

# Quantifying regional left ventricular function using spatio-temporal tracking techniques




Gary Jacob

Medical Vision Laboratory,  
Robotics Research Group,  
Department of Engineering Science,  
University of Oxford.

Trinity Term, 1999

This thesis is submitted to the Department of Engineering Science, University of Oxford, in partial fulfillment of the degree of Doctor of Philosophy. This thesis is entirely my own work, and except where otherwise indicated, describes my own research.



## **Quantifying regional left ventricular function using spatio-temporal tracking techniques**

### **Abstract**

Increasingly, diagnosis of cardiac disease, relies on computer processing of images to aid decision making. In this thesis, we use echocardiography, which is the most widely used cardiac imaging modality to study the motion of the left ventricle.

Currently, clinical reporting of echocardiography examinations is operator-dependent and largely qualitative. Commercially available software does not track the left ventricle. Also, it does not provide quantification of regional function.

This thesis establishes a framework for the quantitative regional analysis of left ventricular function. The endocardial and epicardial contours are automatically tracked during the cardiac cycle. A quantitative measure of regional endocardial wall excursion and myocardial thickening, based on a 16-segment model of the heart, is then obtained based on these boundaries.

The new tracking framework is based on Kalman filtering which makes a single prediction as to the position of the boundary on the next frame. We develop a measurement model for the endocardial border, the tissue/blood interface, and the epicardium, the tissue/tissue interface. Having tracked the endocardial and epicardial boundaries, we introduce an interpretational space which provides clinically meaningful regional quantitative measures of left ventricular function. We illustrate all the concepts on one example.

We apply the ideas developed to stress echocardiography, in a small retrospective clinical test.

---

## Acknowledgements

---

I'd firstly like to thank, my supervisor, Alison Noble, for her direction and assistance in the development of this work. I'm very grateful for the many hours she spent proof reading this thesis over the last few months. My sincerest thanks go to Andrew Blake, for his guidance and support along the way. To Mike Brady, the father of the lab, thanks for the moral support along the way.

To Brian Ripley, for the extremely helpful discussions about the work. Francis Marriott, your time and patience was much appreciated. To Anthony Davison, for keeping me here.

To my collaborators at the John Radcliffe Hospital. Andrew Kelion and Adrian Banning, saying thanks is not enough. Andrew, the many hours you spent helping me along the way won't go forgotten. You rekindled my belief that clinicians really do want to help patients.

To sum up the lab, I turn to Arsène Wenger <sup>1</sup>, with a reference to the Arsenal team. "Each person brings from his own culture the positive side," he explains, "which all comes together in the service of efficiency. That is the beauty. It is almost magical."

Dale, our beanies will live long in memory. We did it! Patrick, you were always ready to answer the 11am/3pm tea call, and not forgetting, explaining my jokes. Seb, our communication will be fondly remembered by all...even if your reply to me often was, "What!?! I don't understand you!"

Big-Mig, for your 11am preaching sermons, for being the intensity-invariant king, and those World Cup predictions...that 40 quid came in handy. Miguel 'The Angel', for the choice of music. Maud, for the delightful cakes on special occasions. Andrew, for the amusing stories.

Margaret, the newsletter queen, keep those wonderful lines coming. Kostas, for your deep stories and searching questions. Guofang, for allowing Albert and I to exploit the possibilities of Jaco-nanchung. To Chris, for providing that final inspiration to me, and doubtless in future, to the lab. Albert, I am proud to say that I survived your proof-reading. An idea for a T-shirt, maybe. Keep defacing those lab photos!

Ralph, for footie chat, and your thorough proof-reading. Cheers mate. 'Oh come-on' Gerardo, for you daily shouts, "Lunch, Balliol, Balliol, anyone,...Balliol" and, "Dinner, Holywell Manor anyone".

Gabriel, Paul and Jacques, for making the lab what it is. I am indebted to everyone who trudged their way along proof reading my thesis: Alison, Andrew, Ralph, Patrick, Chris, Gerardo, Albert, Guofang and Margaret.

From the visual dynamics group, I am very grateful to Michael Isard for answering my seemingly endless questions about the tracker. And in his absence, to Ben North and Benedicte Bascle, for their help.

---

<sup>1</sup>Arsène Wenger, a football coach for Europe, The Economist, 27th February 1999.

This thesis is about heart function. Sadly, during my thesis, two colleagues, David Notley and Richard Eagle, lost their lives to heart disease.

St Catherine's, "a would-be fashionable Oxford College <sup>2</sup>", and, "the world's coolest campest architecture", according to Wallpaper Photographic magazine <sup>3</sup>, not to mention, "its right-on reputation", as The Sunday Times style magazine put it <sup>4</sup>. What it is to be famous, and a great place to while away the years.

Michael Spence, Dean of St. Catherine's, I owe you more than I can repay back. Who would of have heard of an Oxford College taking on a 22 year old Junior Dean without an interview. I have learnt much from you. I hope that I've repaid back your faith in me. I'm grateful for the unconstrained way you let me go about my job. Being Junior Dean has changed me inexplicably for the better. Nothing really matters, does it?

To name but a few of the many people I have met during my time in Oxford. Rachel, for being there for me, and providing the one ray of light, in those dark early days. Mikey, 'slacker', for the many chats. Mike for football chat and reading my many sports letters. To Emily, for dragging me to India, and those timely cups of tea. To Charlotte, my whinings grew longer at the end. Simon Goodwin for 'Jacob Time', and Si Fallis, the medic to-be, for his words of wisdom on my thesis. Thanks also to Tom, Miranda and Michael Hart.

Bob Hiorns and Ian Craig — thanks for giving me that first opportunity to tutor. To Dave Popplewell, my 'tutoring PR officer', who recommended me to numerous Colleges. Your advice about tutoring was always spot on. Simon Hunt and Nick Emler, a special mention. You were only too ready to answer my emails. You put back my faith 'that there really *were* tutors in Oxford who were here to teach'. A special mention to two of my earliest students, Anna and Helen, who concreted my belief in the power, not necessarily positive, that tutors can have.

I am grateful to the Engineering and Physical Science Research Council for sponsoring me for the past three years. Thanks also to the various funding bodies that provided sponsorship for my conference trips. India was an experience.

To my Mum, Dad and family, those timely foods parcels were just the ticket.

Gary Jacob  
30th June 1999  
Oxford.

---

<sup>2</sup>Top Catz, The Times Diary, The Times, 7th December 1998.

<sup>3</sup>Wallpaper 100 Special Issue, November/December 1998, pages 75–80.

<sup>4</sup>The high and matey, The Sunday Times Style Magazine, 6th September 1998.

---

# Contents

---

<b>1</b>	<b>Introduction</b>	<b>1</b>
<b>2</b>	<b>Background</b>	<b>7</b>
2.1	Cardiac Imaging . . . . .	7
<b>3</b>	<b>Literature Review</b>	<b>17</b>
3.1	Visual tracking . . . . .	17
3.1.1	Flow-based . . . . .	18
3.1.2	Model based . . . . .	19
3.2	Computer Vision . . . . .	25
<b>4</b>	<b>Shape-spaces and Dynamical Tracking</b>	<b>31</b>
4.1	B-splines as tracking templates . . . . .	32
4.2	Tracking strategy . . . . .	34
4.3	Shape-spaces . . . . .	35
4.3.1	Definition of a shape-space . . . . .	36
4.4	Probabilistic shape-spaces via Principal Component Analysis (PCA) . . . . .	39
4.5	Predictive Dynamical Model . . . . .	42
4.5.1	Dynamical Model . . . . .	44

