailed distribution and the plots are graphed in units of mm^2 .

*The validation was performed by Dr. John Murray, Cardiologist, Mater Misericordiae Hospital, Dublin, Ireland.

The epi-cardium area was assessed using the linear regression and Bland-Altman plots. It shows a slightly lower percentage error for both the average signed and the rms errors. This can be attributed to the increased overall area of the manually traced contours. Linear analysis, Figure 4.10(a), gives a value of $r = 0.94$ while Figure 4.10(b) gives a similar regression value of $r = 0.95$ which is slightly lower than that produced for the endo-cardium. This lower correlation is a result of low contrast on the lateral side of the heart making the segmentation of the epi-cardium border difficult. In this case our algorithm connects two end-points of robust segments, how these segments are connected can incorporate *a priori* information [83]. Manual segmentation is also problematic in areas of low gradient and is dependent on the users own interpretation of ‘what looks appropriate’. Reproducibility was again assessed with the Bland-Altman plot, figure 4.10(b). Again, both methods produced similar results, both bands of two times the standard deviation are similar and not as tight as those achieved in the blood pool segmentation. There is not a significant difference between both methods as robust gradient information is used when available and both approaches are only applied in areas that are lacking gradient information. Both plots show no bias from the zero error or skew in the data. Although, the second approach which uses a prior database of contours does produce a larger number of outliers for the smaller apical regions where the outer wall may be undefined and approximation is difficult. Using this approach, more appropriate segmentations are achieved when compared to full manual segmentations. However, these methods still have the limitation that they are only working on slice data and not incorporating volume or temporal information.

Table 4.1: Mean Percentage Errors \pm 1SD for manual versus automatic

	Average Signed Error	RMS Error
Ejection Fraction	1.593 ± 0.82	3.176
Endocardium Areas	-3.623 ± 5.14	4.765
Epicardium Areas	-0.556 ± 4.29	3.75

Table 4.2 gives the Euclidean point to curve error in mm’s for all images through a heart sequence. It gives the minimum and maximum distance between the manual and automatic segmentation contours. The average distance, standard deviation (SD) and root-mean-square (RMS) are also given. The results for the epi-cardium boundary point to curve errors are shown in Table 4.3 and illustrated in figure 4.11.




Figure 4.8: The rendered images of (a) the end-diastole *lv* cavity, (b) the end-systole *lv* cavity, (c) and (d) the diastolic myocardium. These volumes are constructed from the true segmentation of the images excluding fat and papillary muscles.




Figure 4.9: Figures (a)-(b) shows scatterline plot of manual segmentation against the automatic segmentation and shows Bland-Altman plot for the left ventricle blood pool areas.

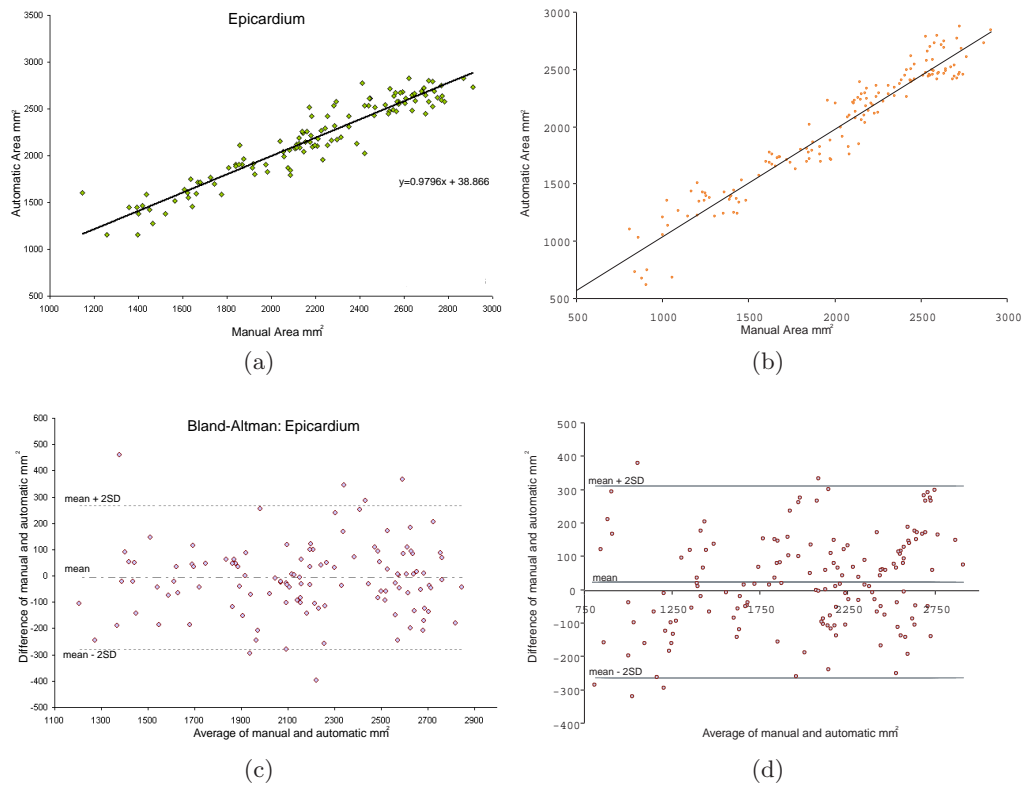


Figure 4.10: (a) illustrates the results using the Robust arc technique and (b) shows the results using the Prior model technique.

Table 4.2: Point to curve Errors between manual and computer segmentation for clustering technique for the endo-cardium boundary segmentation(mm)

Method	Endo-cardium		
	Average (mm)	Std. Dev. (mm)	RMS (mm)
3D k-means Clustering	0.69	0.88	1.12

Table 4.3: Point to curve Errors between manual and automatic segmentation for the epi-cardium boundary(mm) segmentation

Method	Epi-cardium		
	Average (mm)	SD (mm)	RMS (mm)
Robust Arc	1.31	1.86	2.14
Prior Model	1.26	1.27	1.94

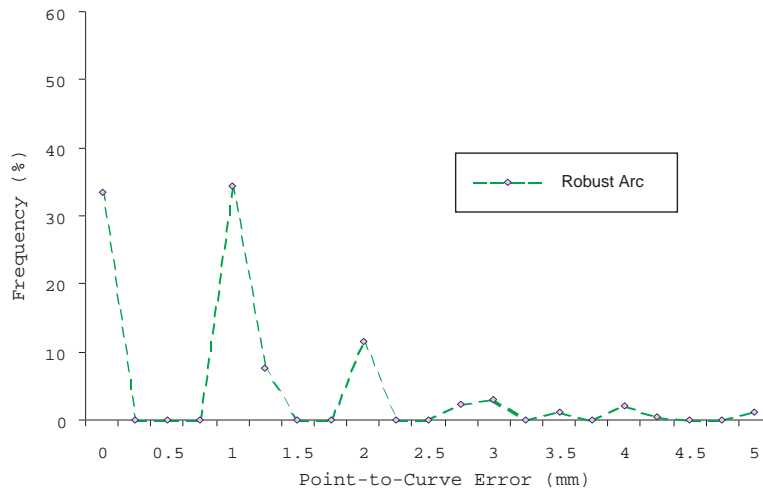


Figure 4.11: Plot shows the error frequency using a point to curve error metric.

4.5 Conclusion

A fully automatic detection and segmentation of the left ventricle myocardium has been detailed in this chapter. Edge preserving data filtering is performed and followed by an unsupervised clustering to successfully segment the left ventricle cavity from short axis MR images of the heart. Once the cavity volume is extracted the ejection fraction can be calculated.

In the second part of the chapter the epi-cardium border is successfully segmented using an edge-based technique. The thickness of the wall is approximated by measuring the thickness of the interventricular septum. The interventricular septum is an anatomically sound feature of the heart and because it is surrounded by blood on both sides it can be robustly segmented. This measurement is then used as an initial estimate for the thickness of the complete wall. A gradient image of the area around the lv is computed and the use of the approximate wall thickness, gradient points potentially belonging to the epi-cardium border are selected. If there are no viable gradients found on the epi-cardium border then the outer wall is estimated using the approximation found using the interventricular septum.

Statistical partitioning of the images allows the extraction of the lv blood pool without the use of prior constraints on shape. Abnormalities in the image data can indicate disease. Model based approaches approximate to the closest plausible instance shape from the training set Point Distribution Model (PDM), but this may not be sufficiently accurate. Also model based approaches that incorporate texture are limited in their use when the texture in the object images varies significantly from those contained in the model training set. The method proposed in this chapter presents a robust, fully automated method to identify the endo-cardium and epi-cardium borders that does not rely on *a priori* knowledge nor does it use shape constraints to find the left ventricle cavity.

Left ventricle segmentation is primarily motivated by the need to clinically diagnose a feature of the heart with potential problems. Models that approximate left ventricular boundaries try to fit variations of boundaries that have already been segmented. The left ventricle is anatomically variant, the scanners are inconsistent and the variations of pathologies found in patients is vast. To build a model to accommodate such diversity would be an immense task. Our algorithm makes no approximations based on observed data but instead produces a true

evaluation of the heart structure by segmenting the true borders in the image. It should be remembered that the aim is not to segment hearts that are part of a model but to assist the cardiologist in the prognosis by delineating the true anatomical features present in the image. Therefore, it is the aim of this thesis to approach the problem from a bottom-up strategy in as far as possible. Image segmentation can be augmented using prior information in the case where no image information is present and also to supervise the segmentation from spilling into other anatomical structures.

Evaluating the endo-cardium and epi-cardium borders using this approach could provide a more appropriate technique for flagging problems like wall thinning and low ejection fraction.

However, while this method provides good results in well imaged data and has been successful in segmenting the left ventricle blood pool in 2D and 3D data and, it is the aim of this thesis to increase the robustness of the segmentation approach by incorporating the entire data presented from the patient scan and remove the heuristic approach by creating a well defined mathematical framework. The aim of this approach is to create a more involved technique which segments both myocardium boundaries as opposed to two separate steps and also facilitate the incorporation of temporal information. The investigation of evolving surfaces, their parameterisation, termination and incorporating advanced information is performed in the next chapter.

Publications associated with this chapter

Journal Publication

Michael Lynch, Ovidiu Ghita and Paul F. Whelan (2005), **Automatic Segmentation of the Left Ventricle Cavity and Myocardium in MRI Data**, *Computers in Biology and Medicine* 36(4):pp389-407.

Conference Publications

Michael Lynch, Ovidiu Ghita and Paul F. Whelan (2004), **Extraction of Epicardial Contours from Unseen Images Using a Shape Database**, *IEEE NSS-MIC 2004 Medical Imaging Conference*, October 16-22, 2004, Rome, Italy.




Figure 4.12: The left ventricle contours obtained using our automatic segmentation method in short axis cardiac MR images. Figures (a)-(f) show images taken at both the end-diastolic phase and end-systolic phase of a gradient-echo sequence. Figures (g)-(i) show images from a spin-echo study.

Michael Lynch, Ovidiu Ghita and Paul F. Whelan (2004), **Comparison of 2D and 3D clustering on Short Axis Magnetic Resonance Images of the left ventricle**, CARS 2004 Computer Assisted Radiology and Surgery, June 23 - 26, 2004 Chicago, USA.

Michael Lynch, Ovidiu Ghita and Paul F. Whelan. **Calculation of Ejection Fraction (EF) from MR Cardio-Images**. Paper in the Irish Machine Vision and Image Processing Conference 2003, Coleraine, Northern Ireland.

Chapter 5

Boundary-Based and Model Driven Segmentation in Multidimensional Data

In this chapter, a review of current boundary based and model based segmentation schemes is detailed and their application to medical image analysis. Particular emphasis is placed on cardiac left ventricle segmentation in MRI [156, 117, 30, 48]. In Section 5.6 the level set framework is described and novel approaches to segmentation with level sets is introduced, in particular the extension to 4D data analysis.

Many boundary based segmentation (also called Active Contours) methods for object segmentation have been developed for use in medical image object extraction. Generally, the aim of boundary based segmentation methods is to deform a closed curve using both intrinsic properties of the curve and image based information to capture the target object [158]. This form of segmentation has many advantages over statistical intensity based partitioning algorithms as boundary shape is one of the key factors in the evolution of the contours. One of the most popular forms of boundary based segmentations are snakes, which were first introduced by Kass *et al.* [68]. From their introduction snakes have received a large amount of interest from the research community and much work has been done on derivations of the original snake. Further work in controlling the snakes propagation was achieved using parametrically deformable models and also by the introduction of *a priori* model driven segmentation with Active Shape and Active

Appearance Models. A Eulerian formulation of the active contour is introduced by means of a level set algorithm. The advantages of this formulation include a more robust mathematical theory, capability to follow topological changes in shape, and other computational advantages like curvature measurement. Work on the level set formulation for segmentation will constitute the main part of this chapter. A number of key issues in the level set are then addressed which include the choice of stopping term, the introduction of *a priori* information, the coupling of two level sets for the extraction of both the epi- and endo-cardium boundary and finally the introduction of an Expectation-Maximisation extension of the level-set framework to fully segment data in 4D (3D + t).

5.1 Active-Contours

Firstly, a 2D simple contour can be defined as $\mathbf{v}(s) = [\mathbf{x}(s) \quad \mathbf{y}(s)]^T$ for $s \in [0, 1]$. The main idea is to deform this contour smoothly to extract certain features in an image [92]. In a segmentation scheme this usually applies to extracting an area of homogeneous signal intensity, this may represent an object in a medical image such as the liver organ or a pool of blood. Therefore the deformation of the curve should flow globally outwards or inwards but should be inhibited from crossing areas of high frequency in the image data.

In this sense, the energy used to deform the boundary is a combination of a smoothing term, relating to the intrinsic properties of the boundary curve $\mathbf{v}(s)$, and an image dependent term, obtained directly from the underlying image data.

$$E = E_{int} + E_{ext} \quad (5.1)$$

5.1.1 Internal Energy

The internal energy aims to smooth the deforming contour, as in most cases in the segmentation of natural objects the boundary is defined as relatively smooth. To this end, the internal energy uses a combination of first derivative to determine tension or elasticity of the local contour and second order differential in order to calculate the bending of the local contour. The resulting values present high energy levels in irregular contours with shape corners and low energy in contours with a smooth transition between evenly separated points. If the contour was to

deform with the deformation energy obtained solely from the internal energies, the contour would achieve a perfect circle.

$$E_{int} = \int_0^1 (\alpha |\frac{\partial \mathbf{v}}{\partial s}|^2 + \beta |\frac{\partial^2 \mathbf{v}}{\partial s^2}|^2) ds \quad (5.2)$$

In Equation 5.2, α and β are weighting factors. In practice β may be set to zero, both to reduce the complexity of the derivation of the curve evolution to a geometric space and also because curve smoothing can be obtained with the first regularisation term alone [22].

5.1.2 External Energy

The external energy uses the image data to stop the deformation at the desired position. Stopping criterion may involve image data intensity, free end of boundary termination, corners or in this case high frequency or high gradient data. The resulting energy should return low values on high gradient points and high values on low gradient points.

$$E_{ext} = -\lambda \int_0^1 |\nabla I(\mathbf{v}(s))| ds \quad (5.3)$$

In Equation 5.3, λ is a user defined weighting function and $I(\mathbf{v}(s))$ is the image intensity. To suppress the influence of noise on the deformation the data may be smoothed using a Gaussian filter, thus becoming $\nabla[G_\sigma * I(\mathbf{v}(s))]$ where σ parameter controls the variance of the Gaussian.

Therefore, the active contour can be described as an energy minimisation problem that seeks to deform a closed contour to rest on high image gradients while maintaining a smooth transition between points. An inflation term may be appended to the energy terms, this can take the value of ± 1 along the normal direction to the curve [31]. This inflation term grows or shrinks the contour from its initial position to aid with the initialisation.

The main advantages of active contours are their extension to 3D, (where they are referred to as active surfaces), their ability to capture a closed structure and the users ability to select different features as stopping terms. In medical imaging, many of the natural anatomical structures are represented by closed smooth active surfaces. It is for this reason that many researchers have investi-




Figure 5.1: Curve propagating with a force 'F', in the normal direction of the local boundary.

gated methods and extensions to employ active contours for the segmentation of medical images, and this will be investigated further in the following section.

There are however disadvantages associated with the snake method. One of the key limitations of the snake algorithm is the problem of initialisation. The active contours aim to deform until the stopping energy overpowers the influence of the intrinsic energies and in some cases the inflation term. Also, the selection of the parameter space and sampling rule also has a large influence on the final segmentation result.

5.1.3 Application of Active Contours

Active contours have been used extensively for segmentation in the field of medical imaging, a full review of deformable models in medical imaging can be found in [92]. McInerney and Terzopoulos [90] apply a 3D dynamic balloon model using triangle-based finite elements to segment the left ventricle from cardiac CT data.

Much attention has been given to improving the snake computational framework, for instance Amini et al. [4] suggests using dynamic programming in order to minimise the energy function. This approach is claimed to produce the optimal local contour by searching all the possible solutions. Geiger [52] describes a non-iterative dynamically programmed method to extract the optimal contour,

providing the initial contour is a close approximation. To speed up this algorithm and to improve robustness, multi-scale images are used. Ronfard [121] introduces region-based energy by building statistically models of the background and object data. These model distributions are used in place of edge information to determine the contour termination.

Chakraborty et al. [25, 26] also introduce region based information into the evolution of the active contour. Molloy and Whelan [98] introduce active meshes that initialise a deformable triangular mesh on corner data in the images and used the forces between nodes to deform the mesh in order to track the data through an image sequence. Sermesant et al. [131] introduce a novel function which performs an affine transformation of a deformable model in order to optimally fit to image data. Jolly et al. [67, 66] employ active contours, semi-automatically initialise on each slice in the short axis view and then propagate through the cardiac cycle. Santarelli et al. [126] introduce a Gradient-Vector-Flow (GVD) snake which proceeds a diffusion filter to segment the inner and outer boundaries of the left ventricle of the heart.

Reuckert et al. [122] applies active contours for localisation of the aorta. Neubauer[1] presented a myocardium segmentation following a manually placed 'skeleton' inside the myocardium. The results are then propagated through all other slices in the volume. Spreeuwers[145] attempts to address the issue of robustness in the presence of erroneous local minima by applying a coupled active contour for the extraction of both the epi- and endo-cardium boundaries simultaneously. Mikic [93] uses optical flow estimates to guide the evolution of the active contour in echocardiographic sequences.

5.2 Parametrically Deformable Models

Staib and Duncan [146, 39] introduce a deformable model based on parametric contours. These models are commonly used when some prior information about the geometric shape of the final contour can be determined. This geometric shape can then be encoded using a small number of parameters. The model is then deformed, maintaining the overall consistency of the global model, by optimising the parameters on the image data. Most commonly, the global model can be defined by a set of analytical curves. Staib and Duncan [146] use elliptic Fourier decomposition for objects with shape irregularities, where a Fourier shape model

is used that represents a closed boundary as a sum of trigonometric function of various frequencies. They then perform an iterative energy minimisation to fit the model to the image data. This method may provide robust localisation of features, where the feature matches the template, however, this technique does not provide an appropriate basis for capturing shape variability and the generic models built using a priori knowledge need to be good approximations of the final segmentation result.

5.2.1 Application of Parametrically Deformable Models to Medical Imaging

Parametrically deformable models have been applied in the segmentation of cardiac MRI images. For instance, Staib and Duncan [147] propose a geometric surface matching. The model uses a Fourier parameterisation which decomposes the surface into a weighted sum of sinusoidal basis functions. In [147], four basis functions are used; tori, open surfaces, closed surfaces and tubes. The surface finding is formulated as an optimisation problem which attracts the surface to strong image gradients in the vicinity of the model.

The main disadvantage of parametrically deformable models is the effects of the choice of coefficients as this determines the complexity of the curve. Placing limits on each coefficient constrains the shape to an extent but not in a systematic way. While these models work well for localisation of the left ventricle, a derived model could not completely hold all the variation of the true left ventricle. These models have problems to define the complex shape of the left ventricle which varies from patient-to-patient and between healthy and dysfunctional ventricles.

5.3 Active Shape Models

Cootes et al. [36] propose a method to fit a shape model to image data. Recently, this has been applied to a wide range of image classification and segmentation problems. This method has had reasonable success in the case where:

- the target object has a well defined shape,
- can be represented with a set of examples and
- can be approximately located within the image.

There are limitations associated with this method where:

- the objects present a high variation in shape,
- grayscale or when the position/size/orientation of the target are not approximately known and
- the models themselves can contain human bias in annotation or error in point correspondence of landmarks.