4.5.2	Learning the dynamical model . . . . .	47
4.6	The measurement model . . . . .	48
4.7	Conclusions . . . . .	51
<b>5</b>	<b>Tracking Endocardial Wall Excursion</b>	<b>52</b>
5.1	Prior Shape Model . . . . .	55
5.1.1	Deformation modes . . . . .	56
5.1.2	How many PCA modes should we use? . . . . .	60
5.1.3	Can we assume an affine mode of deformation? . . . . .	61
5.1.4	Comparing shape models by visual inspection . . . . .	62
5.2	Component plots . . . . .	65
5.3	Improving feature detection . . . . .	66
5.3.1	Localisation based on local phase information . . . . .	70
5.3.2	Comparing feature detection schemes . . . . .	73
5.4	Conclusions . . . . .	76
<b>6</b>	<b>Interpretational shape-spaces</b>	<b>77</b>
6.1	Heart Function . . . . .	78
6.1.1	Global systolic function . . . . .	78
6.1.2	Regional systolic function . . . . .	83
6.2	The 16-segment model of the heart . . . . .	84
6.3	Interpretational shape-spaces . . . . .	87
6.3.1	Segment variation in wall excursion . . . . .	93
6.4	Segmental scoring of wall excursion . . . . .	96
6.4.1	Scoring scheme normalisation . . . . .	97
6.5	Tracking in a clinical space . . . . .	97
6.6	Clinical plots . . . . .	100
6.6.1	Colour kinesis for wall excursion . . . . .	101
6.6.2	Velocity coding of wall motion . . . . .	103

6.7	Conclusions . . . . .	104
<b>7</b>	<b>Myocardial Wall Thickening</b>	<b>106</b>
7.1	Myocardial tracking strategy . . . . .	108
7.2	Tracking epicardial boundary motion . . . . .	109
7.3	Epicardial shape-space estimation . . . . .	111
7.4	Fitting spline templates and choice of epicardium shape-space . . . . .	113
7.4.1	Choice of search scale . . . . .	118
7.5	Measurement model to localise epicardium features . . . . .	120
7.6	Localisation based on a wavelet ridge detector . . . . .	121
7.6.1	Choice of wavelet . . . . .	122
7.6.2	Assimilating information using wavelets . . . . .	124
7.7	Quantifying segmental wall thickening . . . . .	127
7.7.1	Segment variation in wall thickening . . . . .	128
7.8	Segmental scoring of wall thickening . . . . .	130
7.9	Comparison of wall excursion and wall thickening . . . . .	131
7.10	Colour kinesis for wall thickening . . . . .	132
7.11	Conclusions . . . . .	133
<b>8</b>	<b>Stress Echocardiography</b>	<b>135</b>
8.1	Background to stress echocardiography . . . . .	136
8.2	The John Radcliffe Hospital dobutamine stress echocardiography protocol . . . . .	137
8.3	Stress echocardiography analysis . . . . .	139
8.3.1	John Radcliffe Hospital study results . . . . .	141
8.3.2	Comparison of tracker . . . . .	142
8.3.3	Difficulties . . . . .	147
8.4	Clinical Evaluation . . . . .	147
8.4.1	Clinician's segmentation . . . . .	148
8.4.2	Contour comparison . . . . .	151

---

8.5	Conclusions . . . . .	159
<b>9</b>	<b>Discussion and Future Work</b>	<b>161</b>
9.1	The work in this thesis . . . . .	161
9.2	Future work . . . . .	163
9.3	Building models for all views . . . . .	163
9.4	Using Condensation . . . . .	163
9.5	Using temporal information for feature detection . . . . .	164
9.6	Epicardial feature detection . . . . .	164
9.7	Shape space for epicardium estimation . . . . .	164
9.8	Quantitative scoring of wall function . . . . .	165
9.9	Towards a tracker-driven user-interface . . . . .	165
9.9.1	Patient dependent tracker framework . . . . .	166
9.9.2	Generic tracker framework . . . . .	167
9.9.3	3-dimensional echocardiography . . . . .	169
<b>A</b>	<b>Glossary of clinical terms</b>	<b>170</b>
<b>B</b>	<b>Anatomical Background</b>	<b>176</b>
B.1	The Heart . . . . .	176
B.1.1	Functions of the cardiovascular system . . . . .	176
B.1.2	Heart chambers and valves . . . . .	177
B.1.3	Path of blood through the vasculature. The Circulation . . . . .	179
B.1.4	Surface anatomy . . . . .	180
B.1.5	Structure of the heart wall . . . . .	180
B.2	Contraction and excitation of a cardiac myocyte . . . . .	181
B.2.1	The myocyte . . . . .	182
B.2.2	Contraction . . . . .	183
B.2.3	The mechanics of the myocardium . . . . .	183

<b>C</b>	<b>Basic principles of ultrasound</b>	<b>185</b>
C.1	Properties of ultrasound . . . . .	185
C.1.1	Acoustic impedance . . . . .	186
C.1.2	Attenuation . . . . .	186
C.1.3	Scattering . . . . .	187
C.2	Image formation . . . . .	187
<b>D</b>	<b>Miscellaneous Matrix Results</b>	<b>189</b>
D.1	Matrix Operators . . . . .	189
D.1.1	The Kronecker product . . . . .	189
D.1.2	The Moore-Penrose (MP) Inverse . . . . .	189
D.2	Matrix Differentiation . . . . .	189
D.2.1	Definition of Matrix differentiation . . . . .	190
<b>E</b>	<b>Derivations and proofs</b>	<b>191</b>
E.1	Proof that $W^\dagger = (W^\top \mathcal{H} W)^{-1} W^\top \mathcal{H}$ . . . . .	191
E.2	Maximum Likelihood Estimation for parameters $A_0, A_1$ and $B$ . . . . .	192
E.2.1	Univariate case . . . . .	192
E.2.2	Multivariate case . . . . .	193
E.3	Myocardial thickening . . . . .	195
E.4	Generic tracker model — parameter estimation . . . . .	196
E.4.1	Constant dynamics . . . . .	196
<b>F</b>	<b>Tracking results</b>	<b>199</b>
F.1	Comparing affine and PCA tracking shape-spaces . . . . .	199
F.2	Comparing feature detection schemes . . . . .	206
<b>G</b>	<b>Stress Echocardiography</b>	<b>211</b>
G.1	Motivation for stress echocardiography . . . . .	211
G.1.1	Traditional stress echocardiography . . . . .	211

---

G.1.2 Pharmacologic stress testing . . . . .	213
G.1.3 Dobutamine stress echocardiography procedure . . . . .	214
G.2 Stress echocardiography analysis . . . . .	215
G.2.1 Qualitative analysis of wall motion . . . . .	215
G.2.2 Quantitative analysis of wall motion . . . . .	216
G.2.3 Safety . . . . .	218
G.2.4 Sensitivity, Specificity and Accuracy . . . . .	218
G.3 Comparison of Dobutamine echocardiography, Dipyridamole echocardiography and Exercise echocardiography . . . . .	219
G.3.1 Inter-observer and Intra-observer agreement . . . . .	220
<b>H John Radcliffe Hospital dobutamine stress scoring sheet</b>	<b>222</b>
<b>I Contour Comparison results</b>	<b>223</b>

---

## Glossary of Notation

---

<b>SYMBOL</b>	<b>MEANING</b>
$\triangleq$	defined to be
$\square$	end of proof
$\top$	superscript denoting the transpose of a vector/matrix
$\otimes$	Kronecker product of two matrices
$\mathcal{R}$	the space of reals
$\hat{x}$	<i>estimate</i> for the quantity $x$
$\tilde{x}$	<i>prediction</i> of the quantity $x$
$\mathbf{0}$	null matrix, null vector
$\mathbf{I}$	identity matrix, identity vector
$\langle \cdot, \cdot \rangle$	inner product for functions/curves
$\ \mathbf{Y}\ $	norm of matrix $\mathbf{Y}$ , norm of vector $\mathbf{Y}$ ,
$\ \mathbf{Y} - \mathbf{Z}\  = (\mathbf{Y} - \mathbf{Z})^\top (\mathbf{Y} - \mathbf{Z})$	norm of $\mathbf{Y} - \mathbf{Z}$
$\text{Min}(a, b)$	the minimum of $a$ and $b$

<b>SYMBOL</b>	<b>MEANING</b>
<b>A</b>	deterministic constant coefficient matrix in discrete dynamical model
$\mathbf{A}_1, \mathbf{A}_2$	constant components of $\mathbf{A}$ ; $\mathbf{A} = \begin{pmatrix} \mathbf{0} & \mathbf{I} \\ \mathbf{A}_2 & \mathbf{A}_1 \end{pmatrix}$
$\mathbf{B}(s)$	vector of B-spline blending functions parametrized by $s$
<b>B</b>	stochastic coefficients in discrete dynamical model
$\mathbf{B}_0$	stochastic coefficients in 2nd order discrete dynamical model
$\mathcal{B}$	metric matrix for B-spline functions, where, $\mathcal{B} = \frac{1}{L} \int_0^L \mathbf{B}(s)\mathbf{B}^\top(s)ds$
<b>C</b>	covariance coefficient $\mathbf{C} = \mathbf{B}\mathbf{B}^\top$
$E(x)$	expectation of a random variable $x$
<b>H</b>	observation measurement matrix
$\mathcal{H}$	metric matrix for curves in shape space $\mathcal{S}$
$L$	number of spans on B-spline curve
$M$	number of frames in a training sequence. $\mathbf{X}_1, \mathbf{X}_2, \dots, \mathbf{X}_M$
$N$	number of sampled image-features along a B-spline curve
$N_B$	number of control points on B-spline curve
$N_Q$	dimension of spline space
$N_X$	dimension of configuration space (shape-space) $\mathcal{S}$
$p(x)$	a probability density function
$P(A)$	the probability of an event $A$
$\mathbf{q}_n$	control point for a spline curve, $\mathbf{q}_n = (q_n^x, q_n^y)^\top$
$\mathcal{P}(t_k)$	covariance of state-vector $\mathcal{X}(t_k)$ at time $t_k$
$\tilde{\mathcal{P}}(t_k)$	covariance of predicted state-vector $\dot{\hat{\mathcal{X}}}(t_k)$
<b>Q</b>	vector of control points
$\mathbf{Q}^x$	vector of $x$ -coordinates of control points
$\mathbf{Q}^y$	vector of $y$ -coordinates of control points

<b>SYMBOL</b>	<b>MEANING</b>
$Q_0$	control points of template curve
$S$	statistical information matrix
$S_Q$	space of spline curves $Q$
$S_X$	shape-space
$s$	spatial parameter for curve
$U(s)$	matrix mapping control point vector $Q$ to image-curve $r(s)$
$U$	metric matrix for B-spline parametric curves, where, $U = I_2 \otimes B$
$W$	shape-matrix mapping from configuration $X$ to control point vector $Q$
$W^\dagger$	pseudo-inverse of mapping $W$
$X$	curve shape-vector
$\bar{X}$	mean value of curve configuration
$\hat{X}(t_k)$	estimated state vector from a tracker
$\tilde{X}(t_k)$	predicted state vector in a tracker
$\bar{X}$	mean value of state $X$
$Z^{-1}$	denotes the inverse of a matrix $Z$

"If I have learnt one thing," he said yesterday, "it is that you must never give up, you must never stop trying to achieve what you dream of. It came as a complete shock to me when they found the problem with my heart because I had had no warning, no symptoms, no pain, but I was not discouraged."

— Nwankwo Kanu.

*Beat goes on for Arsenal's enigma*

Oliver Holt

The Times

Saturday 1 May 1999.

# CHAPTER 1

---

## Introduction

---

Put yourself in the position of the trainer of the most successful football team in the world. In attack, you have the most prolific goal-scorer in the world. He can't stop scoring. Then suddenly, he misses a few goal chances. And then some more. And then more still. The supporters and Chairman want to know what's happening. As the trainer, it is your job to find out the problem for his poor performance. Day after day you watch him in training, from lots of different angles. Is he kicking the ball too weakly? Does he have an injury? Now imagine it's winter. There is thick fog about, which is clouding your vision.

An echocardiographer performs precisely the same job every day. The player is the patient. His leg, a patient's heart. The injury, heart disease. The fog, the image noise which makes the echocardiographer's job harder. Kicking the ball, the pumping of blood around the body. A goal miss, the death of a portion of heart muscle as a result of inadequate blood supply (an infarct).

Diagnosis of cardiac disease increasingly relies on computer processing of images to aid decision making. The work in this thesis is concerned with echocardiography, i.e. using ultrasound to image the cardiac chambers. Examples of echocardiographic images are shown in Figure 1.1 (the quality of this image is very good; typically, echocardiographic images are considerably noisier than this, as shown in Figure 1.1 (d)).

The use of ultrasound imaging allows for the detection of functional abnormalities in the cardiac chambers and valves. For example, ischemia, which represents inadequate blood supply

to a part of the heart, manifests as a change in the kinetics of the heart. This results in a reduced contractility in the ischemic region, but to compensate for this, non-ischemic regions work harder. The net result is that the global function of the heart appears totally normal. Interest must therefore be focused on looking at regional, rather than global, heart function. In particular, the clinician examines the left ventricle, the principal powerhouse of the heart, which pumps oxygen-rich blood throughout the body. The left ventricular wall is principally made-up of muscle called *myocardium*. The myocardium is protected by an outer covering, called the *epicardium*, and by a protective inner lining, called the *endocardium*. During a cardiac cycle, the myocardium produces a muscular contraction that forces blood from the left ventricle to the body.

To examine the heart, the clinician uses quantitative measures of left ventricular function. But, to get them, it requires the clinician to take time consuming manual measurements. Automating this process would be very beneficial to the clinician, but this is a very difficult problem, which has yet to be solved. We tackle this in this thesis. We derive two clinically meaningful measures based on the movement of the endocardium (endocardial wall excursion) and thickening of the myocardium (myocardial thickening), which are intended for use in a clinical setting. An example of these measurements for the patient in Figure 1.1 (a) is shown in Figure 1.1 (b,c). This patient has a myocardial infarct in the apical region of the heart (this is the uppermost two segments in Figure 1.1 (b)), which can be correctly picked out by the reduced movement and thickening of this segment of the left ventricle. Although the endocardial wall excursion in the upper-left segment appears normal, the segment shows reduced myocardial thickening. This is because the movement of the adjacent segment has masked the abnormality. This shows that myocardial thickening is a more accurate method for the detection of ischemia than endocardial wall excursion.

We apply the approach to stress echocardiography. In stress echocardiography, the patient is exercised (physically or by use of drugs), so that the heart demands more oxygen. In the case of coronary artery disease, the supply cannot match the demand because the 'narrowed' artery is reducing the supply of blood to parts of the heart.

In the rest of this chapter, we overview the thesis and stress the several important contributions of it in the development of quantitative regional left ventricular function.

We begin with an introduction to cardiac imaging (Chapter 2), and compare and contrast the current status of a commercially available software tool for following endocardial wall excursion with our own method. We then review the literature for general tracking and cardiovascular tracking (Chapter 3). Our goal in these chapters is to cast this research in light of what has

previously been accomplished in these fields. The tracking framework is introduced in Chapter 4. The main contributions of this thesis are presented in Chapters 5 to 8.