Firstly, manual delineation of the object in a sample set of images is performed. From the manually drawn contours, positional landmarks are extracted in the form $\mathbf{x} = [x_1, x_2 \dots x_n, y_1, y_2 \dots y_n]^T$ for each of the 2D images in the training set. The principle behind landmarking may be conceptually simple, but in practice is a cumbersome and time consuming job. The tracer must manually position, sometimes hundreds, of markers along the traced contours, with constant referral to previous annotations to ensure correspondence. This becomes increasingly more difficult as more and more data presents itself from 3D and 4D medical scans. Some work in automatic landmarking has been researched in [170, 129, 50, 169, 135]. Once the landmark points have been selected, they are then aligned commonly with Procrustes shape distance metric with respect to scale, position and orientation. As stated, point correspondence is one of the limitations for model based approaches and Hamarneh [60] addresses this problem by represented in the shapes by descriptors obtained after the application of Discrete Cosine Transform (DCT).

To model the shape variation, the classical statistical approach of eliminating redundancy in the database is achieved through Principal Component Analysis (PCA) or *Karhunen-Loeve* transform. PCA performs a variance maximising rotation of the original variable space, this is best illustrated graphically in Figure 5.2 where the two principal axes of a two dimensional data set is plotted and scaled according to the amount of variation that each axis explains [149, 45]. The axes are also ordered according to their variance, meaning the first axis contains the highest variation. In practice the PCA is performed as an eigenanalysis of the covariance matrix of the aligned shapes.

The overall idea behind ASMs is to generate a shape instance using the data obtained from the training set of shape landmarks. This can be seen in Equation 5.4 where \mathbf{x} is the new shape instance and $\bar{\mathbf{x}}$ is the mean shape (see Equation 5.5).

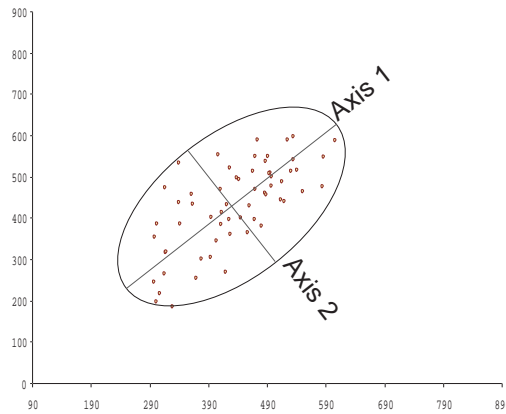


Figure 5.2: Principal axes. A 2D example where axis 1 and axis 2 are the first two eigenvectors.

$$\mathbf{x} = \bar{\mathbf{x}} + \Phi_s \mathbf{b}_s \quad (5.4)$$

$$\bar{\mathbf{x}} = \frac{1}{N} \sum_{i=1}^N \mathbf{x}_i \quad (5.5)$$

The matrix $\Phi_s = [\Phi_1 \dots \Phi_t]$ is made up of the eigenvectors corresponding to the t largest eigenvalues λ_i , where t is the number of modes. \mathbf{b} is a vector defining the set of parameters of the deformable model and is defined in Equation 5.6.

$$\mathbf{b}_s = \Phi_s^T (\mathbf{x} - \bar{\mathbf{x}}) \quad (5.6)$$

There are some disadvantages associated with ASMs, mainly their lack of robustness in the presence of high gradients not associated with the target object, their dependence on initialisation close to the target object, time consuming database construction and the inherent problem of model generality versus accuracy.

5.3.1 Application of ASMs to Medical Imaging

In 1994, Cootes [34] published his work on localisation of medical features using ASMs and used the left ventricle in echocardiographic sequences. Hamarneh and

Gustavsson [60] also apply the ASM to echocardiographic sequences to locate the left ventricle and then in a second phase uses the active contours described in Section 5.1 to accurately determine the true boundaries of the left ventricle. van Ginneken et al. [168] uses a non-linear kNN-classifier instead of the more commonly used linear Mahalanobis distance metric to steer the active shape segmentation scheme to optimal local features. Duta and Sonka [45] improve the ASM by constraining the deformation of the shape model to appropriate shapes defined by the segmentation task, in their case the segmentation of brain images in MRI. Rogers and Graham [120] perform a robust parameter estimation to improve tolerance of outliers in the model and improve the ASM search.

5.4 Active Appearance Models

In order to address some of the ASMs lack of tolerance to grayscale variation of the unseen data, Cootes et al. [35] introduce Active Appearance Models (AAMs). AAMs build on ASMs by including shape and textural information about the manually delineated training data. Textural information is defined as the pixel intensity values across the object and these values are stored in a vector $\mathbf{g} = [g_1, g_2, \dots, g_m]^T$ where m denotes the number of pixels contained within the object surface. Alignment of the texture shapes is achieved through image warping, one such method of image warping is Piece-wise affine using Delaunay triangulation (refer to [149] for more details). This is followed by normalisation with respect to illumination of the images before the PCA is constructed as described in Section 5.3 or in more detail in [149]. A single instance from the texture model can then be extracted as,

$$\mathbf{g} = \bar{\mathbf{g}} + \Phi_g \mathbf{b}_g \quad (5.7)$$

In order to combine the shape and texture models, the shape and model parameters \mathbf{b}_s and \mathbf{b}_g can be combined using a third PCA to make the representation more compact.

There are many advantages to the method. For instance,

- due to the training phase, the segmentation is very task specific,
- once initialised, convergence is fast,
- AAMs are non-parametric and

- in certain situations, are robust against noise.

There are also some limitations associated with the AAM

- the model must contain distinct features, unpredictable objects such as pathologies cannot be handled,
- the annotation of the training set is an arduous task,
- the results are inherently dependent on close initialisation to the target object
- the size and variation of the training set can restrict the AAM from converging on the correct solution and
- the AAM assumes point correspondence of the training data.

5.4.1 Application of AAMs to Medical Imaging

AAMs have received much attention in medical imaging in recent years. Stegmann [148, 149] performed a segmentation of the left ventricle of the heart using the AAM on 2D perfusion images. In [151], Stegmann and Larsson use a clustering method of the texture variation to create a set of texture subspaces, which could represent the phases of bolus passage in cardiac perfusion MRI. Mitchell et al. [96] [94] demonstrate the results when a 3D AAM\ASM combination is performed on the left ventricle of the heart in cardiac data in MRI (see Figure 5.3) and ultrasound images. The model is created using manually traced contours on 2D slices and extended in the z direction using linear interpolation between slices. Van der Geest [165, 166] investigates the semi-manual use of AAMs for the segmentation of the myocardium in MRI data over the entire cardiac cycle. Firstly, the contours are initialised on one image and the model iterates over the entire cardiac cycle until convergence. Finally, manual readjustment of the final model fittings can then be performed.

Bosch et al. [17] examine the use of Active Appearance Motion Models (AAMMs) in MRI and echocardiographic. AAMMs introduces a time factor into Active Appearance Models which aims to minimise the appearance-to-target differences. Lelieveldt et al. [79] and Sonka et al. [143] also use AAMMs in segmentation of cardiac 2D+time MRI sequences. The major advantage of this method over AAMs is the error feedback parameters are calculated for the full image sequence ensuring a segmentation consistent with cardiac motion.




Figure 5.3: Original image (left) following segmentation using AAM (middle) and method described in Mitchell et al. [95] (right).

5.5 Atlas Based Segmentation

Atlas based approaches are parameter free deformations of *a priori* models to extract the target object in an unseen image (for a full review of model based approaches see [48]). In this case, prior knowledge about the shape and intensity values of the object are incorporated. Unlike parametrically deformable models, which use geometric shapes to model the desired shape, atlas based approaches construct the model from manually segmented data.

5.5.1 Application of Atlas Based Methods in Medical Imaging

Kaus et al. [69] use coupled triangular surface meshes to segment the epi- and endo-cardial contours. Prior knowledge is encapsulated from the manually segmented data using a point distribution model as well as the grey level appearance within the myocardium. Lorenzo-Valdés et al. [82] construct a probabilistic atlas of manually segmented temporally aligned data. Automatic segmentation is achieved by registering the atlas on the data, using the atlas as the initial values for a Expectation-Maximisation (EM). The EM is then iterated until convergence before a final classification step using Markov Random Fields (MRF) and Largest Connected Components (LCC). Lelieveldt et al. [80] proposes a method for thoracic volume segmentation by building a model of the anatomical structures contained in the thoracic cavity. The method uses blended fuzzy implicit surfaces and a solid modelling technique called constructive solid geometry (CSG). Initialisation of the model with respect to position, orientation and scaling is one limitation of Lelieveldt's argument.

5.6 Level-set Method

Level sets were first introduced by Osher and Sethian [108], following previous work in Sethian's Ph.D. thesis [132] on flame propagation. Like snakes, the theory behind this boundary-based segmentation is largely based on work in partial differential equations and the propagation of fronts under intrinsic properties such as curvature [133]. While level-set methods can be applied to a host of image processing problems, for example image restoration, inpainting, tracking, shape from shading and 3D reconstruction, segmentation is the main focus of this work. An extensive review of level-set methods is given by Suri [157] and also by Angelini et al. [5]. It can also be thought of as transforming the earlier work of Kass et al. [68] on active contours from a Lagrangian to a Eulerian formulation. Like active contours, the deformation of the level set is seen as a gradient flow to a state of minimal energy, providing the object to be segmented has clearly identifiable boundaries [22, 21, 23, 87, 86].

However, by extending the dimensionality of the problem to $N + 1$, where N is the initial dimension of the problem, some advantageous properties can be exploited. The formulation of the problem is conceptually simple. The evolving curve, or front Γ , evolves as the zero level-set of a higher dimensional continuous function ϕ .

$$\begin{aligned} \frac{\partial \phi}{\partial t} + F|\nabla \phi| &= 0 \\ \phi(s, t = 0) &= \textit{given} \end{aligned} \tag{5.8}$$

This function deforms with a force F that is dependent on both curvature of the front and external forces in the image. The force acts in the direction of the normal to the front. The initial position for the contour is given, so therefore the function ϕ can be constructed.

The use of level-sets for the segmentation of the cardiac muscle in MRI is appropriate for the following reasons:

- one can perform numerical computations involving curves and surfaces on a fixed Cartesian grid without having to parametrize these objects (Eularian, non-marker based solution),
- it becomes easy to implicitly track shapes which change topology, for ex-




Figure 5.4: Level-set representation of the evolution of a circle.

ample when a shape splits in two, develops holes, or the reverse of these operations,

- intrinsic geometric properties of the front, such as the curvature and the normal, can be easily calculated and
- the method may be extended to higher dimensions.

However, there are some issues associated with the basic level set formalisation. In the latter stages of this chapter, the author attempts to address these issues:

- the algorithm is computationally expensive,
- the front may leak through boundaries of low gradient information,
- the level set function requires initialisation close to the target object and
- the evolution does not use prior shape or texture based information.

Level-set segmentation has also been successfully applied to other medical imaging modalities as described in Appendix B.

5.6.1 Level Set Formulation

The fundamental objective behind level-sets is to track a closed interface $\Gamma(t)$, for which $\Gamma(t) : [0, \infty) \rightarrow R^N$, as it evolves in the data space. The interface is represented by a curve in 2D and a surface in 3D or the set of points that are on the boundaries of the region of interest Ω . The theory behind level-set segmentation is largely based on work in partial differential equations and the propagation of fronts under intrinsic properties such as curvature [108, 133]. Level-set theory aims to exchange the Lagrangian formalisation and replace it with Eulerian, initial

valued partial differential equation evolution. By extending the dimensionality of the problem to $N+1$, where N is the initial dimension of the problem, some advantageous properties can be exploited. Representing the boundary as the zero level set instance of a higher dimensional function ϕ , the effects of curvature can be easily incorporated. ϕ is represented by the continuous Lipschitz function $\phi(s, t = 0) = \pm d$, where d is the signed distance from position s to the initial interface Γ_0 (see Equation 5.9). The Lipschitz condition implies that the function has a bounded first derivative. The distance is given a positive sign outside the initial boundary ($D \setminus \Omega$), a negative sign inside the boundary ($\Omega \setminus \partial\Omega$) and zero on the boundary ($\partial\Omega$).

$$\phi(s, t = 0) = \begin{cases} -d & \forall s \in \Omega \setminus \partial\Omega \\ 0 & \forall s \in \partial\Omega \\ +d & \forall s \in R^n \setminus \Omega \end{cases} \quad (5.9)$$

From this definition of ϕ , intrinsic properties of the front can be easily determined, like the normal $\vec{\mathbf{n}} = \pm \frac{\nabla\phi}{|\nabla\phi|}$ and the curvature $\kappa = \nabla \cdot \frac{\nabla\phi}{|\nabla\phi|}$.

Also from this definition, ϕ can be considered as a function in two different ways. Firstly, ϕ can be considered as a static function $\phi(s)$ that is evaluated at particular instances or isovalues, this leads to the formulation of the Eikonal equations and is discussed in more detail in the Fast Marching section (Section 5.6.5). Alternatively, ϕ can be described as a dynamic function $\phi(s, t)$ that evolves through time, and the closed contour or front is the special case where the value of $\phi(s, t)$ equals zero. Using this definition, it can also be said that at any time t_0 the set of points that define a curve can be represented as the function $\phi(s, t_0) = 0$. It is also clear that as the curve evolves through time, the function ϕ also evolves. Consider a point $s(t)$ on the contour that is evolving through time, we constrain the value of that point in the level-set function to be $\phi(s(t), t) = 0$. By chain rule,

$$\frac{\partial\phi}{\partial t} + \nabla\phi \cdot s'(t) = 0 \quad (5.10)$$

Define the force, $F = s(t) \cdot \vec{\mathbf{n}}$ to be the force moving the point $s(t)$ in the normal direction $\vec{\mathbf{n}}$. If $\vec{\mathbf{n}}$ is replaced with $\vec{\mathbf{n}} = \pm \frac{\nabla\phi}{|\nabla\phi|}$, the equation takes the form of a Hamilton-Jacobii as expressed in Equation 5.8. If the force term is rewritten as $F = F_0 + \epsilon\kappa$ to include an advection force F_0 to move the curve and a curvature

based term $\epsilon\kappa$ to regulate the evolution, the evolution of ϕ now becomes,

$$\phi_t = -F_0|\nabla\phi| + \epsilon\kappa|\nabla\phi| \quad (5.11)$$

Classic finite difference schemes for the evolution of this equation tend to overshoot and are unstable. Sethian [134] has proposed a method which relies on a one-sided derivative that looks in the up-wind direction of the moving front to control the outward expansion, and thereby avoids the over-shooting associated with finite differences (see Equation 5.12) while the second derivative can be approximated using central differences. Level-set theory uses a combination of derivative approximations to enable smooth curvature evolution.

$$|\nabla\phi| = \sqrt{\max(D_{i,j}^{-x}, 0)^2 + \min(D_{i,j}^{+x}, 0)^2 + \max(D_{i,j}^{-y}, 0)^2 + \min(D_{i,j}^{+y}, 0)^2}, \quad (5.12)$$

where, for example $D_{ij}^{+x} = \frac{\phi(i+1,j) - \phi(i,j)}{\Delta x}$ and $D_{ij}^{-x} = \frac{\phi(i,j) - \phi(i-1,j)}{\Delta x}$.

Caselles et al. [22] and Malladi et al. [87] used the above theory to independently formalise the implicit minimisation of the classic energy function used in snake evolution, seen in Equation 5.1, for the extension to level set theory.

$$\min \int g(|\nabla I|, I_\sigma) |\Gamma'(s)| ds \quad (5.13)$$

This minimisation includes a stopping term $g(|\nabla I(\Gamma(s))|)$ where g is a stopping function (reciprocal or exponential) based on gradient of pixel intensities and curvature term $\Gamma(s)$ based on the intrinsic properties of the curve and calculated by

$$\int_S |\Gamma'(s)|^2 ds = \int_S g(|\nabla I \Gamma(s)|) ds \quad (5.14)$$

From [22] it can be shown that the Euler-Lagrange gives a minimising curve that is of the form.

$$\frac{d}{dt} \Gamma(s) = g(|\nabla I|) \kappa \vec{\mathbf{n}} - (\nabla g \cdot \vec{\mathbf{n}}) \vec{\mathbf{n}} \quad (5.15)$$

The term $\nabla g \cdot \vec{\mathbf{n}}$ adds a naturally occurring attraction force vector normal to the surface introduced by Yezzi et al. [177] and κ is the curvature term. By representing the boundary as the zero level set instance of a higher dimensional function ϕ as described in Equation 5.9, the effects of curvature can be easily

incorporated [86, 21].

$$\frac{\partial \phi}{\partial t} = g(|\nabla I|)(c + \epsilon \kappa)|\nabla \phi| + \beta(\nabla g \cdot \nabla \phi) \quad (5.16)$$

In this segmentation scheme, a non-zero internal advection or ballooning force, c , is added to the evolution to evolve the either outward ($c = 1$) or inward ($c = -1$) depending on whether the initialisation curve is enclosing or enclosed by the target object for segmentation. β and ϵ are user defined parameters that control the effects of attraction to gradients and curvature respectively and are chosen experimentally. Reducing the β parameter slows down the convergence time as the front is not attracted to edges, however increasing the parameter may have the effect of causing the evolution to jump past appropriate gradients leading to spilling of the curve into other areas. The parameter ϵ controls the smoothness of the contour or surface. Reducing the value of this parameter allows the algorithm to converge on less smooth object boundaries.

Curvature Term

From differential geometry any shape (no matter how complex) collapsing as a function of its curvature κ will evolve to a circle before disappearing [59]. Using this relationship, a force $F = -\kappa$ is defined to always shrink a contour to a point. This is a favourable quality for advancing fronts for segmentation, as it can be shown that this minimises the contour length. As discussed earlier, using the partial differential equations perspective, intrinsic geometric properties such as the curvature and normal can be easily calculated. For example, for a 2D propagating front, the curvature κ can be found using partial differentials of the function ϕ ,

$$\kappa = \nabla \cdot \frac{\nabla \phi}{|\nabla \phi|} = \frac{\phi_{xx}\phi_y^2 - 2\phi_y\phi_x\phi_{xy} + \phi_{yy}\phi_x^2}{(\phi_x^2 + \phi_y^2)^{3/2}} \quad (5.17)$$

The normal can undergo a jump at corners, and this issue is addressed in the work of Sethian and Stain [134] where the normal is normalised.

However, in 3D there are two measures of curvature, the mean and Gaussian curvature. The mean curvature (κ_H), is connected to the physical evolution of soap bubbles and the heat equation.

$$\kappa_H = \nabla \cdot \frac{\nabla \phi}{|\nabla \phi|} \quad (5.18)$$

Gaussian curvature (κ_G), has also been used to model physical problems such as flame propagation. It has been shown that a convex curve evolves to a point under curvature evolution, but it can also be shown that evolution of non-convex surfaces can be unstable.

$$\kappa_G = \frac{\nabla \phi^T \text{Adj}(\mathbf{H}(\phi)) \nabla \phi}{|\nabla \phi|^2} \quad (5.19)$$

where $\mathbf{H}(\phi)$ is the Hessian matrix of ϕ , and $\text{Adj}(\mathbf{H})$ is the adjoint of the matrix \mathbf{H} [173].

Neskovic and Kimia's [106] propose a measure of curvature which involves both mean and Gaussian. In this approach, the direction of flow is obtained from the Mean curvature while the magnitude of the flow is dictated by the Gaussian curvature. This is appropriate as the Mean curvature alone can cause singularities when evolving.

$$\kappa_{nes} = \text{sign}(\kappa_H) \sqrt{\kappa_G + |\kappa_G|} \quad (5.20)$$

Stopping Criterion

The evolution force F is an energy minimisation problem where the speed approaches zero at positions of high gradients to exert a halting to the front propagation. To this end, two diffusive stopping criteria have been proposed. The first and most common stopping term is a reciprocal of the gradient of the image intensity signal convolved with a Gaussian smoothing mask G_σ , where σ is the variance of the Gaussian mask.

$$g(|\nabla I|) = \frac{1}{1 + |\nabla G_\sigma \cdot I(s)|^p}, \quad p \geq 1 \quad (5.21)$$

The convolution with a Gaussian eliminates the effects of noise on the image. Other methods of noise removal, such as non-linear or anisotropic which were discussed in Chapter 3 can be used in place of the Gaussian to improve the results.

Alternatively, if the circumstances require the stopping term to fall to zero faster than the reciprocal function, the following definition can be applied. This may allow the front to overshoot the object boundary in the presence of weak gradients or can cause errors in the presence of noise. Therefore a new stopping term is devised that incorporates texture. This is performed by means of a Gaussian membership function used to determine whether the voxel is inside or outside the target object. This membership function is constructed using the texture analysis of the object region after initialisation. A Gaussian membership function is chosen as MRI response in tissue can be modeled as a Gaussian distribution [76, 70].

$$g(|\nabla I|) = \exp^{-|\nabla G_\sigma \cdot I(s)|} \quad (5.22)$$

where \exp is the exponential function.