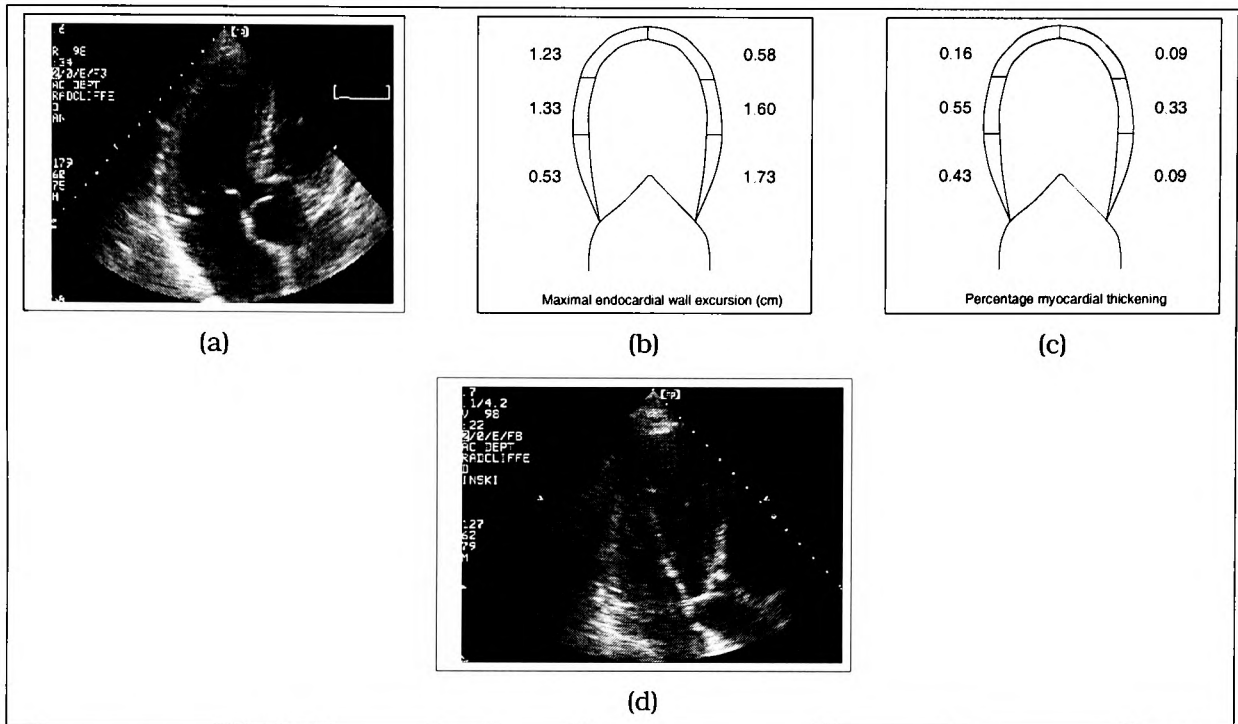


Figure 1.1: **Quantitative measures of left ventricular heart motion based on echocardiography.** Echocardiography is a non-invasive diagnostic method that utilizes high frequency sound to assess anatomical and functional abnormalities of the cardiovascular system. The movement and thickening of the left ventricular boundary are important indicators of the performance of the left ventricle. We are able to automatically measure these. (a) An echocardiographic image acquired from one of the latest ultrasound machines; images can be stored digitally. Note that displaying the echocardiographic image at a reduced size, as in the figure, reduces the visual effect of speckle and clutter noise. (b) Endocardial movement (in cm). Note the reduced wall motion in the upper right segment. This observation is exactly in accordance with the diagnosis that this patient has a myocardial infarct in the apical region. In fact the infarct extends into the upper left segment, but the movement of the adjacent segment has masked the abnormality. (c) Thickening, expressed as a percentage of thickening in diastole (expansion). Note the reduced wall thickening in the upper two segments. This shows that myocardial thickening is a more accurate method for the detection of ischemia than endocardial wall excursion. (d) A stress echocardiography image of a patient at rest.

### Endocardial wall tracking

Chapter 4 introduces the tracking framework used in this thesis. A visual 2-dimensional dynamic model-based tracker with shape constraints (Blake *et al.*, 1995) is adapted to the problem of left ventricular function. The chapter also outlines the three key tracking ingredients: an appropriate model of shape; modelling and learning the object's motion to guide a prediction and using specific image feature detectors to enhance the measurement process, which in turn updates the contour's position.

Chapter 5 applies the ideas introduced in Chapter 4 to the challenging problem of tracking endocardial wall motion. We show that this provides a robust way of tracking the endocardial

wall motion. Figure 1.2 (a) shows an example of tracking the endocardium.

### Quantifying endocardial wall excursion

Chapter 6 develops a framework for clinically interpreting tracking sequences. We introduce an interpretational space to decompose the tracking into clinically meaningful regional time-varying components. The method quantifies the regional endocardial wall excursion of the left ventricle, which can then be used for automatic quantitative endocardial wall scoring (Figure 1.1 (b)). We show that we are able to find an abnormal region of the left ventricle.

We also provide clinically useful plots of wall function, for example as shown in Figure 1.3 (a,b).

### Myocardial tracking

Chapter 7 introduces the very challenging problem of tracking myocardial thickening.

We use the estimation of the endocardium as a starting point in estimating the epicardium.

Features along the epicardium are located using a wavelet-based ridge detector (Figure 1.2 (b))

We then use the tracked endocardial and epicardial borders to quantify percentage myocardial thickening (Figure 1.1 (c)).

We also provide clinically useful plots of wall function, for example as shown in Figure 1.3 (c).




Figure 1.2: **Tracking the left ventricular boundary.** (a) Tracking the endocardium. (b) Tracking the endocardium (red) and epicardium (yellow).




Figure 1.3: **Clinical plots.** Pictorial ways of assessing regional left ventricular motion.

(a) Inter-frame velocity of the endocardium. Low velocities are coded in blue, and faster velocities in red, with a natural scale in-between.

(b) Colour kinesis chart for endocardial excursion. Colour kinesis keeps track of the estimated endocardium and colour codes the moving endocardial border to clearly show, on every frame, and in real time, sequential stages of wall motion throughout the cardiac cycle. The present contour position is coded in blue, the oldest in red, with a natural scale in-between.

(c) Colour kinesis chart for myocardial thickening. We have extended the idea to wall thickening. For every point along the myocardium, the change in the thickening between frames is calculated. The present contour position is coded in blue, the oldest in red, with a natural scale in-between.

### Stress echocardiography

Finally, in Chapter 8, we apply the ideas developed in Chapters 5 to 7 to stress echocardiography. This enables us to evaluate the performance of the combined work in a retrospective clinical test.

Stress echocardiography provides an ideal medium to apply the tracker, because it is intra-, rather than inter-patient, variation that is important here. To assess the response of the heart to an increasing demand for oxygen, a patient is infused with varying levels of a drug (dobutamine). Automatically quantifying the change in heart response between levels has never been done before.


One of the central problems in the area of medical image analysis is the validation of results from a segmentation algorithm. The performance of the tracking techniques described in this thesis is evaluated against ground truth, which was taken to be a clinician's segmentation. A number of quantitative measures of regional comparison are proposed.

### Discussion and future work

Chapter 9 concludes the thesis. The methodology is reviewed and its shortcomings are discussed, along with possibilities for future work.

To improve readability and complement these chapters, we provide a number of appendices and two glossaries, which endeavour to free the main text of results and proofs.

Where some degree of clinical background is needed, we have indicated this with a 'heart' sign, and where some mathematics is needed we have indicated this with a 'sigma' sign. In each case, without loss, the reader can rejoin the discussion with an 'all clear' (an idea stolen from Knuth's  $\text{\TeX}$  manual).



At times, the reader may find a reminder of the notation useful. This is provided by a glossary on page vii. Readers with little background in clinical medicine / cardiovascular imaging may find some of the medical terms unfamiliar.

A glossary of medical heart terms is given in Appendix A.

The basic anatomy and physiology of the heart are detailed in Appendix B.

The basic fundamentals of ultrasound image formation and acquisition are detailed in Appendix C.

Mathematical proofs are given in Appendices D and E.