5.6.2 Non-gradient based curve propagation

Image segmentation and classification has also been approached by incorporating level sets into the partitioning of images based on intensity values. These methods have also been called Region-competition snakes and are deformable models that are governed by local probabilities that determine if the snake is inside or outside the structure to be segmented. Chan and Vese [27] show how the Mumford-Shah functional can be used in a level set framework. The Mumford-Shah functional aims to partition the image I into a smooth approximation f set of regions separated using contours, S .

$$E(S, f) = \nu(S) + \alpha \int_{\Omega} (f - I)^2 dx + \beta \int_{\Omega_S} |\nabla f| dx \quad (5.23)$$

The problem is approached as a energy function which tries to minimise its variables (a) the length of the set of contours $\nu(S)$, (b) the deviation from the original image $\alpha \int_{\Omega} (f - I)^2 dx$ and (c) the smoothness within each region $\beta \int_{\Omega_S} |\nabla f| dx$.

Another approach is developed in a level set framework by assuming a two class problem of an image I defined on Ω . The problem is then posed as follows:

$$E(C) = \int_{insideC} |I - C_0|^2 d\Omega + \int_{outsideC} |I - C_1|^2 d\Omega \quad (5.24)$$

where C is the front, and (C_0, C_1) are the average intensity values for inside and outside the curve C . While this methods addresses boundary leakage and ini-

tialisation problems it assumes a low class of intensity features and grey scale homogeneity across the object. Yezzi et al. [177], Tsai et al. [163], Cohen and Kimmel [32], Deschamps et al. [41] and Angelini et al. [6] adopt variations on this approach to segmentation in medical images.

In [111], Paragios and Deriche unifies both region and boundary information in a level set framework. Following on from [109], Paragios incorporates an intensity based component taken from the grey scale distributions of cardiac features and a prior shape model to deform a coupled level set over the endo and epi-cardium of the heart. Taron et al. [161] perform a variational technique for the segmentation of the *Corpus Callosum* of the brain. They use estimated uncertainties of the registration when applying model priors to the segmentation process.

5.6.3 Introduction of A-Priori Knowledge

Leventon [81] introduced *a priori* knowledge by building a prior model that was embedded in a level set formalisation and evaluating its modes of variation using PCA analysis. This has been the basis for much work in level set formulation incorporating shape priors into the propagation. Due to the model being defined in Eulerian space, it circumvents the problem of point correspondence encountered in the previous sections.

Tsai et al. [163] provide some work, leading from the initial work performed by Leventon and perform segmentation on cardiac images in 3D. In [162], Tsai et al. construct a model of a priori shapes as the zero level set of a number of separate segmented images. The database of level sets are then classified into a user defined number of statistical shape classifications using an Expectation Maximisation algorithm. This method was applied to medical images where congenital brain malformation of the cerebellums was used to create a two class (healthy/diseased) classification scheme. As mentioned, Paragios et al. [110, 112] use a shape model built from previously segmented data to guide the segmentation of his level set.

5.6.4 Coupling of Level Sets

Zeng et al. [181] first introduced the idea of coupled level sets for segmentation of the cortex of the brain. The coupled level set can use the constant thickness or

distance between the level-sets as a constraint to avoid spilling or over segmentation. The ideas introduced by Zeng were extended by Paragios [109] who applied a similar coupling constraint for the segmentation of the myocardium of the heart.

5.6.5 Initialisation using Fast Marching

In order to overcome the 'myopic' characteristics of level set propagation, Sethian [133] introduced a Fast Marching methods. This is the unique case of the level set theory where the force F is always greater than zero, and this propagates a monotonically advancing front. The formula takes the form of the Eikonal Equation 5.25, a nonlinear, static Hamilton-Jacobi equation. If the 2D case is considered again, a set is created $T(x, y)$ that defines the time at which the front Γ crosses the position (x, y) . T satisfies the equation,

$$|\nabla T|F = 0 \quad (5.25)$$




Figure 5.5: Front propagation using Fast Marching. Adapted from Sethian [134].

The evolution is iteratively assessed by solving the roots of the quadratic equation of the Eikonal equation and sorting the values of T with respect to size.

This can be shown graphically in Figure 5.5. The value of the force at each point can be calculated using the upwind scheme approximations to derivatives of the function ϕ .

The stopping term is based on the diffusion of the gradient and can be calculated as

$$F_{ij} = e^{-\alpha \nabla I_{ij}} \quad (5.26)$$

The fast-marching approach gives an approximate segmentation and is used for the evaluation of the initial contour for the dynamic level-set method.

5.6.6 Narrow-band Methods

In order to increase the computational efficiency of the algorithm, Adalsteinsson and Sethian [2] extensively review narrow-band methods. The main disadvantage of formulating the problem in Eulerian space as opposed to the Lagrangian space is the increase memory and computational expense of propagating the front across the full matrix of the image. To eliminate this issue, a narrow band (2D), or narrow tube (3D), around the front is defined and it is in this narrow band that the ϕ values are updated at each iteration. The narrow-band is first initialised by including all data points within a certain bandwidth of the front, this can be achieved by using the values of the ϕ . As explained, at each iteration, only the values of ϕ within the narrow band are updated. With each iteration the front points are evaluated to see if they are close to the edge of the narrow band. If yes, the narrow band is re-initialised otherwise the algorithm iterates as normal. It has been shown in [2] that these boundary conditions do not adversely affect the motion of the level-set. Implementation of this narrow band method can greatly improve speed of execution and some level set approaches prove real-time execution [37].

5.7 Initialisation

To counteract the 'myopic' characteristics of these deformable models, the initialisation process is very influential and is performed in MRI data as follows. Firstly, it is known that the endocardium boundary can be characterised by the high contrast between the blood and the heart muscle in standard (TruFISP) cine imaging of the heart. This characteristic is used when a fast marching algo-

rithm is applied to find a fast efficient initialisation for the blood following the manual insertion of a seed point. The fast marching approach is driven by a force $F_s = e^{-\alpha \nabla I_s}$, which has a diffusive effect aimed at halting the fronts progress at regions of high gradient. This fast-marching approach falls short of the gradient defining the transition from blood to muscle. Therefore the contour found by the fast marching algorithm is used as the initial curve of the level-set algorithm to find the endocardium boundary. The results from the Fast Marching initialisation are illustrated in Figure 5.6.




Figure 5.6: Results show the initialisation (marked in white) from a seeded Fast Marching algorithm. The method was applied to perform a robust initial estimate of left ventricle cavity of the heart on four separate datasets displaying a high variability of left ventricle shape.

To find the epi-cardial boundary the endocardium initialisation is dilated slightly and the inner gradients are masked. Both curves are given a positive advection force to propagate outwards. It is known that both the endo- and epi-cardium boundaries of the left ventricle are approximately circular, therefore the ϵ is given a high significance in the evolution, the evolution is illustrated in Figure 5.7. High curvature constraints, the distance inhibitor and the *a priori* constraints all act to limit the epi-cardium front from joining the inner front or spilling in areas of low gradient, like the liver or the lungs.

5.8 Coupling Force between Fronts

To further control the level-set evolution we employ a coupling function between two level-sets. The coupling adds an extra constraint by introducing a second level-set that is dependent on the first and coupling the level-sets with an inhibitor function, which allows the curve to change direction of growth. This is




Figure 5.7: The images above show evolution of the front at four different iterations (a) *iteration* = 0, (b) *iteration* = 5, (c) *iteration* = 10 and (d) *iteration* = 15.




Figure 5.8: Segmentation results of the same slice at three separate phases through the hearts cycle, (a) end-diastolic, (b) mid-diastolic and (c) end-systolic.

achieved without any extra computational expense as the distance between any point to the level-set boundary is the value of ϕ at that point, see Equation 5.9. The piecewise inhibitor function, which is used as the interaction between the two level-sets, is defined below, where d is the preferred distance between the curves and w controls the slope between inward and outward growth. The result $\eta_2(\phi_1)$ changes value from $+1$ to -1 , which changes the direction of the evolution for ϕ_2 between inwards and outwards. In practice the values of d and w are taken from the scaled *a priori* model.

$$\eta_2(\phi_1) = \begin{cases} -1 & \text{for } \phi_1(s) < -d - w \\ \sqrt[3]{\left| \frac{\phi_1(s) - d}{w} \right|} & \text{for } -d - w < \phi_1(s) < d + w \\ 1 & \text{for } \phi_1(s) > d + w \end{cases} \quad (5.27)$$

For this segmentation scheme, it is assumed that the gradient between the blood pool and the endo-cardium boundary is significantly high to halt the evolution of the level-set. Also it is known that in some cases there is little or no




Figure 5.9: Graph of the inhibitor function where the values of $d = 6$ and $w = 4$.

gradient information between the epi-cardium boundary and the lungs or liver. Therefore, the level-set segmenting the epi-cardium boundary is controlled by the endo-cardium level-set using the inhibitor function described.

5.9 Improved Stopping term

To illustrate the improved performance of the advanced stopping term, the following phantom images were created and tested. Two situations are described as illustrated in Figures 5.10 and 5.11, the first where low gradient information is present between two regions and the second where the grayscale difference between two regions is low. The stopping term, is defined as:

$$g = \frac{1}{1 + \frac{\nabla I}{I_\sigma}} \quad (5.28)$$

uses a combination of the gradient and change in texture. The change in texture (I_σ) is calculated after the initialisation with the fast marching algorithm described in Section 5.7. Within the initialised region the mean μ and variance σ of the voxels are calculated. From these values, a Gaussian is constructed and the $I_\sigma(s)$ is calculated as,

$$I_\sigma(s) = \frac{1}{\sqrt{2\pi\sigma^2}} e^{-\frac{(x-\mu)^2}{2\sigma^2}} \quad (5.29)$$

where x is the value of the voxel at each position s in the image. The value of I_σ is normalised between 0-1.

The user defined parameters ϵ and β represent the influence of the curvature and attraction to gradient on the evolving boundary. In the following tests, we want to evaluate the influence of the improved stopping term, so the value of ϵ is given less significance to reduce the influence of curvature on the evolution. In the segmentation of the left-ventricle boundaries, the value of ϵ is given a higher significance as we know the boundaries approximate circles. Similarly, β controls the attraction of the level-set boundary to gradients that are normal to the curve. Again, this value is given a reduced weighting in the proceeding tests. The results shown in Figures 5.10 and 5.11 demonstrate the improved robustness against boundary leaking between regions.




Figure 5.10: The original phantom image with a diffused segment (a) and the Sobel edge image to illustrate the gradient information (b). The second row shows the evolution with the existing $g = \frac{1}{1+\nabla I}$ at iteration 0, 25 and 50 while the third row shows the evolution with our proposed approach where $g = \frac{1}{1+\frac{\nabla I}{I\sigma}}$ at iteration 0, 25 and 50.

5.10 Introduction of Priors Models

A priori information is incorporated with a probability density function (PDF), which is defined as

$$P(s) = \frac{\sum_S \sum_{i=1}^N f_i(s)}{N} \quad (5.30)$$




Figure 5.11: The original phantom image with a close region (a) and Sobel edge image to illustrate the gradient information (b). The second row shows the evolution with the existing $g = \frac{1}{1+|\nabla I|}$ at iteration 0, 25 and 50 while the third row shows the evolution with our proposed approach where $g = \frac{1}{1+\frac{|\nabla I|}{\sigma}}$ at iteration 0, 25 and 50.

where f_i is the outline of the epi and endo cardium boundaries used for training, N is the number of training examples and s defines the image coordinates. The model is built from a set of hand segmented boundaries, a probability density function is created of both the endo-cardium and epi-cardium boundaries that are then interpolated in the z direction, scaled and aligned in the xy direction.

The PDF is constructed by aligning the binary manually segmented boundary images and summing the boundary elements. This is done for both the endo-cardium boundary and the epi-cardium boundary. It is incorporated into the evolution in a global context, after each iteration the value ρ_t is evaluated as,

$$\rho_t = \sum_{C \in S} \phi(t)_s * P_s \quad (5.31)$$

where $\phi(t)_s$ is the value of ϕ at time t at the position s and P_s is the probability density at position s and this value is summed over the narrow band C which is a subset of the image space. The parameter ρ_t is calculated at each iteration is then normalised between the bounds -1 and 1 as it can have negative and positive values. This is as a result of ϕ also having positive values outside the contour and negative values inside the contour. This means ρ_t will have a more positive

value if the current contour is inside the prior model and more negative if the contour is outside the prior model.

In order to obtain the full evolution equation for the level-set we have to incorporate both the coupling function and the *a priori* knowledge into Equation 5.32.

$$\frac{\partial \phi}{\partial t} = g(\nabla I)(c + \epsilon \kappa)|\nabla \phi| + \beta(\nabla g \cdot \nabla \phi) \quad (5.32)$$

Firstly, the output from the coupling function is either 1 or -1 and we want it to change the direction of the curve evolution. From Equation 5.32 we can see that the advection force defines the direction of the evolution, therefore we incorporate the coupling function by multiplying it with the advection force c . This has the result of changing the direction of the contour, depending on the results from the coupling function. In this sense, both the epi and endo cardium boundaries are tied together. We also assume that the boundary between the left ventricle blood pool and the myocardium has a stronger gradient term than that of the epi-cardium boundary and the liver or lungs. Therefore, this term is applied to the evolution of the level-set surface designed to extract the epi-cardium. Hence, based on the parameters of the coupling function which can be automatically obtained using the distance between the blood pools the outer surface is prohibited from spilling into other organs beyond a certain distance from the endo-cardium boundary.

The *a priori* model is designed to disregard inappropriate gradients and to give significance only to gradients that are situated close to previously manually segmented boundaries. For this reason, we incorporate the *a priori* information in the attraction term from Equation 5.32. As explained, this is taken on a global sense whereby we define for both the inner surface and the outer surface whether or not they are inside or outside the PDF of previously segmented images. Thus, the complete evolution for the coupled level-set is defined as,

$$\phi_{t+1} = \phi_t + g(\nabla I)(c\eta + \epsilon \kappa)|\nabla \phi| + \frac{\beta}{1 + \rho_t}(\nabla g \cdot \nabla \phi) \quad (5.33)$$

where η is the result of the coupling function between the level-sets and is defined in Equation 5.27 and ρ_t is the *a priori* knowledge and is defined in Equation 5.31. The results in Figure 5.13 illustrate the performance using four unseen datasets




Figure 5.12: Images show the probability density functions from a priori hand segmented images. Figure (a) shows the combined contours while (b) and (c) show the endo- and epi-cardium boundaries respectively. Darker gray tone defines a higher probability of the boundaries.




Figure 5.13: The images above show the segmentation using our method on the four previously unseen datasets.

5.11 Extension to 4D

Cardiac data is increasingly available in 3D + time, therefore it is believed that the best approach for a complete data driven segmentation is to apply an appropriate technique to the complete data presented from a patient scan. Due to the increasing amount of data that is available in 4D and growing resolution, some researchers have attempted to address the segmentation problem. Many have evaluated the result of sequential approaches, where from a robust initial segmentation (maybe manually assisted) forms the initialisation for subsequent volumes throughout the cardiac cycle.

While the level set formulation lends itself easily to extension in multidimensional data analysis, the author found few researchers have investigated the application of level set to analysis of 4D data. Fritscher et al. [51] aim to apply full 4D information into boundary driven and region-competition geodesic con-

tours. In initial work, PCA analysis is performed on signed distance maps to create models, the mean of these models serve as the initialisation step in a level set segmentation. More generally, in the earlier 4D segmentation work [9, 91], the temporal dimension was considered in a sequential approach where the segmentation from the previous time frame served as the initialisation for the current time frame. Rueckert and Burger [123] also used this sequential approach where the shape of $(t + 1)$ was a deformation of the shape in time frame (t) . The deformation is achieved using energy minimisation of the deformable template in a Bayesian formulation. Sun et al. [155] create a non-linear dynamic model learned from training data. A manual tracing of the first image in the sequence is used to create a posterior density estimate of the lv at each time frame. A curve evolution is then performed with the maximum posterior estimate. McEachen and Duncan [89] perform tracking of the left ventricle by performing point correspondence of points from time t to time $t + 1$ and assume a small degree of motion between time frames. Based on these assumptions, smooth transition of the parametric contours is achieved using an optimisation algorithm. Paragios [109] introduced an energy into his variational level set approach that enforced a consistency of intensity through the temporal cycle. A transformation is calculated between time I_t and I_{t+1} based on a bounded error function, where I_t represents the intensity value at time t . In Montagnat and Delinette [99] in 2005, the deformable model is influenced by introducing time-dependent constraints. These consist of prior temporal knowledge through either temporal smoothing or trajectory constraints.

Segmentation in 4D should perform a segmentation of the 3D volumes and use information in the time domain. To this end, a number of approaches are proposed with the advantages and disadvantages of each discussed.

- **Sequential Approach**, consists of naively using the results from time sequence t as the initialisation for time sequence $t + 1$. This approach assumes no prior knowledge about the temporal dynamics of heart. The only assumption is that the cardiac muscle boundaries do not exhibit large movements between time sequences.
- **Temporal subtraction**, can give some indication as to the direction of movement of the cardiac boundaries. Again, this does not utilise prior knowledge about the global dynamics of the heart and may be overly sen-

sitive to noise and artifacts. Some optical flow approaches may eliminate these limitations and are being investigated in [10].

- **Temporal Smoothing**, basically constitutes performing the segmentation of the 3D volumes in parallel while forcing the boundaries to move in a physically consistent way using temporal smoothing. In its simplest form, temporal smoothing could be achieved using an averaging function, $\Gamma_t = \frac{\Gamma_{(t-1)} + \Gamma_{(t+1)}}{\Delta t}$, where Γ_t represents the boundary curve at time t .
- **Temporal consistency** of intensity values across the left ventricle cavity and the left ventricle myocardium and was employed by Paragios and Deriche [111]. Again, artifacts in the left ventricle cavity due to the dynamics of the blood through the cardiac cycle may restrict the application of this method.
- **Database of Prior Image Models**, built from a selection of images at particular temporal instances, may be registered to the unseen image. Like many database models, this approach relies on building generic models that are applicable to a wide range of heart morphology. Variations in cardiac morphology caused by individual anatomical features or disease may not be accounted of in such models.
- **Prior Temporal Parameterised Model** proposes to model the dynamics of the cardiac cycle and further refine this model as the parallel segmentation is performed on the 3D volumes. Unlike database models constructed in image space, broader classification of the cardiac boundaries movement through the entire cycle can be applied to all variations of heart morphology. Exploiting the construct of the ϕ function in level set segmentation (see Equation 5.9) enables fast function fitting that may be incorporated into the update of ϕ .

5.12 Applying level set on 3D+t data

From the options above, segmentation of the 4D data should be approached in a parallel sense using temporal constraints to infer prior knowledge in an effort to control the boundary deformation away from erroneous spilling or over segmentation. To this end, a novel approach to control a level set deformation is proposed. The control is achieved by means of prior knowledge about the deformation of the cardiac muscle through a complete cardiac cycle. In the majority of cases,




Figure 5.14: Volume, in pixels, of left-ventricle cavity over the cardiac cycle.

the temporal volume change of the cardiac cavity over the complete cardiac cycle can be illustrated as show in Figure 5.14 where the phase starts at end-diastolic, decreases in volume during the systolic phase until it reaches end-systole before returning to end-diastole during its diastolic phase.

The next question to pose is how this information about the overall shape of the cardiac phase can be implemented in a loosely fitting way to the deformation of the level set.

5.12.1 Modelling the temporal movement

From Figure 5.14, the cardiac cycle can be approximated using an inverted Gaussian curve. Values for the general Gaussian defined in Equation 5.34, A , B , μ and σ are found by fitting a Gaussian curve to the volume data extracted using the Fast Marching algorithm from each time frame. Gaussian fitting is achieved using least squares approximation. Non-linear fitting is unstable due to the low number of volumes in the temporal resolution (~ 25). For nonlinear least squares fitting to a number of unknown parameters, linear least squares fitting may be applied iteratively to a linearized form of the function until convergence is achieved. However, it is often also possible to linearize a nonlinear function at the outset and still use linear methods for determining fit parameters without resorting to iterative procedures.

$$y(x) = A + Be^{-\frac{(x-\mu)^2}{2\sigma^2}} \quad (5.34)$$

This fitted Gaussian represents the model for the dynamics of the cardiac cycle over a single heartbeat. It follows that the displacement of the endo-cardium boundary can also be modelled using this fit. In this regard, the deformation of the boundary surface of the level-set is constrained by this Gaussian model. Exploiting the inherent definition of the level-set function ϕ as the distance function of a single position from the evolving surface, the incorporation of the Gaussian model is straightforward and can be applied in a non-rigid sense to every point within the narrow-band.

This is further illustrated in Figure 5.15, where a 2D image is taken and a single point is selected within the narrow-band. From the definition of ϕ , the value at this point is the distance from that point to its closest point on the zero level-set boundary. In the illustration, the boundary contracts and then expands again in much the same way as the left ventricle boundary evolves from end-diastole to end-systole and back again to end-diastole. As this evolution takes place the value at the position grows and shrinks as the distance to the boundary increases and decreases, this evolution can be modelled using the Equation 5.34 and the parameters B , μ and σ determined from the fast marching initialisation. The value of A represents the offset of the Gaussian model. Figure 5.16 illustrates the model applied to the long axis view.




Figure 5.15: Change of a single point on ϕ as the boundary evolves over the cardiac cycle in the short axis view.




Figure 5.16: Change of a single point on ϕ as the boundary evolves over the cardiac cycle in the long axis view.

In this way, the evolution of the zero level-set boundary can be constrained to contract and expand under Gaussian motion, where the saddle point is the temporal position given by μ and deformation occurs at a rate σ . Initialisation




Figure 5.17: Volume, in pixels, of left-ventricle cavity over the cardiac cycle with fitted model using an Adaptive Gaussian Model.

of the Gaussian model parameters are determined after a primary segmentation of the left ventricle cavity using a Fast Marching method.

Level Set influenced by an Adaptive variance Gaussian

In order to model the dynamics estimated using the Fast Marching algorithm, an adaptive Gaussian model is developed. Similar to the general Gaussian model given in Equation 5.34, the aim is to improve the models fit on the initialised data. This results in the deformation of the boundary that maintains closely the temporal dynamics of the initial segmentation using the Fast Marching algorithm and therefore the model resembles the shape of the raw data and does not resemble the Gaussian curve. In practice, this model is created by a least squares fitting of a Gaussian model where the variance, σ is calculated separately at each temporal position, in essence this means that the least squares error is close to zero at each temporal position. This is illustrated in Figure 5.17 where the model curve mirrors the real data.

Models created from initialisation may not represent the final segmentation of the target object. It places too much confidence in the initial model created using the fast marching approach. For an example using the worst case scenario, if the fast marching algorithm falls into a local minima inside the left ventricle blood pool at one particular time sequence then the temporal model incorporates this. Using the curvature constraint, the level-set algorithm can overcome this error, however, the temporal model that is created may not allow the level-set to

deform greatly from the model created from the initialisation. Therefore, a new approach is proposed, which uses the information obtained from the initialisation step but iteratively updates this model based on the evolving level-set. This creates a smoothing effect on the level-set surfaces over the cardiac cycle but also redresses poor initialisation.

Level Set influenced using Expectation-Maximisation

In order to address the limitations associated with the Adaptive variance model described in the previous section, a novel approach is introduced which iteratively updates the initial parameters of the model. This acts as a form of Expectation-Maximisation (EM) algorithm. The EM algorithm is a two step approach which aims to fit some model to data, and is particularly useful where there is unknown or incomplete data. In the case of cardiac boundary segmentation, the observed data is defined as the value of the level-set function ϕ at a particular position over the entire cardiac cycle. The unknown or missing data is a final Gaussian model which is inferred on a single point in the grid over the complete cardiac cycle. This application to each point on the grid has the advantage that the model is fitted non-rigidly and can allow for less or no deformation, which is the case in diseased hearts.

The EM algorithm takes initial parameters for the model, in this case the information obtained from the Fast Marching segmentation of the left ventricle cavity, and performs an expectation or fit of the data at a particular spatial position over the entire temporal data. These model parameters are stored in an array for each grid point. Then during the maximisation step when the level-set is updated, the information about point position with respect to its expected values are calculated. The results from this expectation stage is the difference or in EM terms, the likelihood, between the model and the observed data. From this expectation calculation, a maximisation is performed to correct for the differences found. This maximisation step is the level-set deformation of the boundary surface. The process is iterative and the parameters for the model are re-evaluated at each iteration.

This addresses many issues associated with the previous method. Firstly, because the parameters for the model using the Fast Marching approach are just used as the initial parameters for the EM algorithm, there is less dependence




Figure 5.18: Volume, in pixels, of left-ventricle cavity over the cardiac cycle with fitted Gaussian model.

placed on these initial parameters as they are re-evaluated at each iteration. Secondly, the iteratively fitting a Gaussian to the data results in giving a Gaussian smoothing of the zero level-set boundary over the temporal cycle.

5.13 Results

In order to assess the validity of this approach, the results of the segmentation using the iteratively optimised algorithm are compared against those obtained from expertly validated* segmentations of the left ventricle. Figures 5.19 display a linear plot and Bland-Altman plot for the areas in 2D of the manually traced boundaries.

Comparative results between the adaptive variance approach and those obtained from the iteratively optimised algorithm can be seen on a point-to-curve error calculation in Table 5.1, showing less error using the optimisation algorithm. This is also confirmed in a linear plot of the blood pool areas when compared against manual segmentation where the Gaussian curve with adaptive variance produced a regression value of 0.71 while the optimised approach yields a regression of 0.77.

The iteratively optimised algorithm also is guaranteed convergent [174, 40, 13] and also reduces the error between the observed data and the model at each

*The validation was performed by Dr. John Murray, Cardiologist, Mater Misericordiae Hospital, Dublin, Ireland.




Figure 5.19: Results of the 4D segmentation of the left ventricle cavity boundary compared against those obtained from manual segmentation.

Table 5.1: Table representing the point to curve error for *Method 1* using the Gaussian curve with adaptive variance and *Method 2* using the Expectation-Maximisation of the Gaussian parameters.

	Endo-cardium		
	Average	Std. Dev.	RMS
Method 1	1.649013	1.584626	2.309887
Method 2	0.844075	0.914422	1.268981

iteration. This means that convergence is faster than using the static model. This is characterised in Figure 5.20 by measuring the error decay between the two methods based on known phantom data.




Figure 5.20: Results of the 4D segmentation of the left ventricle cavity boundary compared against those obtained from manual segmentation.

5.13.1 Testing under different motion approximation

In order to show the generality of the method, an implementation of the 4D segmentation was performed using a different prior temporal model. In this experiment, the temporal function is given a linear function. In Figure 5.21 a cube is expanded using a linear function. This is illustrated better in Figure 5.22 which graphs the volume acquired using the Fast Marching algorithm over time. In this graph, the fitting of a linear function to the data is also given.



Figure 5.21: Selected images from a 4D sequence demonstrating a linear volume expansion.

5.13.2 Coupled Approach

Coupling of two level-sets can also be achieved in a coherent and thorough way by employing two Gaussian models, as illustrated in Figure 5.23. Again, in a non-rigid sense each point on the grid has associated with it the parameters for two




Figure 5.22: The volume data from the sequence shown in Figure 5.21 with linear function fitted.




Figure 5.23: Estimation using prior knowledge of the Epi-cardium and Endo-cardium deformation through the cardiac cycle using inverse Gaussian curves.

Gaussian models representing the evolution of the epi-cardium and endo-cardium boundary. The evolution of the epi-cardium boundary is less pronounced and therefore the Gaussian model is shallower. Results from a coupled segmentation are illustrated in Figure 5.24 for different phases and slices.

5.14 Conclusions

In this chapter, deformable contours for feature extraction in medical imaging were introduced and discussed. An overview of current methods employed in the segmentation of the left ventricle of the heart was performed.




Figure 5.24: Results from a coupled 4D segmentation of a cardiac sequence for diastolic, systolic and mid-phase for a basal (top row), mid-slice (middle row), and apical slice (bottom row).

A novel formulation for the segmentation of the left ventricle is developed using a coupling of two level-set surfaces representing the endo- and epi-cardium boundaries. This was then extended to incorporate prior knowledge about left ventricle anatomy from manually segmented images encoded in a probabilistic model. This method provides adequate results in mid and basal slices where spilling is avoided by adding the additional constraints imposed by the prior knowledge. However, this method encountered difficulty in data representing high variation and in particular in the irregular shapes present near the apical regions. In these approaches, strengthening the *a priori's* influence on the evolution may result in loss of segmentation detail, patient abnormalities, muscle dysfunc-

tion etc. Investigating ways of improving accuracy without removing generality are addressed in the following methodology. A new supervision is proposed that does not encode the prior knowledge based on information from the image space.

A new general solution to left ventricle segmentation from 4D MRI data is presented. Temporal information obtained from the initialisation based on a fast-marching segmentation is encoded in a parametric model. The model is based on non-rigid deformation of the left ventricle boundaries over time using prior knowledge about cardiac dynamics. After each evolution of the level-set algorithm, the model is optimised to the data using an expectation-maximisation to reduce to target to object error. This approach has the following advantages; firstly, it provides a temporal smoothing over the cardiac cycle that is consistent with the motion of the cardiac muscle, secondly it constrains the boundaries from spilling in the event that a particular time instance lacks appropriate gradient information and finally, the temporal model is defined on each grid voxel within the narrow-band, this has the advantage that it can incorporate longitudinal contraction and expansion along the short axis into the model. This unique property of the temporal model can be realised due to the formulation of the level-set.

Excellent results are obtained when compared to expertly assisted segmentations of the boundaries. This method also gives comparable performance against other methods described in literature, for example Kaus et al. [69] report a mean error of 2.45 ± 0.75 mm for the end-diastolic phase and 2.84 ± 1.05 mm for end-systolic phase using a deformable model technique.

This method did not perform as accurately against the manual segmentation when comparing results to those illustrated in Chapter 4. In this application, supervision was achieved in the evolution of the boundaries by incorporating knowledge both in the temporal and space domain. Manual segmentation or the statistical partitioning techniques described in earlier chapters do use temporal information when segmenting the left ventricle. In this way, we believe that the 3D+t approach provides more accurate results, ensuring the cardiac boundaries evolve in a smooth fashion more consistent to the physical motion of the muscle. By incorporating the 4D data, we can remove inconsistencies in signal intensity values by smoothing the values over the high resolution temporal and spatial data.

The results are illustrated for a coupled surface segmentation where the left ventricle inner and outer boundaries are tracked in a computationally efficient

way using two separate models of temporal motion.

Finally, this technique represents a framework for incorporating temporal information into the evolution of an evolving surface. Also, demonstrated is a variation of this approach where temporal information is applied using a linear temporal model as the prior information. This may be associated with tracking the movement of passing objects. The complexity of the temporal model is not a limiting factor in this methodology and further applications of this technique are discussed in the following chapter.

Publications associated with this chapter

Journal Publications

In Submission

Michael Lynch, Ovidiu Ghita, Paul F. Whelan. **Segmentation of the left ventricle in $3D + t$ MRI data using an optimised non-rigid temporal model.** Submitted to IEEE Transactions in Medical Imaging, March 2006.

Michael Lynch, Ovidiu Ghita, Paul F. Whelan. **Left-Ventricle myocardium segmentation using a Coupled Level-Set with A-Priori knowledge.** Submitted to Computerized Medical Imaging and Graphics, November 2005.

Chapter 6

Conclusions and Further Developments

In this concluding chapter of the thesis, an overview of methods developed for the segmentation and tracking of the left ventricle myocardium is discussed. With particular emphasis on the aims and challenges outlined in Chapter 1, the motives for choosing particular paths in research are examined. The relevant results from each of the processes are also discussed in relation to the objectives. In the final part of the chapter, the prospect of further work is investigated in relation to the application of the proposed methods in different scenarios and also the advancement of the developed methodologies.

6.1 Summary

Diagnosis of cardiac disease can be achieved through the accurate measurement of cardiac function [103, 128]. In order to extract the most relevant clinical measurements from the heart, the thoracic cavity must be imaged and the cardiac muscle of the left ventricle needs to be segmented. MR imaging gives relatively high spatial and temporal resolution of the beating heart without the need for ionising radiation. The imaging of the heart is fast, non-invasive, painless and entails minimum discomfort to the patient.

In order to increase the accuracy, speed and repeatability of the functional measurements of the cardiac data, much research has focussed on the image anal-

ysis tasks involved in the segmentation of the cardiac muscle of the left ventricle. In this thesis, novel methods are employed in the segmentation of the left ventricle myocardium. By increasing the dimensionality of the solution thus expanding the amount of data being processed a more involved technique is developed that incorporates the three dimensional image data plus the temporal data obtained from the MRI scanner.

The problem is addressed in a systematic approach, first dealing with the inherent noise associated with the medical imaging procedures. A performance characterisation of the main diffusive based non-linear filters is provided both in 2D and 3D. The performance is evaluated using two measures, firstly the filters ability to smooth the noise in homogeneous areas and secondly the filters facility to preserve strong edges in the image using edge strength and edge spread as the criteria. The evaluation was performed in MRI data of varying protocols. From these measurements an appropriate filter is chosen as a tool to accurately remove unwanted noise from the images.

When the unwanted artifacts have been removed from the input data, statistical partitioning is successfully employed to automatically segment the image into appropriate anatomical structures based on signal intensity in both 2D and 3D data. A novel localisation of the left ventricle blood pool is achieved using shape descriptors before segmentation of the outer wall of the left ventricle myocardium is accomplished using gradient information and prior knowledge.

To fully utilise all the data presented from a single patient scan, methods were investigated for the introduction of temporal information into the segmentation process. Temporal information is useful, as predictions of spacial deformation can be used to increase robustness segmentation. Level-set theory is introduced as a numerically stable method of evolving a surface in 3D based on intrinsic properties of the surface and external forces obtained from the image. In this thesis, a successful extension of Malladi and Sethians [86] formalisation for shape recovery is employed which incorporates a texture component and a probabilistic model of previously segmented cardiac boundaries to avoid the surface spilling into other anatomical structures in the presence of low gradient. Employing the idea of a coupled level-set introduced by Zeng et al. [181], the inner and outer wall of the left ventricle are segmented simultaneously using coupled surfaces that interacts using a coupling function.

Exploiting the Eulerian formalisation of the level set, the extension to complete 4D segmentation introduces a parametric model of left ventricle deformation over a cardiac cycle to aid the segmentation. This model is then iteratively refined using an optimisation algorithm. The model is re-parameterised for each position on the grid within a narrow-band of the evolving surface or surfaces, giving it a non-rigid deformation to take account of areas of the cardiac muscle that do not demonstrate significant spatial deformation, for example in the case of diseased tissue.

Each of the methods introduced have been tested on synthetic images and real patient scans. Performance is evaluated by comparing results against expertly* assisted manual delineation of the cardiac contours. In the next section, the strategies employed and advantages of these methods over existing methods commonly used in the cardiac segmentation will be discussed.

6.2 Contributions

In assessing the research conducted in this project toward the goal of cardiac image analysis, it is clear that a number of significant contributions have been made as well as other minor contributions. One of the objectives of the project is to integrate all the data available from a single patient scan into the segmentation process in an appropriate and functional manner. A full characterisation is attained at each stage in the development of the hypothesis. The major contributions of this thesis are as follows:

- A novel method for the segmentation of 4D information using prior knowledge about temporal deformation is introduced in a level-set framework. This prior knowledge is then iteratively optimised through the segmentation process.
- Produced a novel formulation for a coupled segmentation scheme, in a level-set framework, using a probabilistic model which segments the myocardium of the left ventricle.
- Developed an improved methodology for cardiac image analysis using statistical data partitioning.

*The validation was performed by Dr. John Murray, Cardiologist, Mater Misericordiae Hospital, Dublin, Ireland.

- Formed a gradient based segmentation of the left ventricle muscle outer wall using prior knowledge.
- Performed a full characterisation of advanced data filtering algorithms in medical images.

There were also some minor contributions resulting from this research:

- Developed a novel seed generator for initialising seed positions for automatic data partitioning algorithms based on histogram analysis.
- Applied the level-set segmentation technique in CT data for the extraction of polyp morphology for colon cancer detection.
- Designed a basic graphical user interface, see Figure 6.1, for visualising data and patient information and a separate back-end repository of algorithms for medical data processing and analysis.

6.3 Discussion

At the start of this thesis, a brief overview of two opposite approaches to segmentation were outlined, bottom-up and top-down approaches. Some examples of how both methodologies have been applied in the field of medical imaging were also given. From this initial discussion, a number of advantages and disadvantages for both were provided.

Firstly, bottom-up approaches offer a general solution without making any assumptions about the data being processed or about the final solution to the problem. Spatial information may be used locally about a small neighbourhood (edge-detectors, region-growing) or may not be used at all (thresholding, signal intensity clustering). These methods perform effectively in well defined data such as in CT data or in data after performing advanced filtering but in the case of poor or noisy data, bottom-up techniques can produce unpredictable and uncontrollable results.

On the other hand, top-down approaches such as template matching, ASMs and AAMs perform the segmentation using purely information that has been used in a training process. For example, template matching uses information in a global sense to minimise the error in order to find the most appropriate fit between the image data and the template. Such methods have demonstrated

robust localisation in the presence of low SNR [150]. Incorporating other metrics into the model such as texture has been shown to minimise the model to target differences. Other methods have been developed for minimising the model to target differences involving alternative approaches to model construction such as PCA, where the principal components of the models variation are utilised in the deformation process. Top-down approaches are limited in their use where the structure of the target object varies significantly from those contained in the training data. For example in cardiac imaging, the general models employed by ASMs\AAMs that are obtained from training sets are limited in their application for accurate segmentation to the variety of heart shapes. Abnormalities in the image data can indicate disease. Model based approaches approximate to the closest plausible instance shape from the training set Point Distribution Model (PDM), but this may not be sufficiently accurate. Also, AAMs cannot deal well with the changes in texture.

Also included in Chapter 2 was a note on how to combine both top-down and bottom-up approaches in order to obtain a more appropriate solution. In this thesis, methods of effectively combining prior information and local image properties are investigated. Following the removal of unwanted noise from the image, the process of partitioning the structural features within the image is achieved using a statistical based clustering algorithm. Localisation of the left ventricle cavity is achieved using prior knowledge about the shape of the structure based on prior knowledge. Once the left ventricle cavity has been successfully localised and extracted, a novel method for the outer wall of the left ventricle cavity is pursued. Approximate knowledge about the myocardium thickness is obtained from the distance between the left and right blood pools, assuming that the right ventricle blood pool is close to the left ventricle blood pool and the separating muscle (interventricular septum) approximates the thickness of the myocardium around the left ventricle. This knowledge is used when extracting local gradient information that may form part of the epi-cardium boundary. By linking appropriate edges together, segments are produced. These segments can then be eliminated with respect to orientation. Where gradient information is lacking, a top-down approach is adopted whereby missing segments are inserted by means of a probabilistic model of previously segmented images.

To further advance the concept of using the top-down approaches to guide bottom-up approaches, the idea of an evolving surface is introduced in Chapter 5. In a level-set framework, prior knowledge about the distance between

the epi- and endo-cardium boundaries as well as a probabilistic model of previously segmented images were used to influence a coupled level set deformation. The probabilistic model is introduced as a cost function, penalising growth away from model instances. Unlike the variational framework proposed by Paragios [110, 111] that uses both probabilistic measures for signal intensity obtained from an expectation-maximisation algorithm and prior shape information encoded in a level-set framework, our method uses high gradient information as the predominant stopping term and can therefore be applicable in situations where variations in grayscale are encountered.

Extending this methodology to $3D + t$ space, the aim was to remove the confidence attributed to the prior knowledge of the anatomical shape of the left ventricle, as it is known to contain a high degree of variation especially in abnormal or unhealthy specimens. It is proposed to model the temporal motion of the heart, as temporal motion in healthy and unhealthy hearts maintain the systole and diastole phases. Using this characteristic, a temporal model is constructed and iteratively updated to guide the local deformation of the level-set algorithm. This method of top-down knowledge about temporal deformation, optimised in order to influence the bottom-up approach gives a significant step towards a robust, elegant and complete solution to the $3D + t$ segmentation problem. The idea of encoding the temporal motion in a parametric model can be applied in different scenarios. In the next section, some possible situations are discussed.

6.4 Further Work

While this work addresses a specific research question, there is further work which can be undertaken in a broader sense as a result of the ideas put forward. In this section a number of areas are proposed which warrant further investigation.

Initialisation of the level-set algorithm could be improved. Fast-marching algorithm does not take curvature terms into its evolution. Further advancement of the fast marching method can improve the initialisation of the temporal model parameters used in the 4D case described in this thesis.

From a theoretical aspect, the level-set formulation is robust and numerically stable. Further work may involve a more involved formalisation of the level-set evolution in order to incorporate the 4D information. Further advances may in-

clude the extension of the temporal model using non-linear approximations with more advanced functions. The temporal motion model may also be encoded in PCA or other method to reduce the dimensionality of the model. These models may be derived using data from prior information based on expertly segmented cardiac images.

In this thesis, a novel method for applying top-down information in a bottom-up approach to segmentation is achieved. The application chosen to demonstrate the ideas proposed in this thesis are in multi-dimensional cardiac data. Application of these ideas in different areas would warrant further investigation. The work may be transferred to perform segmentation in the right ventricle or the measurement of valve regurgitation may also be achieved. Modelling temporal characteristics using more advanced functions can be utilised outside of the medical domain. Measuring growth in plants may be one application of this technique [11].

6.5 Concluding Remarks

In this work, a thorough investigation into multidimensional image analysis of cardiac data in MRI has been performed which was the main contribution of this research. The primary steps involved advancing the framework from a purely bottom-up approach based on statistical analysis to a more involved approach based on surface propagation using increasing dimensional data and incorporating top-down information to aid the segmentation. This is achieved in a novel and intuitive fashion. Optimisation of the algorithms performance from a computational expense point of view was performed but advanced developments in this area was not one of the main goals for this project. Additional research has been investigated outside the topic and contribute to minor advances in research. These are explained in detail in Appendix A and B.