So as not to overwhelm the main body of the thesis, various tracking results are shown in Appendix F.

A review of stress echocardiography is provided in Appendix G and the dobutamine stress echocardiography scoring sheet used at the John Radcliffe Hospital is shown in Appendix H.

Results obtained from an evaluation of the tracking techniques described in this thesis against a clinician's segmentation are given in Appendix I.

# CHAPTER 2

---

## Background

---

Cardiac patient diagnosis typically begins with an assessment of the patient's history and a physical examination, usually followed by an electrocardiogram (ECG). Initial hypotheses are then evaluated using an array of diagnostic tools, the most common of which is cardiac imaging (ultrasound, nuclear medicine PET, SPECT angiography and magnetic resonance imaging). Through cardiac imaging methods, a clinician typically attempts to visualize cardiac anatomy and (less often) tries to derive quantitative descriptors. Anatomical assessment includes determination of the size, shape and structure of the cardiac chambers, valves, great vessels and coronary arteries.

### 2.1 Cardiac Imaging

The most widely used cardiac imaging methods, for example, ultrasound and nuclear medicine, offer information about anatomy and function. There has also been some recent work on assessment of myocardial perfusion, metabolism and tissue characterisation (Skorton *et al.*, 1996, page 5).

The field of cardiac imaging has grown substantially over the last five years, and there are now numerous and classic books on cardiac imaging, for example, Morganroth *et al.* (1983); Buda & Delp (1985); Kotler & Steiner (1986); Marcus & Braunwald (1991); Feigenbaum (1994) and Panidis (1996). Most, if not all of these, have a clinical and qualitative bias.

The importance of wall thickness in evaluating the mechanical performance of the left ventricle was recognised as early as late last century by Woods (1892). In 1935, Tennant & Wiggers (1935) described the movement of an area of the left ventricular myocardium (by recording the movement of a lever arm attached to the myocardium) in a dog following ligation of the coronary artery supplying the area. However, it was not until the late 1960's and early 1970's that the abnormal regional performance of the human left ventricle became a major interest and research in clinical cardiology (Baxley & Reeves, 1971). There are two reasons for this historical sequence. Firstly, a reasonable complete hemodynamic evaluation of many of man's cardiac abnormalities had to await the advent of cardiac catheterization and angiography. Even then, diseases causing the most gross alteration in the heart's anatomy and pathophysiology, such as valvular or congenital defects, were in general the easiest and hence the first studied. Secondly, cardiac surgery progressed, and it was demonstrated that the resection of non-contracting scars can improve the condition of some patients with chronic heart failure following myocardial infarction.

In addition, there has been an increasing interest in obtaining measurements of the left ventricle, for example, chamber dimensions, area, volume and the ejection fractions. Interest has been driven to automate these measurements. This has led to a large amount of work aimed at developing algorithms for image processing, segmentation and tracking of cardiac images (Zhang & Geiser, 1982; Mailloux *et al.*, 1987; McEachen II & Duncan, 1997; Weng *et al.*, 1997; Gorce *et al.*, 1997; Declerck *et al.*, 1997).

The principal way of imaging the heart is based on ultrasound. The use of high-frequency sound (3–5MHz) enables the clinician to detect functional abnormalities in the cardiac chambers and valves.

Echocardiography is non-invasive <sup>1</sup> and has no demonstrable side effects. For these reasons, and those listed later, echocardiography is currently the most widely used imaging technique for the assessment of cardiovascular anatomy and function. A review is provided by Wells (1993). Four examples of echocardiographic images of the left ventricle are shown in Figure 2.1. These examples convey the varying degrees of difficulty the clinician has in making their assessment. They also show the challenging task we face in trying to model left ventricular function. Images are dependent on the ultrasound machine and the patient. The image quality from older machines can be poor, and the data may be stored on video (with all its degradation effects). Figure 2.1 (a) shows an example of this.

Improved images are possible using newer machines, but more often than not, images are still stored on video. Only the most recent machines allow the clinician to store images digitally. Even so, costs rule this out for everyday use. No machine can alleviate the difficulty with trying

---

<sup>1</sup>If no contrast agents are used.




Figure 2.1: **Echocardiographic images. Examples.** Echocardiography is a non-invasive diagnostic method that utilizes high frequency sound to assess anatomical and functional abnormalities of the cardiovascular system. Four images are shown which convey the varying degree of difficulty the clinician has in making their assessment. They also show the challenging task we face in trying to model heart function. (a) A image digitised from a video recording. (b) A 'good' image acquired from one of the latest ultrasound machines; images can be stored digitally. (c) A stress echocardiography image of a patient at rest. (d) The same patient as in (c) at a peak dose of dobutamine. Note that displaying the echocardiographic image at a reduced size, as in the figure, reduces the visual effect of speckle noise.

to image different patients. People are different, as are their bodies. For example, one of the latest ultrasound machines <sup>2</sup>, and that used for part of the work in this thesis, can produce an image of Figure 2.1 (b). But the same machine can have difficulty with a different patient as Figure 2.1 (c) shows. [Note that displaying the echocardiographic image at a reduced size, as in the figure, reduces the visual effect of speckle noise.] Moreover, the same machine can have difficulty with the same patient under two different imaging protocols. Figure 2.1 (c) represents an image of a patient, at rest, during a stress echocardiography examination. Correspondingly, Figure 2.1 (d) shows the same patient at peak stress. The work in this thesis needs to aim to be robust enough to cope with these problems.

<sup>2</sup>Hewlett-Packard SONOS 5500.

To complement the clinician's visual examination, time-consuming manual measurements are recorded, for example, the left ventricular area, volume, ejection fraction and dimensions. The clinical demand is to automate this process. Interest has therefore been centred on segmenting the left ventricular boundary (Zhang & Geiser, 1982; Chu *et al.*, 1988; Han *et al.*, 1991; Wilson *et al.*, 1993; Giachetti, 1998).

As methods have developed for finding the boundaries, it has become evident that despite their work, this is a challenging problem, not least because of the poor signal-to-noise ratio (SNR) of an ultrasound image. Further, the left ventricular cavity is often corrupted by clutter noise (general noise inside the left ventricle chamber), which differs from speckle noise, even though some of the properties of these echoes may sometimes be similar to those of speckle. The distinction is in the origin of the echoes. Clutter makes finding the endocardial border — which represents the blood tissue interface — a challenging problem.

Finding the epicardium — which represents a tissue-tissue interface — is even more difficult than finding the endocardium, because unlike the endocardium, there are no obvious features to look for. This explains the difficulty clinicians have in visually finding the epicardium. Segmenting the left ventricle, is one of the goals of this project. Since the functional dynamics of the left ventricle is crucial, we will also need to track the motion of the endocardium and the epicardium.

Ultrasonic image noise comes principally from three sources: speckle, first order reverberations and multiple reverberations (Schistad & Taxt, 1991). Speckle gives an image a granular structure and is a result of interference of ultrasound waves within each resolution cell being imaged. Speckle can be characterized as multiplicative noise because the average speckle signal corresponds to the average scattering of the object upon which speckle is superposed. Burckardt (1978) derives the distribution of speckle modelling tissue as a collection of scatters with different wavelengths. The speckle results from the coherent accumulation of random scatterings within one resolution cell (an area which represents the smallest observable resolvable detail). Speckle is not random in the sense that when the same static object is scanned twice under the same conditions, exactly the same pattern results (Wagner *et al.*, 1983; Wells & Halliwell, 1981). Several groups have considered methods for speckle reduction Wells & Halliwell (1981); Dickinson (1982); Bamber & Daft (1986); Bamber (1993) and Evans & Nixon (1996).

Moreover, the images can contain significant imaging artifacts, due to such things as, reverberations, side lobes, shadowing and enhancement and near field clutter (Feigenbaum, 1994, pages 22–26). Obtaining echocardiographic images is made more difficult because ultrasound does not pass through bone or air. This makes detecting subtle functional abnormalities of the cardiac chambers and valves more challenging.