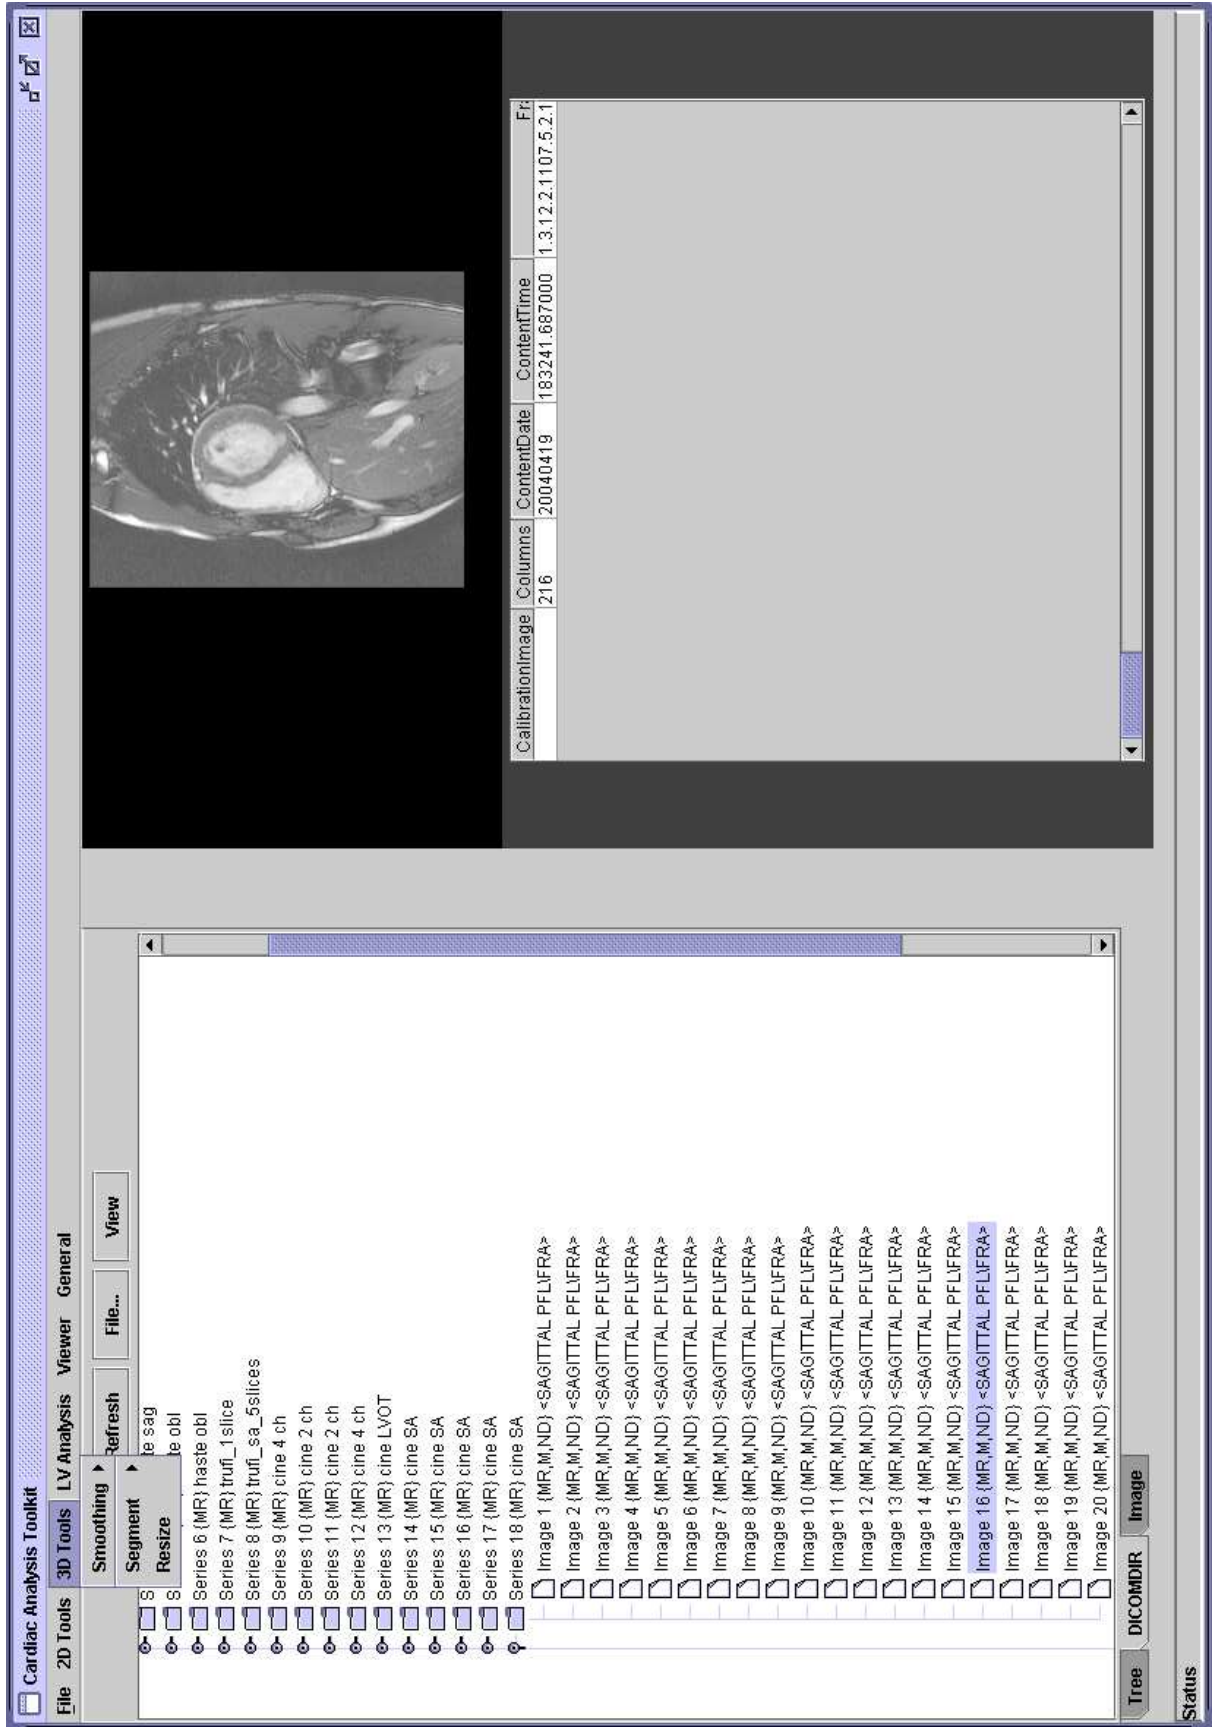


Figure 6.1: Screen capture for the graphical interface.

Appendices

Appendix A

Application of the Expectation-Maximisation Algorithm to Medical Images

This appendix details the Expectation-Maximisation (EM) for partitioning image using pixel intensity values. A novel approach for the initialisation of parameters is detailed using analysis of the intensity histogram of the image.

The application of the EM algorithm for the partitioning of medical images into anatomical structures has been documented, particularly in brain segmentation in MRI [47]. The EM algorithm shows robust and repeatable performance in the segmentations of heart, brain and abdominal images. The EM algorithm is locally convergent [174, 40, 13] so we have introduced an automatic seeding method that uses local maxima in the intensity histogram. In this appendix the novel initialisation of the EM algorithm is investigated and analysis is presented. Also results against manual initialisation and apply the algorithm to some common medical image processing tasks are demonstrated.

A.1 EM Algorithm

The EM algorithm [40, 14] attempts to classify data using a soft membership function as a weighted sum of a number of Gaussian distributions called a Gaussian Mixture Model (GMM). The generation of this GMM is achieved through an EM




Figure A.1: Image intensity histogram overlaid with an illustration of the associated Gaussian Mixture model.

technique, which aims to find the maximum likelihood estimate for an underlying distribution from a given data set when the data is incomplete. Its advantage over the k -means clustering technique [42] is its ability to provide a statistical model of the data and its capability of handling the associated uncertainties. Consider the general case of a d -dimensional random variable $X = [x_1, x_2, x_3, \dots, x_d]^T$ and suppose it follows a k -component finite mixture distribution. Its probability density function (pdf) could be written as,

$$p(x|\theta) = \sum_{m=1}^k \alpha_m p(x|\theta_m) \quad (\text{A.1})$$

where k is the number of mixtures, α_m is the mixing parameter for each of the Gaussian's in the GMM and $p(x|\theta_m)$ is the probability that variable x belongs to class θ_m and is defined in Equation A.2.

$$p(x|\theta_m(\mu_m, \sigma_m)) = \frac{1}{\sigma\sqrt{2\pi}} e^{-\frac{(x-\mu)^2}{2\sigma^2}} \quad (\text{A.2})$$

where $\theta_m = \{\mu_m, \sigma_m\}$ are the Gaussian's parameters. This can be displayed graphically in Figure A.1. The value of α_m is defined as,

$$\alpha_m \geq 0, \text{ and } \sum_{m=1}^k \alpha = 1 \quad (\text{A.3})$$

The algorithm is built on an iterative scheme and consists of two steps. The first, the E-step, calculates the expected log-likelihood function for the complete data, defined by Q using the estimates for the parameters $\hat{\theta}(t)$. X defines the input data and Y defines the output classified data.

$$Q(\theta, \hat{\theta}(t)) \equiv E[\log p(X, Y|\theta)|X, \hat{\theta}(t)] \quad (\text{A.4})$$

The second, M-step, uses the maximized values of this result to generate the next set of parameters.

$$\hat{\theta}(t+1) = \arg \max_{\theta} Q(\theta, \hat{\theta}(t)) \quad (\text{A.5})$$

The algorithm iterates between (A.4) and (A.5) until convergence is reached. It is important to note that local convergence of the EM algorithm is assured since $\hat{\theta}$ is smaller at each iteration [174, 40, 13].

The updates for the parameters for the GMM are the mixture values α_m and the Gaussian's parameters $\theta_m = \{\mu_m, \sigma_m\}$. These can be calculated from Equations A.6, A.7 and A.8.

$$\alpha_m^{new} = \frac{1}{N} \sum_{m=1}^k p(m|x_i, \hat{\theta}(t)) \quad (\text{A.6})$$

$$\mu_m^{new} = \frac{\sum_{m=1}^k x_i p(m|x_i, \hat{\theta})}{\sum_{m=1}^k p(m|x_i, \hat{\theta})} \quad (\text{A.7})$$

$$\sigma_m^{new} = \frac{\sum_{m=1}^k p(m|x_i, \hat{\theta})(x_i - \mu_m^{new})(x_i - \mu_m^{new})^T}{\sum_{m=1}^k p(m|x_i, \hat{\theta})} \quad (\text{A.8})$$

A.1.1 Seed Generation

To address the initialisation step a novel approach to collect relevant seed points for cluster centers based on histogram analysis is developed. A histogram of the image data is constructed, n_j , where n is the number of pixels contained in the bin with value j . This histogram is then divided into M evenly distributed bins. This value M is manually set, typically to a higher number than the number of perceived relevant regions in the image. For the images shown in this appendix, the value of M was set experimentally to 25. From each bin, the highest peak in the histogram is assigned to a seed center, C_m .

$$C_m = \arg \max_j (n_j) \quad (\text{A.9})$$

These M seed centers are then clustered together using their closeness in the grayscale space until the desired number of seeds, k , is reached. The clustering is an iterative process where clusters are joined together by evaluating the Euclidean distance between the cluster centers.

A.2 Results

The described scheme was applied to gated MRI short-axis images of the heart, MRI coronal brain slices and a section from a whole body MRI showing the lower abdomen. The results are compared against those obtained when the cluster means and variances are manually extracted from the image. From Figure A.2 and Table A.1, it is clear that using the automatic seed initialisation gives a better distribution of initial seeds across the data. Table A.1 presents the manually selected means of the Gaussians and automatically selected means using the method described above. Also, the Gaussian means following the EM algorithm has been applied are presented.

To evaluate the performance of the described algorithm, the EM segmentation algorithm is applied to each of the MRI datasets. As mentioned previously, the algorithm is locally convergent and therefore initialisation of the algorithm is crucial to the final solution. A comparison is made between the results obtained using the automatically seeding process and the results obtained when the initial seeds for the EM segmentation are chosen manually. To achieve this, areas are selected in each of the images that attempt to represent the most significant regions. This is objective and related to the purpose of the segmentation but the overriding motivation is to pick regions that are clinically significant and also have a high degree of variation between regions. In each of the images given, 6 regions were manually selected. In these selected regions the mean pixel intensity values and the variance of the pixel intensity values are calculated. These manually selected values are used as the initial θ_m 's, where $1 \leq m \leq 6$ in the EM algorithm, the mixing parameters α_m were each set to $\frac{1}{m}$.

Figure A.2 illustrates the strategy applied to short axis images from a cardiac MRI study. The areas manually selected are shown in Figure A.2 (b) and the




Figure A.2: Figures show the short axis view of cardiac MRI (*a*) shows the original image (*b*) indicates the manually selected areas (*c*) represents the results after applying the EM using the manually picked initialisation and (*d*) is the result after applying the automatic seed picking.

resultant segmentation after applying the EM segmentation using these initial parameters is shown in Figure A.2 (*c*). The final Figure A.2 (*d*) shows appropriate results after the automatic parameter selection, in particular the results show a better distribution within the grayscale distribution of the analysed image. Figure A.3 shows a coronal slice from a T1-weighted head MRI. Again the automatic segmentation method performs well in differentiating the white matter from the gray matter. Figure A.4 shows a coronal slice from an abdominal section of a full body MRI.

Table A.1: Changes in cluster means in the Cardiac data

	Manual μ 's	Manual μ 's after EM	Automatic μ 's	Automatic μ 's after EM
$\mu(0)$	57.31914	55.2806	57	31.33457
$\mu(1)$	125.366	112.0961	137	125.284
$\mu(2)$	194.0437	151.1044	167	171.6872
$\mu(3)$	19.84193	16.74244	12	17.75531
$\mu(4)$	225.1899	112.8278	255	254.2933
$\mu(5)$	28.87568	28.43651	92	79.93145

It is clear from Tables A.1, A.2 and A.3 that the described automatic seed picking algorithm demonstrates better performance when compared to the manual selection technique. This is evident from the lower differences between initialised seeds and the final values after optimisation through the EM algorithm.

Most medical images obtained from MRI are 3D and in some cases 4D, but because the described algorithm works on the data histogram (hence, intensity values) and is not dependent on spatial position, therefore as a result the algorithm can be applied equally successfully to any dimensioned data. This is illustrated in Figure A.5 where the algorithm is successfully applied in 3D MRI images. This aspect is examined further in Section A.3 where the results are used in conjunction with a diffusion based filtering [54, 115] to extract some clinically relevant regions from the images.

It is worth noting that statistical classification of pixels is a more appropriate way to segment medical images as the standard region growing technique will fail to produce appropriate results in images that exhibit a low signal to noise ratio (SNR). Also, medical images generally show good separation between significant regions as this is one of the aims in the acquisition. This is application dependent some common medical applications are investigated in the following section.

A.3 Applications in Medical Imaging

One of the key indicators of cardiac health is left ventricle *ejection fraction*, a measure of the volume of blood pumped from the left ventricle with each heart-




Figure A.3: Figures show an coronal slice from a brain MRI (a) shows the original image (b) indicates the manually selected areas (c) represents the results after applying the EM using the manually picked initialisation and (d) is the result after applying the automatic seed picking.

beat[48]. Cardiac cine MRI is a standard procedure where 3D volume images are acquired at gated temporal positions through the cardiac pumping cycle. Such images are frequently taken using gradient echo imaging, which exhibits a relatively high differentiation between the blood and the myocardium. Figure A.6 shows the end-diastole segmented left ventricle blood-pool after the application of the EM algorithm to identify the left ventricle cavity. Figure A.6(e) is a rendered volume of the blood pool inside the cavity of the left ventricle when the muscle is at its end-diastole phase.

Table A.2: Changes in cluster means in the brain data

	Manual μ 's	Manual μ 's after EM	Automatic μ 's	Automatic μ 's after EM
$\mu(0)$	164.6	123.922	116	117.66
$\mu(1)$	131.18	120.03	96	97.8356
$\mu(2)$	2.3	2.03	13	2.07
$\mu(3)$	66.59	33.01	44	27.48
$\mu(4)$	90.1	94.49	73	70.836
$\mu(5)$	164.21	194.81	153	140.6223




Figure A.4: Figures show a coronal slice from a section of a full body MRI (a) shows the original image (b) indicates the manually selected areas (c) represents the results after applying the EM using the manually picked initialisation and (d) is the result after applying the automatic seed picking.

The classification of brain MRI's white matter, gray matter, cerebrospinal fluid and in some cases lesions, is a fundamental first step for surgical planning, radiotherapy planning and the identification of brain disease [180]. Illustrated in Figure A.7 is a segmentation of white matter of the brain.

The accurate measurement of body fat from whole-body MRI images is becoming an increasingly important metric as high body fat level is recognised to play a

Table A.3: Changes in cluster means in the whole body data

	Manual μ 's	Manual μ 's after EM	Automatic μ 's	Automatic μ 's after EM
$\mu(0)$	170.92	169.4365	183	178.41
$\mu(1)$	42.29	44.45	52	50.484
$\mu(2)$	3.84	4.177	5	4.27
$\mu(3)$	123.61	118.868	151	153.720
$\mu(4)$	95.35	82.99	124	121.496
$\mu(5)$	57.2	55.897	92	85.687

significant role in a variety of serious health problems [18]. MRI is the modality of choice due to its repeatability and high spatial resolution. Figure A.8 illustrates the results from one section of a whole-body MRI dataset where the fat tissue has being segmented out of the volume.

The developed method shows appropriate results with respect to the gray scale values for all datasets. From these results we can conclude that this approach offers robust, reproducible and accurate estimation of the initial parameters for the EM algorithm and the segmentation scheme described is capable of providing useful clinical measurements when applied to a large range of medical datasets.

Journal Publications

Michael Lynch, Dana Ilea, Kevin Robinson, Ovidiu Ghita, Paul F. Whelan. **Automatic seed initialisation for the Expectation-Maximisation algorithm and its application in 3D medical imaging**. Journal of Medical Engineering and Technology. (Accepted – awaiting publication)

Conference Publication

Michael Lynch, Ovidiu Ghita and Paul F. Whelan (2005), **Automatic Seed Picking Algorithm for Region-Based Segmentation of Cardiac MRI Images**, European Society of Cardiac Imaging, ESCR 2005, October 2005, Zürich, Switzerland.




Figure A.5: 3D space partitioning using EM: Images show a single slice of a 3D dataset from (a) the original volume, (b) after segmentation with the EM algorithm.

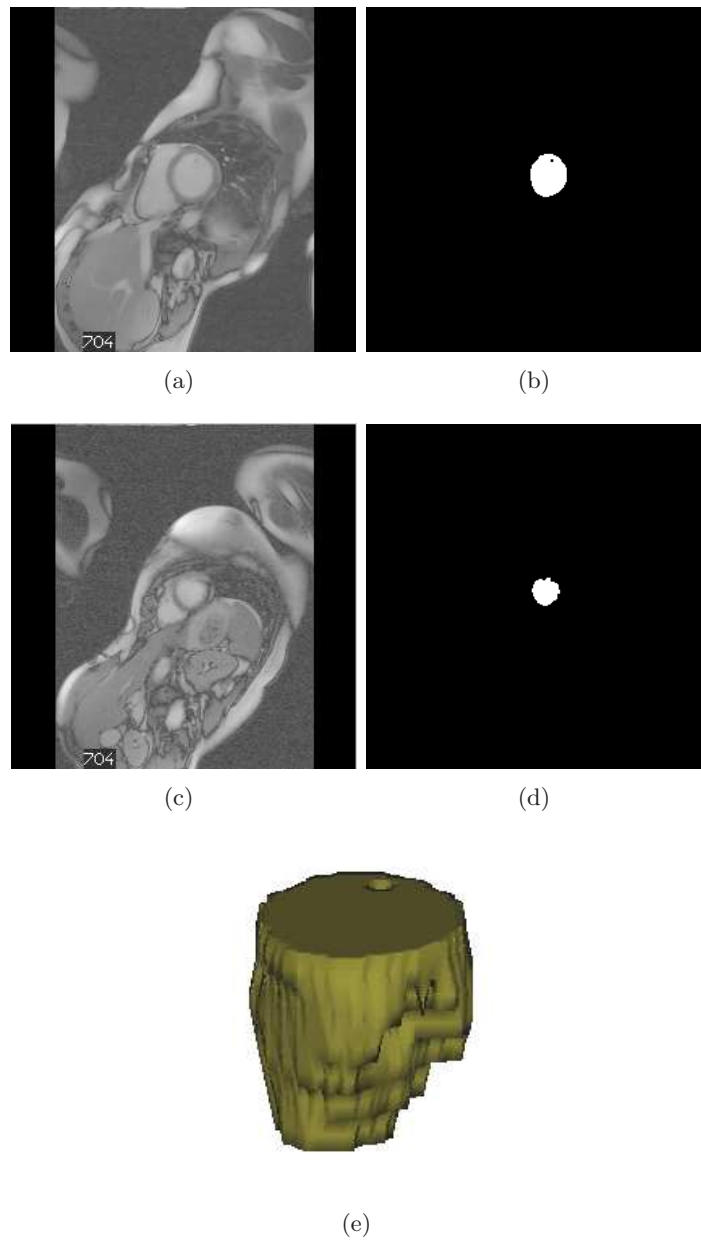


Figure A.6: Images show slices 1 ((a) and (b)) and 4 ((c) and (d)) from the original volume (left) and with left ventricle blood cavity segmented (right) and (e) shows the rendered volume of the segmentation.




Figure A.7: Images show slices 1 ((a) and (b)) and 14 ((c) and (d)) from the original volume (left) and with segmented white matter (right) and (e) shows the rendered volume of the segmentation.




Figure A.8: Images show slices 2 ((a) and (b)) and 6 ((c) and (d)) from the original volume (left) and with body fat segmented (right) and (e) shows the rendered volume of the segmentation.

Appendix B

Level-set Segmentation for Candidate Polyp extraction in CTC

The extraction of candidate polyps from Computer Tomography Colonography (CTC) is a primary and important step in candidate polyp classification, where polyps are a precursor to colon cancer. Such a classification step is necessary due to the high frequency of false positive polyp detections which are apparent in previous computer aided diagnostic techniques. Previous work in this area uses curvature constraints on candidate polyps to establish morphology [176]. This type of classification encounters difficulty when determining folds, a naturally occurring instance in the colonography exam. In this work, we have used surface normal intersection to determine possible polyp candidates, we then proceed to segment the polyp using a level set curve evolution algorithm to extract an accurate segmentation of the polyp features. Results are presented using point to surface error and the reduction in false positives after the extracted surfaces were classified using a statistical classifier.

Much of the previous work in polyp extraction uses local curvature and shape constraints to determine polyp candidates and to establish morphology [154, 178, 72]. This type of classification encounters difficulty when determining folds, a naturally occurring instance in the colonography exam. Yao et al. [176] proposed a segmentation of method which used a knowledge guided deformable model to extract the surface of the polyp and compared it to manual segmentation of

experts. The knowledge was provided by the curvature of the deformable model and the signal intensities of the pixels surrounding the polyp. The segmentation was performed in 2D and the 2D images were combined together to create the local 3D volume.

B.1 Convex Surface Extraction

Initially, the colon is segmented using a seeded 3D region growing algorithm that was applied to segment the air voxels, which assures the robust identification of the colon wall. In some situations the colon is collapsed due to either insufficient insufflation or residual water. In order to address this issue we have developed a novel colon segmentation algorithm that is able to correctly identify the colon segments using knowledge about their sizes and location within the body in all imaging conditions. After the identification of the colon wall, for each colon wall voxel the surface normal vector is calculated using the Hummel-Zucker operator [182]. The normal vectors sample the local orientation of the colonic surface and the suspicious candidate structures that may resemble polyps are extracted using a simple convexity analysis. In this regard, the colonic suspicious surfaces have convex properties and are determined using the 3D histogram and Gaussian distribution of the Hough points (full details about this developed algorithm can be found in [29]). This method is able to correctly identify all polyps above 3mm but it is worth nothing that this is achieved at a cost of high level of false positives. In order to reduce the level of false positives, the surface is extracted using a level-set method and the results are classified using a statistical morphological features.

B.2 Level-Set Initialisation. Fast-Marching Algorithm

As previously outlined in Chapter 5 formulation of the problem is conceptually simple. The evolving curve or front Γ , evolves as the zero level-set of a higher dimensional function ϕ . This function deforms with a force F that is dependent on both curvature of the front and external forces in the image. The force acts in the direction of the normal to the front.

$$\begin{aligned} \phi_t + F|\nabla\phi| &= 0 \\ \phi(x, y, t = 0) &= \textit{given} \end{aligned} \tag{B.1}$$

The implementation employed is a standard two step approach which includes




Figure B.1: Flow-chart of proposed algorithm

a fast-marching initial step to speed up the segmentation. Fast marching is a special case of the above equation where $F(x, y) > 0$. Let $T(x, y)$ be the time that the front Γ crosses the point (x, y) . The function $T(x, y)$ then satisfies the equation;

$$|\nabla T|F = 1 \quad (\text{B.2})$$

which simply says that the gradient of the arrival time is inversely proportional to the speed of the surface. The T function is evaluated using the diffusion and attraction to pixels within the front. The front grows out from its initial position to points with the smallest value of $T(x, y)$. The $T(x, y)$ function is then updated and continued until the front does not grow.

B.3 Level-Set Analysis

The theory behind level-set segmentation is largely based on work in partial differential equations and the propagation of fronts under intrinsic properties such as curvature [108, 133, 41, 74]. By extending the dimensionality of the problem to $N+1$, where N is the initial dimension of the problem, some advantageous prop-

erties can be exploited. Representing the boundary as the zero level set instance of a higher dimensional function ϕ , the effects of curvature can be easily incorporated. ϕ is represented by the continuous Lipschitz function $\phi(s, t = 0) = \pm d$, where d is the signed distance from position s to the initial interface Γ_0 (see Equation B.3). The distance is given a positive sign outside the initial boundary ($D \setminus \Omega$), a negative sign inside the boundary ($\Omega \setminus \partial\Omega$) and zero on the boundary ($\partial\Omega$).

$$\phi(s) = \begin{cases} -d & \forall s \in \Omega \setminus \partial\Omega \\ 0 & \forall s \in \partial\Omega \\ +d & \forall s \in R^n \setminus \Omega \end{cases} \quad (\text{B.3})$$

From this definition of ϕ , intrinsic properties of the front can be easily determined, like the normal $\vec{n} = \pm \frac{\nabla\phi}{|\nabla\phi|}$.

Since curvature of the polyp should be a pertinent factor in the segmentation evolution, particular emphasis is given to this measure. The mean curvature (H), is connected to the physical evolution of soap bubbles and the heat equation. While smooth, it may not necessarily be convex and can lead to singularities.

$$H = \nabla \cdot \frac{\nabla\phi}{|\nabla\phi|} \quad (\text{B.4})$$

Gaussian curvature (K), has also being used to model physical problems such as flame propagation. It has being shown that a convex curve evolves to a point under curvature evolution, but it can also be shown that evolution of non-convex surfaces can be unstable [7].

$$K = \frac{\nabla\phi^T \text{Adj}(\mathbf{H}(\phi))\nabla\phi}{|\nabla\phi|^2} \quad (\text{B.5})$$

where $\mathbf{H}(\phi)$ is the Hessian matrix of ϕ , and $\text{Adj}(\mathbf{H})$ is the adjoint of the matrix \mathbf{H} .

Due to the characteristic curvature features of polyps it is proposed to use Neskovic and Kimia's [106] measure of curvature, which involves both mean and Gaussian. In this approach, the direction of flow is obtained from the Mean curvature while the magnitude of the flow is dictated by the Gaussian curvature. This is appropriate as the Mean curvature alone can cause singularities and

extracts the strictly convex surface of the polyp candidate.

$$\kappa = \text{sign}(H)\sqrt{K + |K|} \quad (\text{B.6})$$

Using this value for κ , the level set is iteratively updated within a defined narrow band around the segmented boundary to increase efficiency. The following equation details the update parameters

$$\phi_{t+1} = \phi_t + k_I(1 - \epsilon\kappa)|\nabla\phi| + \beta\nabla I \cdot \nabla\phi \quad (\text{B.7})$$

where ϵ and β are user defined parameters (see Table B.1), κ is the curvature term defined in Equation B.6 and k_I is the gradient dependent speed term and is given by $\frac{1}{1+|\nabla I|}$. The third term, $\nabla I \cdot \nabla\phi$ represents the attractive force vector normal to the front. The level-set segmentation is performed in 3D.

Possible polyp candidate centres are calculated over the entire data set by calculating the normal vectors at each voxel on the lumen wall. Polyp candidates are defined as regions of high convexity, therefore the centres for possible polyp candidates are located at points that contain high concentration of normal intersections [29].

The level set is initialised at the polyp candidate centres and grows outwards until a boundary is encountered. The convex surface is maintained by placing a high influence on the curvature parameter (see Figure B.2). Once the level-set has converged or completed its iterations, the surface of the polyp candidate is taken as all boundary points that have an associated gradient. This ensures that just the lumen surface is extracted.




Figure B.2: Extracted polyp surface (dotted) using the levelset approach based on curvature.

B.4 Classifier

Once the true surface of the polyp candidates has been extracted, they are passed to a classifier to determine whether they are polyps or folds. The classifier is a statistical model of known polyps and folds and uses statistical features of the candidates morphology such as least squares ellipsoid fitting error, normalised distribution of the surface curvature and the Gaussian sphere radius [29]. These features are used to classify the candidate polyp surfaces into polyps or folds using a feature normalised nearest neighbour classification scheme [55]. The classifier was trained with 64 polyps and 354 folds that were selected as true positives by a radiologist.

B.5 Results

The segmentation algorithm described above was performed on 10 full CTC data set, converted to isotropic dimensions using cubic interpolation. Visual representations of the segmentation are shown in Figure B.3 and the extracted surface renderings are shown in Figure B.5. Table B.1 lists the user defined parameters used in the level-set algorithm. From this table it can be seen that curvature is given a large influence to maintain the convexity of the polyp candidate surface. The narrow bandwidth is given a small value of 10 to increase the efficiency of the update.

A classifier, trained on expertly categorised unseen data, is then used to determine whether the extracted surface is classified as polyp or non-polyp. Small folds in the colon lumen are the main cause of detecting a false positive. It can be clearly seen in Figure B.5 that fold surface is extracted is saddle shaped and thus can be easily classified using its shape characteristics.

Table B.2 shows the measured point-to-curve error between the automatic segmentation results against those found from a manual segmentation of the small number of polyp candidates. Indicated on the table are the average error, standard deviation of the error and the root-mean-square of the error. This error is measured in pixels where each pixel has sub-millimeter dimensions.

Table B.3 gives the results from 10 datasets (9 patients) containing 31 polyps. From the high number of polyp surface candidates, a relatively low number are detected. The results show a sensitivity of 100% for all polyps $>10\text{mm}$. Normally, in a clinical situation, polyps below 5mm have less clinical significance. One cause for our method missing smaller polyps, are their low curvature difference between the polyp and the colon wall, therefore some colon wall is taken into




Figure B.3: Images above show the segmentation of the convex polyp candidate. The bottom left image shows the segmentation of a fold.

Table B.1: Control parameters used in the level-set segmentation.

Control Parameters	Values
Fast-Marching Iterations	3
Level-set Iterations	10
Level-set ϵ	0.5
Level-set β	0.08
Level-set Narrow bandwidth	10

Table B.2: Point-to-curve errors between manually segmented data and our method.

Error	Average	Std. Dev.	RMS
	0.298	0.587	0.661

Table B.3: Performance Analysis for Polyp Classification. True positive (TP) and False Positive (FP).

Size	Detected	Missed
$\geq 10\text{mm}$	10	0
5–10mm	9	1
$\leq 5\text{mm}$	2	20

the candidate surface (see Figure B.6 and Figure B.4). The false positives per dataset was calculated to be 1.3, which compares favorably with figures reported to literature.




Figure B.4: Extracted polyp surface (dotted) for a small polyp, note the inclusion of healthy colon lumen.




Figure B.5: Images above show the polyp candidate renderings of the extracted surface. Figures (a)-(c) show correctly classified polyps, where Figures (d)-(f) show correctly classified folds.




Figure B.6: One of the $\leq 5\text{mm}$ polyps misclassified due to the inclusion of colon wall in the surface extraction.

Publications associated with this chapter

Conference Publication

Michael Lynch, Tarik Chowdhury, Ovidiu Ghita and Paul F. Whelan (2005), **Determining Candidate Polyp Morphology from CT Colonography using a Level-Set Method**, European Medical and Biological Engineering Conference EMBEC 2005, November 2005, Prague, Czech Republic.

Appendix C

Mathematical Background

C.1 LMS Circle

Using the Least Squares solution a circle is fitted around a collection of points, P_i , with images coordinates, (x_i, y_i) for $i = 1, 2 \dots N$.

A circle is defined by three parameters. These parameters are the coordinates of its centre (x_0, y_0) and its radius r . The equation of a circle can be written isolating these three parameters as follows:

$$\begin{pmatrix} 2x_i & 2y_i & 1 \end{pmatrix} \begin{pmatrix} x_0 \\ y_0 \\ r^2 - x_0^2 - y_0^2 \end{pmatrix} = \begin{pmatrix} x_i^2 + y_i^2 \end{pmatrix}$$

In order to find these three unknowns a linear least squares solution is obtained where:

$$A = \begin{pmatrix} 2x_1 & 2y_1 & 1 \\ 2x_2 & 2y_2 & 1 \\ 2x_3 & 2y_3 & 1 \\ \dots & \dots & \dots \\ 2x_N & 2y_N & 1 \end{pmatrix}, b = \begin{pmatrix} x_1^2 + y_1^2 \\ x_2^2 + y_2^2 \\ x_3^2 + y_3^2 \\ \dots \\ x_N^2 + y_N^2 \end{pmatrix}$$

The best fitting circle for the points P_i is the least squares solution to $[x_0 \quad y_0 \quad r^2 - x_0^2 - y_0^2]^T = (A^T A)^{-1} A^T b$ where $(A^T A)^{-1} A^T b$ can be written as:

$$\begin{pmatrix} 4 \sum x_i^2 & 4 \sum x_i y_i & 2 \sum x_i \\ 4 \sum x_i y_i & 4 \sum y_i^2 & 2 \sum y_i \\ 2 \sum x_i & 2 \sum y_i & N \end{pmatrix}^{-1} \begin{pmatrix} 2 \sum x_i^3 + 2 \sum x_i y_i^2 \\ 2 \sum y_i^3 + 2 \sum x_i^2 y_i \\ \sum x_i + \sum y_i^2 \end{pmatrix}$$

The error of this least squares solution can be calculated as the difference between the area of the fitted circle and the area of concentric circles passing through the data points with the equation $e_{circle} = \| A[x_0 \ y_0 \ r^2 - x_0^2 - y_0^2] - b \|$

C.2 LMS Ellipsoid

To determine the left ventricle cavity after the application of 3D clustering, the error between each segmented shape and a fitted ellipsoid is found. The radii of the ellipsoid are calculated using the eigenvalues of the covariance matrix from the lists of points that define the surface of the shape.

$$C = \begin{bmatrix} \sigma_x & \sigma_{xy} & \sigma_{xz} \\ \sigma_{xy} & \sigma_y & \sigma_{yz} \\ \sigma_{xz} & \sigma_{yz} & \sigma_z \end{bmatrix} = \begin{bmatrix} \sum_{n=0}^{N-1} \frac{(x-\bar{x})^2}{N} & \sum_{n=0}^{N-1} \frac{(x-\bar{x})(y-\bar{y})}{N} & \sum_{n=0}^{N-1} \frac{(x-\bar{x})(z-\bar{z})}{N} \\ \sum_{n=0}^{N-1} \frac{(x-\bar{x})(y-\bar{y})}{N} & \sum_{n=0}^{N-1} \frac{(y-\bar{y})^2}{N} & \sum_{n=0}^{N-1} \frac{(y-\bar{y})(z-\bar{z})}{N} \\ \sum_{n=0}^{N-1} \frac{(x-\bar{x})(z-\bar{z})}{N} & \sum_{n=0}^{N-1} \frac{(y-\bar{y})(z-\bar{z})}{N} & \sum_{n=0}^{N-1} \frac{(z-\bar{z})^2}{N} \end{bmatrix} \quad (C.1)$$

Based on work by Pearson, principle component analysis (PCA) chooses the first ellipsoid axis as the line that goes through the centroid, but also minimizes the square of the distance of each point to that line, see figure C.1. The line is a correlation of the points along the data's principle axis. Equivalently, the line goes through the maximum variation in the data.

The second PCA axis also must go through the centroid, and also goes through the maximum variation in the data, but with a certain constraint: It must be completely uncorrelated (i.e. at right angles, or 'orthogonal') to PCA axis 1. The ellipsoid is an extension of this PCA to 3D finding the three principal axes.

C.3 Splines

A closed natural cubic spline is fitted around the points on the epi-cardium [144]. The spline is used to close the epi-cardium contour by connecting all the points




Figure C.1: The two principle axes of a two dimensional data set are plotted and scaled according to the amount of variation that each axis explains.

on the curve in a smooth way.

Splines are piece-wise polynomials of degree n ($n = 3$ in the case of cubic splines) with the pieces smoothly joined together. The joining points of the polynomial pieces are called control points which need not be evenly spaced. These control points are defined as a collection of points P_i where $i = 1, 2, 3 \dots N$ and N is the number of points. It works by fitting a cubic curve between each pair of points in the collection. Smoothness of the curve is maintained by forcing the first and second derivative of the end point of one curve to equal the start of the next curve. This is achieved by solving a system of simultaneous equations. The equation is illustrated below:

$$f_i(x) = a_i + b_i u + c_i u^2 + d_i u^3$$

$$0 \leq u \leq 1$$

$$1 \leq i \leq n$$

Where i is the amount of points on the curve and u is the number of steps in between each point. The coefficients of the cubic equation are,

$$a_i = x_n$$

$$b_i = \frac{dx_n}{dP}$$

$$c_i = 3(x_{n+1} - x_n) - 2\frac{dx_n}{dP} - \frac{dx_{n+1}}{dP}$$

$$d_i = 2(x_n - x_{n+1}) + \frac{dx_n}{dP} + \frac{dx_{n+1}}{dP}$$

The derivatives used in to smooth the curve are computed as follows:

$$\begin{pmatrix} D[0] \\ D[1] \\ \cdot \\ \cdot \\ \cdot \\ D[n] \end{pmatrix} = \begin{pmatrix} 4 & 1 & & & & 1 \\ 1 & 4 & 1 & & & \\ & 1 & 4 & 1 & & \\ & & & \dots & & \\ & & & & 1 & 4 & 1 \\ 1 & & & & 1 & 4 \end{pmatrix}^{-1} \begin{pmatrix} 3(x_1 - x_n) \\ 3(x_2 - x_0) \\ \cdot \\ \cdot \\ 3(x_n - x_{n-2}) \\ 3(x_0 - x_{n-1}) \end{pmatrix}$$

Bibliography

- [1] R. Wegenkittl A. Neubauer. A Skeleton-Based Inflation Model for Myocardium Segmentation. In *Proceedings of the 16th International Conference on Vision Interface*, Halifax, Canada, June 2003. [cited at p. 79]
- [2] D. Adalsteinsson and J. A. Sethian. A fast level set method for propagating interfaces. *Journal of computational Physics*, 118(2):269–277, May 1995. [cited at p. 25, 95]
- [3] M. B. Al-Daoud and S. A. Roberts. New methods for the initialisation of clusters. *Pattern Recognition Letters*, 17(5):451–455, 1996. [cited at p. 54]
- [4] A. Amini, T. Weymonth, and R. Jain. Using dynamic programming for solving variational problems in vision. *IEEE Transactions in Pattern Analysis and Machine Intelligence*, 12(9):8555–8567, 1990. [cited at p. 25, 78]
- [5] E. D. Angelini, Y. Jin, and A. F. Laine. *Handbook of Medical Image Analysis: Advanced Segmentation and Registration Models*, chapter State-of-the-Art of Levelset Methods in Segmentation and Registration of Medical Imaging Modalities. Kluwer Academic Publishers, 2004. [cited at p. 86]
- [6] E. D. Angelini, R. Otsuka, S. Homma, and A. F. Laine. Comparison of ventricular geometry for two real time 3d ultrasound machines with three dimensional level set. In *Proceedings of the IEEE International Symposium on Biomedical Imaging*, pages 1323–1326, Arlington, VA, USA, 2004. ISBI. [cited at p. 25, 93]
- [7] E. Angelopoulou and L. B. Wolff. Sign of gaussian curvature from curve orientation in photometric space. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 20(10):1056–1066, October 1998. [cited at p. 144]
- [8] L. Axel. Cardiac MRI. In *Biomedical Imaging: Macro to Nano, 2004. IEEE International Symposium on*, pages 1212–1214, 2004. [cited at p. 27]
- [9] E. Bardinet, L. Cohen, and N. Ayache. Tracking and motion analysis of the left ventricle with deformable superquadrics. *Medical Image Analysis*, 1(2):129–149, 1996. [cited at p. 103]