To put the work in this thesis into context, currently, the state-of-the art Hewlett-Packard 5500 ultrasound machine (and that partly used in this thesis) provides an automatic boundary detection scheme called Acoustic Quantification (AQ) (Mor-Avi *et al.*, 1995). AQ is based on automated boundary detection (ABD) (Keren *et al.*, 1995; Kimball *et al.*, 1996; Pérez *et al.*, 1997) which discriminates the endocardial blood interfaces or boundaries using discontinuities in the radio-frequency (RF) signal returning from the tissue. Either the whole image or a region of interest (ROI) is selected. Pixels with discontinuities above a user-defined threshold are marked as boundary points. In practice, this is very sensitive to the choice of threshold, and hence the process depends on the judgement of different sonographers. Pixels labelled as boundary points are then joined together to form a connected boundary or boundaries (the latter because discontinuities in the RF signal occur all over the image) as shown in Figure 2.2.

Although some post-processing is done in the ultrasound machine, the AQ boundary appears rather jagged and coarse and takes no account of the fact that measurements taken from an image are noisy.

Moreover, although AQ is supposed to locate the endocardial border by connecting the labelled pixels, there is no guarantee that the connected boundary will in fact be the endocardial border of the left ventricle. Often, the connected boundary may extend into the left atrium or right ventricle, rendering measurements useless (Keren *et al.*, 1995).

More fundamentally, although AQ can sometimes give good results, it is very reliant on good image quality. In our work, we do not have the luxury of the RF signal. However, in comparison to AQ, the segmentation scheme we develop incorporates the a-priori information we have about the shape deformation and the motion of the left ventricle. AQ uses neither. This prior information makes our approach to segmentation more robust to noise and clutter in the image (Chapter 5). Moreover, we treat image measurements as the true underlying measurement corrupted with noise, and therefore the endocardial (and epicardial) border is estimated by fitting a smooth curve to the measurements.

Because AQ looks at the discontinuity in the radio-frequency (RF) signal, it is only feasible for it to find the endocardium (blood-tissue boundary). It cannot find the epicardium boundary, and hence myocardium. It therefore cannot track myocardial thickening, which is often of more interest to the clinician than the endocardial excursion. We address finding the epicardium in Chapter 7. Further, AQ only provides global measurements as part of the on-line automatic calculations. In Chapters 6 and 7 we shall show that we can provide regional measurements of wall excursion and myocardial thickening.




Figure 2.2: **Acoustic Quantification.** Acoustic Quantification (AQ) is based on automated boundary detection which discriminates boundaries using discontinuities in the radio-frequency (RF) signal returning from the tissue. Pixels with discontinuities above a user-defined threshold are marked as boundary points. Pixels labelled as boundary points are then joined together to form a connected boundaries. (a) Echocardiographic image. (b) Corresponding AQ image. Note how in the top figure, the connected boundary extends into the left atrium. In the bottom figure, the choice of threshold affects the labelled boundary.

Nowadays, to complement classic echocardiographic techniques, several newer approaches are being developed. For example, Doppler echocardiography, transesophageal echocardiography, stress echocardiography, contrast echocardiography, ultrasound tissue characterisation and three-dimensional echocardiography (Feigenbaum, 1994; Panidis, 1996).

In Chapter 8 we will apply the ideas developed in this thesis to stress echocardiography. Stress echocardiography aims to induce the appearance of a heart abnormality.

In stress echocardiography, the exercise is governed by a carefully defined protocol; this may be walking on a treadmill until a target heart rate is achieved, or it may be receiving precise amounts of a drug. Either way, the care and attention directed towards the quantitative nature of the protocol, is not matched by the analysis. This is essentially visual and thus qualitative.

This is because, no prior work has been able to easily quantify regional heart function. A test of this kind, seems to contradict the purpose of any experiment, where changes are made to the input variables of a process or system so that the reasons for the changes in the output variables may be observed. To provide accurate and reproducible analyses, both parts of the diagnostic process should be quantitative in nature.

We tackle this anomaly in Chapter 8, in the context of a small scale clinical study based on dobutamine stress echocardiography. To do this, we use the framework developed in Chapters 5 to 7 of this thesis for quantitatively assessing regional heart function.

Other imaging modalities have been used to image the heart, as shown in Table 2.1. For instance, nuclear medicine (such as Multi-gated acquisition (MUGA) scanning) (Harbert & Da Rocha, 1984; Juni & Buda, 1985), (fast) computed tomography (CT) (Brundage, 1985; Gorce *et al.*, 1997), positron emission computed tomography (PET) (Grover & Schelbert, 1985) and magnetic resonance imaging (MRI), of which there are two techniques, conventional MR (Aisen & Buda, 1985) and tagged MR (Kerwin & Prince, 1998).

However, we shall only concern ourselves with echocardiography. The practical reason for this is the data available to us. Throughout the project, we have worked very closely with cardiologists, at the John Radcliffe Hospitals in Oxford. By far and away, echocardiography is the most heavily used technique for cardiac imaging at the hospital.

Practicalities also dictate that this situation is mirrored up and down the UK. As imaging equipment goes, an ultrasound machine is relatively cheap at £100,000. Contrast this with an MRI machine of the order of £2million. Maintenance costs of the MRI machine also outweigh those for ultrasound. Further, ultrasound machines are more accessible than MRI machines, because there are considerably more of them. Ultrasound machines are also portable, which is important in emergencies, and they can be set up quickly. Unlike a MRI machine, an ultrasound machine can be used on patients with a pacemaker.

On a more theoretical level, it makes no sense to image a fast moving dynamic object, such as the mitral valve, with a modality with, at present, poor temporal resolution. More importantly, if you consider that an average cardiac cycle may last around one second, and the echocardiographer will typically look at each of the four standard views for around one minute (60 cycles), at present, this is totally impractical using MRI. It is also not viable for a hospital! Echocardiography also allows the clinician to observe many sequences of images in real time, as opposed to a sequence of ECG gated images, as in the case of MRI, MR tagging and nuclear imaging. To emphasize the point, hundreds of echocardiographic examinations are routinely done weekly at the JR hospital. Only a mere fraction of these would be possible using MRI. This situation gets worse, when you

Modality	Dimension	Nature	Observations
Echocardiography	2D + T, 3D + T	Anatomy	Ventricular function Cavities Valvular function
Ventriculography (with contrast agents)	2D + T, 3D + T		Function and anatomy of ventricles
Coronary angiogram (with contrast agents)	2D, 2D + T	Anatomy	Function and anatomy of arteries
Conventional MRI	2D, 3D, 3D + T	Anatomy Anatomy Function Anatomy	Ventricular anatomy Muscular movement Perfusion Great vessels
Tagged MRI	2D, 2D + T	Function	
CT	2D, 3D	Anatomy	

Table 2.1: **Cardiac imaging modalities.** Cardiac imaging modalities and their characteristics: the dimensions of the images, the nature of the modality and an example of the examination. Note that in the case of 3-dimensions (3D), the images are acquired by compounding a series of 2-dimensional images.

consider that the MRI machine takes a few minutes to set up (localisation of scans).

However, the spatial resolution of MRI is better than ultrasound, and the image does not suffer from the same imaging artifacts. More important, is that ultrasound images are operator dependent and subject dependent.

In practice, conventional MRI images of the heart are obtained by syncing the MR machine to an ECG of the patient; an image sequence is then built up from images acquired at known different phases of the cardiac cycle. The technique is often called 'ciné mode', because the heart is apparently visualized in real-time, although they do not have the time resolution of true real-time images. Moreover, as Webb (1988, page 425) points out, because lines of data are obtained in time at each cycle of the heart, they necessarily include an averaging effect due to the small changes in the pattern of the heart cycle during the course of the study. By definition, also the acquisition time is prolonged. One potential advantage offered by phase contrast MR is the added 3-dimensional velocity information available with each intensity image. Despite the problem with using MRI, recent interest has focused on using MR for analysing heart function (Meyer *et al.*, 1996; McEachen II & Duncan, 1997; Rueckert *et al.*, 1997; Weng *et al.*, 1997). Although these techniques offer some understanding of the heart dynamics (Stuber *et al.*, 1998), their clinical use is inherently limited.

In tagged MR, manipulation of the magnetic resonance pulse sequence selectively alters the magnetization of specific ventricular segments, causing them to display an intensity different from that of the surrounding myocardium. The tags move with the otherwise featureless tissue, to give a visual indication of motion in the images. Typically, tagged MR produces thin sheets of magnetically presaturated tissue orthogonal to the image plane. These tag surfaces result in a series of tag lines. The problem is then to generate motion estimates of the left ventricle using the tagged images (Park *et al.*, 1996a,b; Lee *et al.*, 1998; Kerwin & Prince, 1998). The advantage of tagged MR is that the corresponding points between frames are known using the tag lines.

There are, however, a number of problems associated with tagged MR: (1) it relies on the patient holding their breath during the acquisition; (2) the tagging itself, only lasts around half a second (not even one cardiac cycle) making feature detection a challenge; (3) tag lines are less and less clear as the sequence is acquired; this generally means that only images in systole are acquired.

An alternative to MRI and ultrasound is to use nuclear cardiology techniques (Huang *et al.*, 1997), such as multi-gated acquisition (MUGA) scanning (Harbert & Da Rocha, 1984). This technique generally uses vasodilators, such as dipyridamole (Picano & Lattanzi, 1991) or adenosine (Marwick, 1997). Vasodilators dilate the normal arteries at the expense of the obstructed arteries, resulting in blood being shunted towards the normal myocardium. The patient is then imaged with a MUGA scan, which is ECG-gated. By measuring blood counts (using a gamma camera) following injection of a radiopharmaceutical, information is supplied about cardiac function. This is then averaged over many cycles. Wall motion abnormalities can then be detected as the blood is shunted away from the obstructed arteries. The main advantage of MUGA scanning is that it is very accurate for measuring global left ventricle function, for example, the ejection fraction (Parker *et al.*, 1972). In particular, MUGA is more accurate than echocardiography at measuring global function, as it does not rely on unrealistic geometrical assumptions. This makes it ideal for recording global measurements of a patient over time; echocardiographic global measures are swamped by larger standard errors.

The most common type of radioisotope study of the heart is called ventriculography, which measures the output of blood from the left ventricle to the body. The disadvantage of MUGA is that it is not capable of imaging myocardial thickening. Further, the poor spatial resolution and sensitivity of the gamma camera means it is difficult to delineate accurately regions of interest. There are also problems with obtaining the correct positioning of the gamma camera.

During the past two decades, computed tomography (CT) has provided a further way of imaging the heart (Brundage, 1985; Gorce *et al.*, 1997). More recently, there has been interest in using 3-dimensional echocardiography to view the valves and chambers of the heart (Rohling *et al.*,

1997; Declerck *et al.*, 1998; Sheehan *et al.*, 1998; Sanchez-Ortiz *et al.*, 1999a)

In this chapter, we have reviewed the various cardiac imaging modalities, focusing on ultrasound which is the imaging modality used in this thesis. The remaining chapters explore how quantitative information can be obtained from echocardiography.

# CHAPTER 3

---

## Literature Review

---

This literature review is set out in two sections. The first section gives an overview of the different approaches used in 2-dimensional curve tracking in 2-dimensional visual image sequences. The second section gives a detailed review of key papers in the cardiovascular tracking and segmentation literature relevant to this thesis. The aim of this chapter, is to put the work in this thesis into context with the previous work in the field.

### 3.1 Visual tracking

Moran *et al.* (1997) points out: *“The field of [visual] tracking and data association was born out of a need to process the vast amount of information generated by the highly sensitive sensors found in modern surveillance systems.”*

Bar-Shalom & Fortmann (1988, page 4) define a track as follows: *“A track is a state trajectory estimated from a set of measurements that have been associated with the same target.”*

More generally, visual tracking is the recovery of the position of one or more objects (known as targets) viewed in a sequence of images over time. A target is tracked if its position is correctly identified in each image.

Broadly speaking, the literature on visual tracking can be approximately split into three classes of method: flow-based methods, feature-based methods and model-based methods.

Flow-based methods generally assume little or no prior knowledge of the object being tracked. The motion is tracked on a pixel by pixel basis, the aim being to find the image deformation which brings pixels originating from the same target into correspondence between images. The structure of the scene is then tracked from this visual flow field. Feature-based methods (Perona & Malik, 1990) use low level image features, such as corners (Noble, 1988). These features are then grouped according to higher level information about the scene. Both of these methods contrast to model based trackers, which use prior information about the object's shape, expected deformation, motion, intensity characteristics or other distinguishing attributes. The task of the tracking is then to estimate the parameters of the object model from the image sequence. Feature-based methods are not discussed at length here.

The advantage of representing an image by contour nodes (or control points) rather than image pixels, such as in Markov Random fields (Geman & Geman, 1984; Ripley, 1988, page 79), is that it incorporates the postulated simple connectivity, and drastically reduces the dimension of the configuration space. This approach is used in this thesis.

### 3.1.1 Flow-based tracking

In flow based tracking (Horn, 1986), the goal is to try to estimate image motion through the examination of the apparent motion of individual pixels.

Optic flow (Koenderink & van Doorn, 1975; Horn, 1986) has long been used as a way both to estimate dense motion fields and to segment areas of consistent flow into discrete objects. Optic flow relies on the fundamental assumption that provided the intensities of the textures of objects in the world are not changing, any change in the image over time will be related to the motion  $(u, v)$  of objects in the image and the spatial derivative of the image. Horn (1986) assumes that the total derivative of the intensity when moving with optic flow is then zero, which leads to the motion constraint equation given by Horn (1986),

$$\frac{\partial I}{\partial x}u + \frac{\partial I}{\partial y}v + \frac{\partial I}{\partial t} = 0. \quad (3.1)$$

In itself, Equation 3.1 does not allow complete recovery of the optic flow field,  $(u, v)$ , throughout the entire image. This is because, there is insufficient information at the edges of moving objects to locally determine if there is motion parallel to the edge — a phenomena referred to as the aperture problem (Horn, 1986; Nalwa, 1993, page 263). In order to solve Equation 3.1, additional constraints are required, for example, the motion of an object is parameterised using a low-dimensional model, for instance, an affine model. An alternate approach is to minimise an

error term based on the left hand side of Equation 3.1 summed over some region, with fixed  $(u, v)$  (Black & Anandan, 1993). Another constraint can be achieved by adding a regulariser to this error term, and assuming that the change in motion is smooth over the an image region as Horn (1986) does, for example. The disadvantage of this approach is that this smoothes object boundaries, which makes location of the discontinuities in the flow field harder. Nagel & Enkelmann (1986) relax the constraint in Equation 3.1, although the approach is still flow-based. A series of robust methods for determining optical flow have been developed, see for example Black & Anandan (1993).

### 3.1.2 Model based methods: snakes

In Computer Vision, important information about the scene in question is often conveyed by lines, curves, shading, texture and motion. Images are often processed to produce new images that are more desirable in some fashion, perhaps to increase the contrast between regions. For example, lines and curves manifest themselves as discontinuities in image brightness. One way to describe this 1-dimensional information is to use an active contour or snake. The term snakes was coined by Kass *et al.* (1987) to describe a parameterised family of curves (or surfaces in 3-dimensions) in the image plane which is deformed in such a way as to locally minimise an energy function.