- [10] J. L. Barron. Experience with 3d optical flow on gated mri cardiac datasets. In *Proceedings in 1st Canadian Conference on Computer and Robot Vision*, pages 370–377, May 2004. [cited at p. 104]
- [11] J. L. Barron and L. Liptay. Measuring 3d plant growth using optical flow. *BioImaging*, 5:82–86, 1997. [cited at p. 123]
- [12] R. H. Bartels, J. C. Beatty, , and B. A Barsky. *An Introduction to Splines for Use in Computer Graphics and Geometric Modelling*. Morgan Kaufmann, San Francisco, CA, 1998. Splines in chapter 'Hermite and Cubic Spline Interpolation', pages 12-13. [cited at p. 64]
- [13] J.A. Bilmes. A gentle tutorial of the EM algorithm and its application to parameter estimation for gaussian mixture and hidden markov models. Technical Report TR-97-021, Berkeley, CA, 1998. [cited at p. 109, 127, 129]
- [14] C. M. Bishop. *Neural Networks for Pattern Recognition*. Oxford University Press, 1995. [cited at p. 23, 127]
- [15] J.M Bland and D.G Altman. Stastical methods for assessing agreement between two methods of clinical measurements. *Lancet*, 1(8476):307–310, 1986. [cited at p. 66]
- [16] E. Borenstein, E. Sharon, and S. Ullman. Combining top-down and bottom-up segmentation. In *Proceedings IEEE workshop on Perceptual Organization in Computer Vision, IEEE Conference on Computer Vision and Pattern Recognition*, Washington, DC, June 2004. CVPR. [cited at p. 25]
- [17] J.G. Bosch, S.C. Mitchell, B.P.F. Lelieveldt, F. Nijland, O. Kamp, M. Sonka, and J.H.C. Reiber. Automatic segmentation of echocardiographic sequences by active appearance motion models. *IEEE Transactions on Medical Imaging*, 21(11), 2002. [cited at p. 24, 84]
- [18] D. Brennan, P.F. Whelan, K. Robinson, O. Ghita, R. Sadleir, J.O' Brien, and S. Eustace. Rapid automated measurement of body fat tissue distribution from whole body mri. *American Journal of Roentgenology*, 185(1), 2005. [cited at p. 135]
- [19] J. Cai, J. C. Chu, D. Recine, M. Sharma, C. Nguyen, R. Rodebaugh, Saxena A., and A. Ali. Ct and pet lung image registration and fusion in radiotherapy treatment planning using the chamfer-matching method. *International Journal of Radiation Oncology Biology Physics*, 43(4):883–891, 1999. [cited at p. 26]
- [20] J. Canny. A Computational Approach to Edge Detection. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 8(6):679–698, November 1986. [cited at p. 22, 37, 61, 64]
- [21] V. Caselles, F. Catté, T. Coll, and F. Dibos. A geometric model for active contours in image processing. *Numerische Mathematik*, 66(1):1–31, October 1993. [cited at p. 86, 90]

- [22] V. Caselles, R. Kimmel, and G. Sapiro. Geodesic active contours. *Int. J. Comput. Vision*, 22(1):61–79, 1997. [cited at p. 77, 86, 89]
- [23] V. Casselles, R. Kimmel, and G. Sapiro. Geodesic active contours. In *Proc. Fifth IEEE International Conference on Computer Vision*, pages 694–699, Cambridge, USA, June 1995. [cited at p. 86]
- [24] M. D. Cerqueira, N. J. Weissman, V. Dilsizian, A. K. Jacobs, S. Kaul, W. K. Laskey, J. A Pennell, D. J. and. Rumberger, T. Ryan, and M. S. Verani. Standardized Myocardial Segmentation and Nomenclature for Tomographic Imaging of the Heart. *American Heart Association Writing Group on Myocardial Segmentation and Registration for Cardiac Imaging*, 105:539–542, 2002. [cited at p. 9, 13]
- [25] A. Chakraborty, L. Staib, and J. S. Duncan. Deformable boundary finding influenced by region homogeneity. In *Proc. Conf. on Computer Vision and Pattern Recognition*, Los Alamitos, CA, 1994. IEEE Computer Society Press. [cited at p. 25, 79]
- [26] A. Chakraborty, L. H. Staib, and J. S. Duncan. Deformable boundary finding in medical images by integrating gradient and region information. *IEEE Transactions in Medical Imaging*, 15(6):859–870, 1996. [cited at p. 25, 79]
- [27] T. Chan and L. Vese. Active contours without edges. *IEEE Trans. Image Processing*, 10(2):266–277, February 2001. [cited at p. 92]
- [28] K. Chen. A Feature Preserving Adaptive Smoothing Method for Early Vision. Technical report, National Laboratory of Machine Perception and The Center for Information Science, Peking University, Beijing, China, 1999. [cited at p. 28, 34]
- [29] T. A. Chowdhury, O. Ghita, and P. F. Whelan. A statistical approach for robust polyp detection in ct colonography. In *27th Annual International Conference of the IEEE Engineering in Medicine and Biology Society*, Shanghai, China, 2005. [cited at p. 142, 145, 146]
- [30] L. P. Clarke, R. P. Velthuizen, M. A. Camacho, J. J. Heine, M. Vaidyanathan, L. O. Hall, R. W. Thatcher, and M. L. Silbiger. Mri Segmentation: Methods and Applications. *Magnetic Resonance Imaging*, 13(3), 1995. [cited at p. 53, 75]
- [31] L. D. Cohen. Note on active contour models and balloons. *CVGIP: Image Understanding*, 53(2):211–218, March 1991. [cited at p. 77]
- [32] L. D. Cohen and R. Kimmel. Global minimum for active contour models: A minimal path approach. *International Journal of Computer Vision*, 24(1):57–78, August 1997. [cited at p. 93]

- [33] D. Comaniciu and P. Meer. Mean shift: A robust approach toward feature space analysis. *IEEE Trans. on Pattern Analysis Machine Intelligence*, 24(5):603–619, 2002. [cited at p. 28]
- [34] T. Cootes, A. Hill, C. Taylor, and J. Haslam. The use of active shape models for locating structures in medical imaging. *Image and Vision Computing*, 12(6):355–366, 1994. [cited at p. 24, 82]
- [35] T. F. Cootes, G. J. Edwards, and C. J. Taylor. Active appearance models. *Lecture Notes in Computer Science*, 1407:484–498, 1998. [cited at p. 83]
- [36] T. F. Cootes, C. J. Taylor, D. H. Cooper, and J. Graham. Active shape models - their training and application. *Computer Vision and Image Understanding*, 61(1):38–59, January 1995. [cited at p. 80]
- [37] C. Corsi, G. Saracino, A. Sarti, and C. Lamberti. Left ventricular volume estimation for real-time three-dimensional echocardiography. *IEEE Transactions on Medical Imaging*, 21(9):1202–1208, September 2002. [cited at p. 95]
- [38] K. K. Delibasis, N. Mouravliansky, G. K. Matsopoulos, K. S. Nikita, and A. Marsh. MR functional cardiac imaging: Segmentation, measurement and WWW based visualisation of 4D data. *Future Generation Computer Systems*, 15(2):185–193, March 1999. [cited at p. 23, 55]
- [39] H. Delingette, M. Hébert, and K. Ikeuchi. Shape representation and image segmentation using deformable surfaces. *Image and Vision Computing*, 10(3):132–144, April 1993. [cited at p. 79]
- [40] A. Dempster, N. Laird, and D. Rubin. Maximum likelihood from incomplete data via the EM algorithm. *Journal of Royal Statistical Society B*, 39:1–38, 1977. [cited at p. 23, 54, 109, 127, 129]
- [41] T. Deschamps and L.D. Cohen. Fast extraction of minimal paths in 3D images and applications to virtual endoscopy. *Medical Image Analysis*, 5(4), December 2001. [cited at p. 93, 143]
- [42] R. Duda and P. Hart. *Pattern Classification and Scene Analysis*. Wiley, NY, 1973. [cited at p. 37, 42, 45, 54, 55, 128]
- [43] M. Dulce and G.H. Mostbeck. Quantification of the Left Ventricular Volumes and Function with Cine MR Imaging: Comparison of Geometric Models with Three-Dimensional Data. *Radiology*, 188(371-376), 1993. [cited at p. 10, 13, 14]
- [44] J. S. Duncan and N. Ayache. Medical image analysis: Progress over two decades and the challenges ahead. *IEEE Transactions in Pattern Analysis and Machine Intelligence*, 22(1):85–106, Jan 2000. [cited at p. 25, 26]
- [45] N. Duta and M. Sonka. Segmentation and interpretation of mr brain images. an improved active shape model. *Medical Imaging, IEEE Transactions on*, 17(6):1049–1062, December 1998. [cited at p. 81, 83]

- [46] A. Falcão, J. Udupa, S. Samarasekera, and S. Sharma. Usersteered image segmentation paradigms: Livewire and livelane. *Graphical Models and Image Processing*, 60(4):233–260, 1998. [cited at p. 2, 22]
- [47] B. Fischl, D. H. Salat, E. Busa, M. Albert, Dieterich. M., C. Haselgrove, A. van der Kouwe, R. Killiany, D. Kennedy, S. Klaveness, A. Montillo, N. Makris, B. Rosen, and A. M. Dale. Whole brain segmentation: automated labeling of neuroanatomical structures in the human brain. *Neuron*, 33(3):341–355, Jan 2002. [cited at p. 127]
- [48] A. F. Frangi, W. J. Niessen, and M. A. Viergever. Three-Dimensional Modelling for Functional Analysis of Cardiac Images: A Review. *IEEE Transactions on Medical Imaging*, 20(1), 2001. [cited at p. 10, 24, 75, 85, 133]
- [49] A.F. Frangi, D. Rueckert, and J.S. Duncan. Three-dimensional cardiovascular image analysis. *IEEE Trans on Medical Imaging*, 21(9):1005–1010, 2002. [cited at p. 25]
- [50] A.F. Frangi, D. Rueckert, J.A. Schnabel, and W.J. Niessen. Automatic construction of multiple-object three-dimensional statistical shape models: application to cardiac modeling. *IEEE Transactions on Medical Imaging*, 21(9):1151– 1166, September 2002. [cited at p. 81]
- [51] K. D. Fritscher, R. Pilgram, and R. Schubert. Automatic cardiac 4d segmentation using level sets. In A. Frangi, editor, *Proc. of Functional Imaging and Modelling of the Heart*, pages 113–122. Lecture Notes in Computer Science, 2005. [cited at p. 102]
- [52] D. Geiger, A. Gupta, L. Costa, and J. Vlontzos. Dynamic programming for detecting, tracking and matching deformable contours. *IEEE Transactions in Pattern Analysis and Machine Intelligence*, 17(3):294–304, 1995. [cited at p. 25, 78]
- [53] G. Gerig, O. Kübler, R. Kikinis, and F.A. Jolesz. Nonlinear anisotropic filtering of MRI data. *IEEE Transactions on Medical Imaging*, 11(2):221–232, June 1992. [cited at p. 28, 37]
- [54] O. Ghita, K. Robinson, M. Lynch, and P. F. Whelan. MRI diffusion-based filtering: A note on performance characterisation. *Computerized Medical Imaging and Graphics*, 29(4):267–277, 2005. [cited at p. 132]
- [55] O. Ghita and P. F. Whelan. A bin picking system based on depth from defocus. *Machine Vision and Applications*, 13:234–244, 2003. [cited at p. 22, 146]
- [56] O. Ghita and P. F. Whelan. Integration of tina into neatvision: Developers guide. Technical Report IST-2001-34512, Dublin City University, September 2003. [cited at p. 24]

- [57] A. Goshtasby and D. A. Turner. Segmentation of cardiac cine mr images for extraction of right and left ventricular chambers. *IEEE Transactions in Medical Imaging*, 14(1):56–64, 1995. [cited at p. 22]
- [58] H. Gray. *Anatomy of the human body, by Henry Gray. 20th ed., thoroughly rev. and re-edited by Warren H. Lewis*. Philadelphia: Lea and Febiger, 20th edition, 1918. [cited at p. iv, 7, 8]
- [59] M. A. Grayson. The heat equation shrinks embedded plane curves to round points. *Journal of Differential Geometry*, 26:285–314, 1987. [cited at p. 90]
- [60] G. Hamarneh and T. Gustavsson. Combining Snakes and Active Appearance Shape Models for Segmenting the Human Left Ventricle in Echocardiographic Images. *IEEE Computers in Cardiology*, 27:115–118, 2000. [cited at p. 25, 81, 83]
- [61] J. A. Hartigan and M. A. Wong. Statistical algorithms: Algorithm AS 136: A K -means clustering algorithm. *Journal of Applied Statistics*, 28(1):100–108, March 1979. [cited at p. 54, 55]
- [62] C. K. Hoh, M. Dahlborn, G. Harris, Y. Choi, R. A. Hawkins, M. E. Philips, and J. Maddahi. Automated iterative three-dimensional registration of positron emission tomographic images. *Journal of Nuclear Medicine*, 34(11):2009–2018, 1993. [cited at p. 26]
- [63] K. H. Höhne and W. A. Hanson. Interactive 3D-segmentation of MRI and CT volumes using morphological operations. *J. Comput. Assist. Tomogr.*, 16(2):285–294, 1992. [cited at p. 54]
- [64] A. K. Jain, M. N. Murty, and P. J. Flynn. Data clustering: a review. *ACM Computing Surveys*, 31(3):264–323, 1999. [cited at p. 54]
- [65] A.K. Jain and R.C. Dubes. *Algorithms for Clustering Data*. Prentice-Hall, Inc., 1988. [cited at p. 23, 54]
- [66] M. P. Jolly. Combining edge, region, and shape information to segment the left ventricle in cardiac mr images. In *MICCAI*, pages 482–490, 2001. [cited at p. 25, 79]
- [67] M. P. Jolly, N. Duta, and G. Funka-Lea. Segmentation of the left ventricle in cardiac mr images. In *ICCV*, pages 501–508, 2001. [cited at p. 25, 79]
- [68] M. Kass, A. Witkin, and D. Terzopoulos. Snakes: Active Contour Models. *International Journal of Computer Vision*, 1(4):321–331, 1988. [cited at p. 24, 75, 86]
- [69] M. Kaus, J. von Berg, W. J. Niessen, and V. Pekar. Automated segmentation of the left ventricle in cardiac mri. *Medical Image Analysis*, 8:245–254, July 2004. [cited at p. 85, 114]

- [70] M. Khalighi, H. Soltanian-Zadeh, and C. Lucas. Unsupervised mri segmentation with spatial connectivity. In *Proceedings of the SPIE Int. Sym. on Medical Imaging*, February 2002. [cited at p. 92]
- [71] S. S. Khan and A. Ahmad. Cluster center initialization algorithm for k-means clustering. *Pattern Recognition Letters*, 25(11):1293–1302, 2004. [cited at p. 54]
- [72] G. Kiss, J. Cleynenbreugel, M. Thomeer, P. Suetens, and G. Marchal. Computer-aided diagnostics in virtual colongraphy via a combination of surface normal and sphere fitting. *European Radiology*, 12(1):77–81, 2002. [cited at p. 141]
- [73] J. J. Koenderink. The structures of images. *Biological Cybernetics*, 50:363–370, 1984. [cited at p. 28, 32]
- [74] J. J. Koenderink and A. J. van Doorn. Surface shape and curvature scales. *Image and Vision Computing*, 10(8):557–564, October 1992. [cited at p. 143]
- [75] T. Kohonen. *Self-Organizing Maps*. Springer, Berlin, 1995. [cited at p. 23]
- [76] D.H. Laidlaw, K.W. Fleischer, and A.H. Barr. Partial-volume bayesian classification of material mixtures in mr volume data using voxel histograms. *Medical Imaging, IEEE Transactions on*, 17(1):74–86, February 1998. [cited at p. 92]
- [77] R. Larsen and K. B. Hilger. Statistical shape analysis using non-euclidean metrics. *Medical Image Analysis*, 7(4):417–423, 2003. [cited at p. 24]
- [78] R. Larsen, M. B. Stegmann, S. Darkner, S. Forchhammer, T. F. Cootes, and B. K. Ersbøll. Texture enhanced appearance models. *Computer Vision and Image Understanding*, 2006. [cited at p. 24]
- [79] B. P. F. Lelieveldt, S.C. Mitchell, J.G. Bosch, R.J. van der Geest, M. Sonka, and J.H.C. Reiber. Time-continuous segmentation of cardiac image sequences using active appearance motion models. In *Proceedings of the 17th International Conference on Information Processing in Medical Imaging*. Lecture Notes In Computer Science, 2001. [cited at p. 84]
- [80] B.P.F. Lelieveldt, R.J. van der Geest, M. Ramze Rezaee, J. G. Bosch, and J. H. C. Reiber. Anatomical model matching with fuzzy implicit surfaces for segmentation of thoracic volume scans. *IEEE Transactions on Medical Imaging*, 18(3):218–230, 1999. [cited at p. 85]
- [81] M. E. Leventon, W. E. L. Grimson, and O. Faugeras. Statistical shape influence in geodesic active contours. In *Proceedings of Computer Vision and Pattern Recognition*, 2000. [cited at p. 93]

- [82] M. Lorenzo-Valdés, G. I. Sanchez-Ortiz, A. G. Elkington, R. H. Mohiaddin, and D. Rueckert. Segmentation of 4d cardiac MR images using a probabilistic atlas and the EM algorithm. *Medical Image Analysis*, 8:255–265, July 2004. [cited at p. 85]
- [83] M. Lynch, O. Ghita, and P. F. Whelan. Extraction of epi-cardial contours from unseen images using a shape database. In *Proceedings of Nuclear Science Symposium and Medical Imaging Conference*. NSS-MIC, 2004. [cited at p. 67]
- [84] T. Mäkelä, P. Clarysse, O. Sipilä, N. Pauna, Q. C. Pham, T. Katila, and I. E. Magnin. A review of cardiac image registration methods. *IEEE Transactions on Medical Imaging*, 21(9):1011–1021, September 2002. [cited at p. 26]
- [85] T. J. Mäkelä, P. Clarysse, J. Lötjönen, O. Sipilä, K. Lauerma, H. Hänninen, E.-P. Pyökkimies, J. Nenonen, J. Knuuti, T. Katila, and Magnin I. E. A new method for the registration of cardiac pet and mr images using deformable model based segmentation of the main thorax structures. In Niessen W. and Viergever M. A. (Eds.), editors, *Proceedings of the 4th International Conference on Medical Image Computing and Computer Assisted Intervention (MICCAI'01)*, pages 557–564. 2001 Springer-Verlag, 2001. [cited at p. 26]
- [86] R. Malladi and J. A. Sethian. Level set methods for curvature flow, image enhancement, and shape recovery in medical images. In *Proc. of Conf. on Visualization and Mathematics*, Berlin, Germany, June 1997. [cited at p. 25, 86, 90, 118]
- [87] Ravi Malladi, James A Sethian, and B C Vermuri. Shape modeling with front propagation: A level set approach. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 17:158–175, 1995. [cited at p. 86, 89]
- [88] H. A. McCann, J. C. Sharp, T. M. Kinter, C. N. McEwan, C. Barillot, and J. F. Greenleaf. Multidimensional Ultrasonic Imaging for Cardiology. *Proceedings of the IEEE*, 76(9):1063–1073, 1988. [cited at p. 13]
- [89] J. C. McEachen and J. S. Duncan. Shape-based tracking of the left ventricular wall motion. *IEEE Trans. on Medical Imaging*, 16(3):270–283, June 1997. [cited at p. 103]
- [90] T. McInerney and D. Terzopoulos. A dynamic finite element surface model for segmentation and tracking in multidimensional medical images with application to cardiac 4d image analysis. *Journal of Computerized Medical Imaging and Graphics*, 19(1):69–83, February 1995. [cited at p. 78]
- [91] T. McInerney and D. Terzopoulos. A dynamic finite element surface model for segmentation and tracking in multidimensional medical images with application to cardiac 4d image analysis. *Computerized Medical Imaging and Graphics*, 19(1):69–83, 1995. Special Issue on Cardiopulmonary Imaging. [cited at p. 103]

- [92] T. McInerney and D. Terzopoulos. Deformable models in medical images analysis: A survey. *Medical Image Analysis*, 1(2):91–108, 1996. [cited at p. 76, 78]
- [93] I. Mikić and S. Krucinski. Segmentation and Tracking in Echocardiographic Sequences: Active Contour Guided by Optical Flow Estimates. *IEEE Transactions on Medical Imaging*, 17(2), April 1998. [cited at p. 79]
- [94] S. C. Mitchell, B. P. F. Lelieveldt, J. G. Bosch, R. Van der Geest, J. H. C. Reiber, and M. Sonka. Segmentation of Cardiac MR Volume Data using 3D Active Appearance Models. *Proceedings of SPIE - The International Society for Optical Engineering*, 4684 I:433–443, 2002. [cited at p. 84]
- [95] S. C. Mitchell, B. P. F. Lelieveldt, R. J. van der Geest, H. G. Bosch, J. H. C. Reiber, and M. Sonka. Multistage Hybrid Active Appearance Model Matching: Segmentation of Left Ventricles in Cardiac MR Images. *IEEE Transactions on Medical Imaging*, 2001. [cited at p. v, 85]
- [96] S.C. Mitchell, J.G. Bosch, B.P.F. Lelieveldt, R.J. van der Geest, J.H.C. Reiber, and M. Sonka. 3-d active appearance models: segmentation of cardiac mr and ultrasound images. *IEEE Transactions on Medical Imaging*, 21(9):1167– 1178, September 2002. [cited at p. 84]
- [97] P. Mitra, S. K. Pal, and M. A. Siddiqi. Non-convex clustering using expectation maximization algorithm with rough set initialization. *Pattern Recognition Letters*, 24(6):863–873, 2003. [cited at p. 54]
- [98] D. Molloy and P. W. Whelan. Active-meshes. *Pattern Recognition Letters*, 21:1071–1080, 2000. [cited at p. 79]
- [99] J. Montagnat and H. Delingette. 4d deformable models with temporal constraints: application to 4d cardiac image segmentation. *Medical Image Analysis*, 9(1):87–100, February 2005. [cited at p. 103]
- [100] E. N. Mortensen and W. A. Barrett. Interactive segmentation with intelligent scissors. *Graphical Models and Image Processing*, 60(5):349–384, Sept. 1998. [cited at p. 2]
- [101] K. M. Mudry, R. Plonsey, and J. D. Bronzino, editors. *Biomedical Imaging*. CRC Press,, 2nd edition, 2003. [cited at p. 7]
- [102] S. G. Myerson, N. G. Bellenger, and D. J. Pennell. Assessment of left ventricular mass by cardiovascular magnetic resonance. *Hypertension*, 39:750–755, 2002. [cited at p. 14]
- [103] E. Nagel, R. Underwood, D. Pennell, U. P. Sechtem, S. Neubauer, S. Wickline, O. M. Hess, M. Schwaiger, and E. Fleck. New developments in non-invasive cardiac imaging: critical assessment of the clinical role of cardiac magnetic resonance imaging. *European Heart Journal*, 19(9):1286–1293, 1998. [cited at p. 117]

- [104] NASPE. Patient and Public Information Center, 1999-2003. [cited at p. 58]
- [105] K. S. Nayak, C. H. Cunningham, J. M. Santos, and J. M. Pauly. Real-time cardiac mri at 3 tesla. *Magnetic Resonance in Medicine*, 51(4):655–660, 2004. [cited at p. 27]
- [106] P. Neskovic and B. B. Kimia. Three-dimensional shape representation from curvature dependent surface evolution. In *IEEE Proceedings for the International Conference of Image Processing*, pages 6–10, November 1994. [cited at p. 91, 144]
- [107] N. M. I. Noble, D. L. G. Hill, M. Breeuwer, J. A. Schnabel, D. J. Hawkes, F. A. Gerritsen, and R. Razavi. Myocardial delineation via registration in a polar coordinate system. *Academic Radiology*, 10(12):1349–1358, December 2003. [cited at p. 61]
- [108] S. Osher and J. A. Sethian. Fronts propagating with curvature-dependent speed: Algorithms based on Hamilton-Jacobi formulations. *Journal of Computational Physics*, 79:12–49, 1988. [cited at p. 86, 87, 143]
- [109] N. Paragios. A variational approach for the segmentation of the left ventricle in cardiac image analysis. *International Journal of Computer Vision*, 50(3):345–362, 2002. [cited at p. 93, 94, 103]
- [110] N. Paragios. A level set approach for shape-driven segmentation and tracking of the left ventricle. *IEEE Transactions on Medical Imaging*, 22(6):773–776, June 2003. [cited at p. 25, 93, 122]
- [111] N. Paragios and R. Deriche. Geodesic active regions: A new paradigm to deal with frame partition problems in computer vision. *Journal of Visual Communication and Image Representation*, 13(2):249–268, 2002. [cited at p. 93, 104, 122]
- [112] N. Paragios, M. Rousson, and Ramesh. V. Knowledge-based registration and segmentation of the left ventricle: A level set approach. In *IEEE Workshop on Applications in Computer Vision*, Orlando, Florida, USA, December 2002. [cited at p. 93]
- [113] A. Pednekar, I. A. Kakadiaris, and U. Kurkure. Adaptive fuzzy connectedness-based medical image segmentation. In *Indian Conference on Computer Vision, Graphics and Image Processing 2002*, Ahmedabad, India, December 2002. [cited at p. 54]
- [114] A. Pednekar, I. A. Kakadiaris, R. Muthupillai, and S. Flamm. Automatic segmentation of the left ventricle from dual contrast cardiac mri data. In *Proceedings of the 4th Annual International Conference of the IEEE Engineering in Medicine and Biology Society*, pages 991–992, Houston, TX, Oct. 2002. [cited at p. 54]

- [115] P. Perona and J. Malik. Scale-space and edge detection using anisotropic diffusion. *IEEE Trans. on Pattern Analysis and Machine Intelligence*, 12(7):629–639, 1990. [cited at p. 28, 32, 33, 132]
- [116] M. Petrou and P. Bosdogianni. *Image Processing: The Fundamentals*. Wiley Publishing, Inc., 1st edition, 1999. [cited at p. 27, 28]
- [117] D. L. Pham, C. Xu, and J. L. Prince. A Survey of Current Methods in Medical Image Segmentation. Technical report, The John Hopkins University, The John Hopkins University, Baltimore, MD 21218, Jan 1998. [cited at p. 53, 75]
- [118] R. Razaee, C. Nyqvist, P. van der Zwet, E. Jansen, and J. Reiber. Segmentation of mr images by a fuzzy c-means algorithm. In *Computers in Cardiology*, pages 21–24. Computers in Cardiology, 1995. [cited at p. 23]
- [119] A. Rodgers and P. Vaughan. The World Health Report 2002. Reducing risks, promoting healthy life, The World Health Organisation, 1211 Geneva 27, Switzerland, 2002. [cited at p. 1]
- [120] M. Rogers and J. Graham. Robust Active Shape Model Search. In *Proceedings of the 7th European Conference on Computer Vision*, pages 517–530, May 2002. [cited at p. 83]
- [121] R. Ronfard. Region-based strategies for active contour models. *International Journal of Computer Vision*, 13(2):229–251, October 1994. [cited at p. 25, 79]
- [122] D. Rueckert and P. Burger. Automatic Tracking of the Aorta in Cardiovascular. *IEEE Transactions on Medical Imaging*, 16(5), October 1997. [cited at p. 79]
- [123] D. Rueckert and P. Burger. shape-based segmentation and tracking in 4d cardiac mr images. In *In Proc. International Conference on Computer Vision, Virtual Reality and Robotics in Medicine*, pages 43–52, Grenoble, France, 1997. CVRMed, Springer-Verlag. [cited at p. 103]
- [124] G.I. Sanchez-Ortiz. Fuzzy clustering driven anisotropic diffusion: Enhancement and segmentation of cardiac mr images. In *IEEE Nuclear Science Symposium and Medical Imaging Conference*, Toronto, Canada, November 1998. [cited at p. 28]
- [125] G.I. Sanchez-Ortiz, J. Declerk, M. Mulet-Parada, et al. Automating 3D Echocardiographic Image Analysis. In *Proceedings of MICCAI 2000*, pages 687–696, Pittsburg, PA, USA, Oct 2000. [cited at p. 13]
- [126] M. Santarelli and V. Positano. Automated cardiac MR image segmentation: Theory and Measurement Evaluation. *Medical Engineering and Physics*, 25, 2003. [cited at p. 79]

- [127] A. Savitzky and M. J. E Golay. Smoothing and differentiation of data by simplified least squares procedures. *Analytical Chemistry*, 36(8):1627–1639, 1964. [cited at p. 28, 30]
- [128] S. Schalla, E. Nagel, H. Lehmkuhl, C. Klein, A. Bornstedt, B. Schnackenburg, U. Schneider, and E. Fleck. Comparison of Magnetic Resonance Real-Time Imaging of Left Ventricular Function with Conventional Magnetic Resonance Imaging and Echocardiography. *The American Journal of Cardiology*, 87(1):95–99, 2001. [cited at p. 14, 117]
- [129] S. Sclaroff and A. P. Pentland. Modal matching for correspondence and recognition. *IEEE Transaction in Pattern Analysis and Machine Intelligence*, 17(7):545–561, 1998. [cited at p. 81]
- [130] U. Sechtem, P. Pflugfelder, and R. et al. White. Cine mr imaging: Potential for the evaluation of cardiovascular function. *Americal Journal of Roentgenology*, 148:239–246, 1987. [cited at p. 14]
- [131] M. Sermesant, C. Forest, X. Pennec, H. Delingette, and N. Ayache. Deformable biomechanical models: Application to 4d cardiac image analysis. *Medical Image Analysis*, 7(4):475–488, December 2003. [cited at p. 79]
- [132] J. A. Sethian. *An Analysis of Flame Propagation*. PhD thesis, University of California, Dept. of Mathematics, University of California, Berkeley, CA, USA, 1982. [cited at p. 86]
- [133] J. A. Sethian. A marching level set method for monotonically advancing fronts. In *Proceedings of the National Academy of Sciences*, volume 93, 1996. [cited at p. 86, 87, 94, 143]
- [134] J. A. Sethian. *Level set methods and fast marching methods*. Cambridge University Press, 2nd edition, 1999. Rev. ed. of: *Level set methods*. 1996. [cited at p. v, 89, 90, 94]
- [135] Y. Shang and O. Dössel. Landmarking method of 3d surface models for construction of 4d cardiac shape model. *Biomedizinische Technik*, 48(1):124–125, 2003. [cited at p. 81]
- [136] E. Sharon, A. Brandt, and R. Basri. Segmentation and boundary detection using multiscale intensity measurements. In *Computer Vision and Pattern Recognition*, pages 469–476. CVPR, 2001. [cited at p. 23, 55]
- [137] P. Shi, G. Robinson, R. T. Constable, A. Sinusas, and J. Duncan. A model-based integrated approach to track myocardial deformation using displacement and velocity constraints. In *Fifth International Conference on Computer Vision*, pages 687–692, 1995. [cited at p. 25]
- [138] B. Sievers, S. Kirchberg, A. Bakan, U. Franken, and H. J. Trappe. Impact of papillary muscles in ventricular volume and ejection fraction assessment by cardiovascular magnetic resonance. *Journal of Cardiovascular Magnetic Resonance*, 6(1):9–16, Jan 2004. [cited at p. 11]

- [139] S. Sinha, U. Sinha, J. Czernin, G. Porenta, and H. R. Schelbert. Non-linear assessment of myocardium perfusion and metabolism: Feasibility of registering gated mr and pet images. *American Journal of Roentgenology*, 36:301–307, 1995. [cited at p. 26]
- [140] B. Smolka, R. Lukac, A. Chydzinski, K. N. Plataniotis, and K. Wojciechowski. Fast adaptive similarity based impulse noise reduction filter. *Real Time Imaging*, 9(4):261–276, 2003. [cited at p. 27, 28]
- [141] P. Soille. *Morphological Image Analysis*. Springer-Verlag, Heidelberg, 2nd edition, 2003. [cited at p. 54]
- [142] M. Sonka, V. Hlavac, and R. Boyle. *Image processing, analysis and machine vision*. PWS Boston, 2nd edition, 1998. [cited at p. 28]
- [143] M. Sonka, S. C. Mitchell, B. P. F. Lelieveldt, J. C. Bosch, R. J. Van der Geest, and J. H. C. Reiber. Active appearance motion model segmentation. In *Second International Workshop on Digital and Computational Video*, pages 64–69, 2001. [cited at p. 84]
- [144] H. Späth. *Spline algorithms for curves and surfaces*. Utilitas Mathematica Pub., 1974. [cited at p. 64, 152]
- [145] L. J. Spreeuwiers and M. Breeuwer. Detection of Left Ventricular Epi- and Endocardial Borders Using Coupled Active Contours. In *Computer Assisted Radiology and Surgery*, pages 1147–1152, June 2003. [cited at p. 79]
- [146] L. H. Staib and J. S. Duncan. Boundary finding with parametrically deformable models. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 14(11):1061–1075, 1992. [cited at p. 79]
- [147] L. H. Staib and J. S. Duncan. Model-based deformable surface finding for medical images. *IEEE Transactions on Medical Imaging*, 15(5):720–731, October 1996. [cited at p. 80]
- [148] M.B. Stegman. AAM-API. <http://www.imm.dtu.dk/aam/>, 2002. C++ implementation of the Active Appearance Models (AAM) Framework. [cited at p. 84]
- [149] M. B. Stegmann. Active Appearance Models: Theory, Extensions and Cases. Master’s thesis, The Technical University of Denmark, Denmark, 2000. [cited at p. 81, 83, 84]
- [150] M. B. Stegmann, B. K. Ersbøll, and R. Larsen. FAME - a flexible appearance modelling environment. *IEEE Transactions on Medical Imaging*, 22(10):1319–1331, 2003. [cited at p. 24, 121]
- [151] M. B. Stegmann and H. B. W. Larsson. Motion-Compensation of Cardiac Perfusion MRI Using a Statistical Texture Ensemble. In *Functional Imaging and Modelling of the Heart (FIMH’03)*, pages 151–161, Lyon, France, June 2003. [cited at p. 24, 84]

- [152] M. B. Stegmann, H. Ólafsdóttir, and H. B. W. Larsson. Unsupervised motion-compensation of multi-slice cardiac perfusion MRI. *Medical Image Analysis*, 9(4):394–410, aug 2005. [cited at p. 24]
- [153] W. H. Strauss, B. L. Zaret, P. J. Hurley, T. K. Natarajan, and B. Pitt. A scintiphotographic method for measuring left ventricular ejection fraction in man without cardiac catheterization. *The American Journal of Cardiology*, 28(5):575–580, November 1971. [cited at p. 20]
- [154] R. M. Summers, C. D. Johnson, L. M. Pusanik, J. D. Malley, A. M. Youssef, and J. E. Reed. Automated polyp detection at CT colongraphy: Feasibility assessment in a human population. *Radiology*, 219(1):51–59, 2001. [cited at p. 141]
- [155] W. Sun, M. Çetin, R. Chan, V. Reddy, G. Holmvang, V. Chandar, and A. S. Willsky. Segmenting and tracking the left ventricle by learning the dynamics in cardiac images. In *IPMI*, pages 553–565, 2005. [cited at p. 103]
- [156] J. S. Suri. Computer vision, pattern recognition and image processing in left ventricle segmentation: The last 50 years. *Pattern Analysis and Applications*, 3(3):209–242, 2000. [cited at p. 25, 75]
- [157] J. S. Suri, K. Liu, S. Singh, S.N. Laxminarayan, X. Zeng, and L. Reden. Shape recovery algorithms using level sets in 2-d/3-d medical imagery: a state-of-the-art review. *IEEE Transactions on Information Technology in Biomedicine*, 6:8 – 28, 2002. [cited at p. 86]
- [158] J. S. Suri, S. K. Setarehdan, and S. Singh, editors. *Advanced Algorithmic Approaches to Medical Image Segmentation: State-of-the-Art Applications in Cardiology, Neurology, Mammography and Pathology*. Springer, Berlin, 2002. [cited at p. 14, 75]
- [159] J.S. Suri, D. Wu, J. Gao, S. Singh, and S. Laxminarayan. Comparison of state-of-the-art diffusion imaging techniques for smoothing medical/non-medical image data. In *Proc. 15th International Conference on Pattern Recognition (ICPR'02)*, Quebec, August 2002. [cited at p. 27]
- [160] K. Tang, J. Astola, and Y. Neuovo. Nonlinear multivariate image filtering techniques. *IEEE Trans. on Image Processing*, 4(6):788–797, 1995. [cited at p. 28]
- [161] M. Taron, N. Paragios, and M. P. Jolly. Modelling shapes with uncertainties: Higher order polynomials, variable bandwidth kernels and non parametric density estimation. In *ICCV*, pages 1659–1666, 2005. [cited at p. 93]
- [162] A. Tsai, W. M. Wells, S. K. Warfield, and A. S. Willsky. An EM algorithm for shape classification based on level sets. *Medical Image Analysis*, 9(5):419–502, October 2005. [cited at p. 93]

- [163] A. Tsai, Jr. Yezzi, A., W. Wells, C. Tempny, D. Tucker, A. Fan, W.E. Grimson, and A. Willsky. A shape-based approach to the segmentation of medical imagery using level sets. *IEEE Transactions on Medical Imaging*, 22(2):137–154, February 2003. [cited at p. 25, 93]
- [164] M. Urschler, H. Mayer, R. Bolter, and F. Leberl. The LiveWire Approach for the Segmentation of Left Ventricle Electron-Beam CT Images. In *26th Workshop of the Austrian Association for Pattern Recognition (AGM/AAPR)*, pages 319–326, 2002. [cited at p. 2, 22]
- [165] R. J. van der Geest, E. Jansen, V. G. M. Buller, and J. H. C. Reiber. Automated detection of left ventricular epi- and endocardial contours in short-axis MR images. *IEEE Computers in Cardiology*, 33(6):33–36, 1994. [cited at p. 84]
- [166] R. J. van der Geest and J. Reiber. *Cardiovascular MRI and MRA*, chapter 5. Lippincott, Williams and Wilkins, Philadelphia, 2003. [cited at p. 84]
- [167] R. J. van der Geest and J. H. C. Reiber. Quantification in cardiac mri. *Journal of Magnetic Resonance Imaging*, 10(5):602–608, Nov. 1999. [cited at p. 25]
- [168] B. van Ginneken, A. Frangi, J. Staal, B. ter Haar Romeny, and M. Viergever. Active shape model segmentation with optimal features. *IEEE Transactions in Medical Imaging*, 21(8):924–933, 2002. [cited at p. 24, 83]
- [169] F. Vos, P. W. de Bruin, J. G. M. Aubel, G. J. Streekstra, M. Maas, L. J. van Vliet, and A. M. Vossepoel. A statistical shape model without using landmarks. In *ICPR (3)*, pages 714–717, 2004. [cited at p. 81]
- [170] K. N. Walker, T. F. Cootes, and C. J. Taylor. Determining correspondences for statistical models of facial appearance. In *Proc. 4th IEEE Int. Conf. on Automatic Face and Gesture Recognition*, pages 271–276. IEEE Comput. Soc., 2000. [cited at p. 81]
- [171] J. Weickert. A review of nonlinear diffusion filtering. *Scale-Space Theory in Computer Vision*, 1252:3–28, 1997. Springer, Berlin. B. ter Haar Romeny, L. Florack, J. Koenderink and M. Viergever (Eds.). [cited at p. 28, 32]
- [172] J. Weickert, B. ter Haar Romeny, and M. Viergever. Efficient and reliable schemes for nonlinear diffusion filtering. *IEEE Transactions on Image Processing*, 7(3):398–410, 1998. [cited at p. 28]
- [173] C. Xu, Jr Yezzi, A., and J.L. Prince. A summary of geometric level-set analogues for a general class of parametric active contour and surface models. In *Variational and Level Set Methods in Computer Vision, 2001. Proceedings. IEEE Workshop on*, 2001. [cited at p. 91]
- [174] L. Xu and M. I. Jordan. On convergence properties of the EM algorithm for gaussian mixtures. *Neural Computation*, 8(1):129–151, 1996. [cited at p. 109, 127, 129]

- [175] F. Yan, H. Zhang, and Kube R. C. A multistage adaptive thresholding method. *Pattern Recognition Letters*, 26(8):1183–1191, 2005. [cited at p. 22]
- [176] J. Yao, M. Miller, M. Franaszek, and R. M. Summers. Colonic Polyp Segmentation in CT Colonography- Based on Fuzzy Clustering and Deformable Models. *IEEE Trans. in Medical Imaging*, 23(11):1344–1352, 2004. [cited at p. 141]
- [177] A Yezzi, S Kichenassamy, A Kumar, P Olver, and A Tannenbaum. A geometric snake model for segmentation of medical imagery. *IEEE Transactions on Medical Imaging*, 16(2):199–209, 1997. [cited at p. 89, 93]
- [178] H. Yoshida, Y. Masutani, P. MacEneaney, D. T. Rubin, and A. H. Dachman. Computerized detection of colonic polyps at ct colongraphy on the basis of volumetric features: A pilot study. *Radiology*, 222(2):327–336, 2002. [cited at p. 141]
- [179] J. N. Yu, F. H. Fahey, H. D. Gage, C. G. Eades, B. A. Harkness, and C. A. Pelizzari. Intermodality, retrospective image registration in the thorax. *Journal of Nuclear Medicine*, 36(12):2333–2338, Dec. 1995. [cited at p. 26]
- [180] A. Zavaljevskia, A.P. Dhawan, W. Gaskil, M. andBall, and J. D. Johnson. Multi-level adaptive segmentation of multi-parameter mr brain images. *Computerized Medical Imaging and Graphics*, 24(2):87–98, 2000. [cited at p. 134]
- [181] X. Zeng, L. H. Staib, R. T. Schultz, and J. S. Duncan. Segmentation and measurement of the cortex from 3D MR images. In *Medical Image Computing and Computer Assisted Intervention, MICCAI*, October 1998. [cited at p. 93, 118]
- [182] S.W. Zucker and R.A. Hummel. A three-dimensional edge detector. *IEEE Transactions on PAMI*, 3(3):324–331, 1981. [cited at p. 142]