In Kass *et al.* (1987), the features of interest were edges, and the energy which governed the evolution of the curves was a combination of an internal energy, which constrained curve flexure, and a global image gradient energy, which was maximal in the vicinity of edges. By choosing an appropriately defined energy term, one can select alternative features of interest, such as texture regions (Irvin & Porrill, 1995). There have been many modifications to the basic snake methodology. For example, to tackle the problem of occluded data, an *expansion term* may be introduced (Cohen & Cohen, 1990), which adds an expansive force in the normal direction along the curve. The contours may be used to track but have no model of the dynamics of image features and can only utilise positional continuity between frames.

Recently, Sclaroff & Isidoro (1998) proposed an extension to active snakes, called active blobs, which uses a region-based approach to non-rigid tracking. The method employs a texture mapped triangular mesh model for tracking deforming shapes in colour images. The non-rigid motion is modelled using a modal representation (Pentland & Horowitz, 1991), and the resulting energy function is minimised to find the registration between a template image and subsequent images.

### Active shape models

Active shape models (ASMs) were introduced by Cootes & Taylor (1992). Often termed, point distribution models (PDM's), the ASM is derived from a principal component analysis (PCA) of the shape variations of an object (or target) over a sequence of images (Cootes *et al.*, 1992). More precisely, corresponding landmark points are selected on a number of training examples from an image sequence. The PDM is then obtained by a PCA on this set of points. Strictly speaking, the same corresponding landmark points must be used for each training example. This poses more of a problem in 3-dimensions, although it can also be a problem in 2-dimensions. For example in the context of a heart, which twists over the cycle, there is no guarantee that the same point will be observable in different 2-dimensional image slices. In practice, using roughly corresponding points will suffice.

The PDM method has been applied to a number of applications, for example, hands (Ahmad *et al.*, 1995), echocardiographic images (Parker *et al.*, 1994; Cootes *et al.*, 1994), brains (Cootes *et al.*, 1994; Hill *et al.*, 1994; Caunce & Taylor, 1998) and faces (Cootes & Taylor, 1993a, 1995; Edwards & Taylor, 1998). Information about the local grey level structure of the target is incorporated into the model fitting procedure by building a simple statistical template of the grey level values along search lines perpendicular to the model (Cootes *et al.*, 1993). The idea of a PDM was extended to build a PDM of this intensity profile normal (Cootes & Taylor, 1993b; Byrne *et al.*, 1994). As Marais (1998, pages 85–87) points out, the problem with this method is that correspondences between profiles are difficult to find — “*as a consequence, the grey level PCA can admit solutions which bear little relation to the underlying anatomy*”. The position of the model is then located by searching for positions which minimise a Mahalanobis distance between the grey-level template and the image. Recently, the PDM has been extended into 3-dimensions by Caunce & Taylor (1998) and Marais (1998). This is a challenging problem because finding corresponding points in a 3-dimensional data set is difficult.

Baumberg & Hogg (1994a,b) build on the work of Cootes *et al.* (1994), discussing a way of automatically generating PDM's from a video sequence. The models are built by subtracting each image from a median filtered background image, and then thresholding to give a binary image of the target. Points are positioned equally around the object, and indexed based on their position relative to the principal axis of the target. Several of these points sets are collected over an image sequence. A PDM is then extracted by performing a PCA on these point sets. Baumberg & Hogg (1995) add learnt dynamics to the model, connecting parts of the PDM together with springs and dampers. The dynamical model allows the tracker to predict the motion of the object.

Snakes and active shape models deal with static image segmentation approach. We now turn to

segmentation approaches based on dynamical models of motion.

### Non-rigid model-based tracking

Many feature trackers assume a rigid or nearly rigid model of the motion of the target. Terzopoulos & Metaxas (1992) use a sophisticated 3-dimensional model of a non-rigid target. The model is based on deformable superquadrics for the global shape with local deformations away from the superquadric surface (Terzopoulos *et al.*, 1991).

Active contours, such as snakes, are applied statically to images. In dynamical applications, such as this thesis addresses, it is important to also build in any prior knowledge about the likely shape and deformations of the object. One way of doing this is to use a dynamic and adaptive contour (Blake & Isard, 1998) to model 2-dimensional non-rigid motion (Curwen *et al.*, 1991; Curwen & Blake, 1992; Blake *et al.*, 1993, 1995).

This improves on the original snakes formulation introduced by Kass *et al.* (1987) by using a B-spline model of the curve (Cipolla & Blake, 1990), which confers implicit smoothness constraints, and by incorporating a learnt model of the dynamics of the image feature (Curwen & Blake, 1992). Built into the dynamic framework, is prior knowledge about the 2-dimensional shape. This is achieved through the template mechanism coupled with the idea of a shape-space basis. An example of this invariant template mechanism is tracking space curves using a subspace based on affine stereo (Koenderink & van Doorn, 1991). Other subspaces, based on principal component analysis (PCA) (Krzanowski, 1988, page 53), enable the tracker to be more finely tuned to the deformation of the object; 3-dimensional rigid motion spaces are possible, too. The dynamical model which is used to guide a prediction of the motion of the object is learnt from a training sequence. This idea of learning models has been considered by several groups (Baumberg & Hogg, 1995; Yacoob & Davis, 1998a). Blake *et al.* (1993) introduce adaptive contours which uses a probabilistic framework giving automatic control of the spatial and temporal scale. This is achieved using a Kalman filter (Bucy, 1969; Gelb, 1974). The method has been applied to a number of applications such as, traffic monitoring (Ferrier *et al.*, 1994), lips (Dalton *et al.*, 1995; Kaucic, 1997), hands (Isard & Blake, 1996), faces (Bascle & Blake, 1998), automatic crop spraying (Wildenberg, 1997) and cardiac motion (Reynard *et al.*, 1994). In this thesis, we will apply the method to tracking endocardial wall excursion (Jacob *et al.*, 1999) and myocardial thickening.

An alternative approach to tracking is provided by Szeliski & Terzopoulos (1991) and Terzopoulos & Szeliski (1992). The elastic 2-dimensional contour model of the image feature is similar to Kass *et al.* (1987), but with added elastodynamics. This model may be used in tracking by deriving

the equations of motion which relate image forces to the motion of the contour.

The application for model-based tracking has grown exponentially over the recent years. For example, Sullivan (1992) and Paragios & Deriche (1998) consider surveillance of known objects, such as road traffic and aircraft servicing at airports. The accurate tracking of the motion of the human body has many applications (Yacoob & Black, 1998).

Application of telecommunications has prompted interest in facial tracking (Turk & Pentland, 1991; Phillips & Smith, 1994; DeCarlo & Metaxas, 1998; Jebara *et al.*, 1998; Bascle & Blake, 1998). The promise of audio-visual speech recognition for deaf and hearing impaired individuals and in noisy environments, has led to an interest in automated lip tracking (Kaucic, 1997; Kaucic & Blake, 1998; Basu *et al.*, 1998). Many tracking algorithms have been applied on moving (Isard & Blake, 1996; North, 1998) or static hands (Grenander *et al.*, 1991).

#### Data association and the Kalman filter

In visual tracking, an important issue is to decide how to assimilate the incoming measurements — called data association. For instance, the basic question of how to decide whether an observation belongs to an object before adding it into the state estimate. A review of data association techniques is given by Rao (1973) and Cox (1993). One method is to validate each measurement, for example, using a  $\chi^2$  test on the significance that the measured observation is the predicted observation. A similar technique is to defer any decision about the measurement until subsequent measurements are collected, which can be used to judge the measurement (Bar-Shalom & Fortmann, 1988, page 158).

A more probabilistic approach is offered by (Bar-Shalom & Fortmann, 1988, page 163) called the Probabilistic Data Association Filter (PDAF). The PDAF weights each observation with the probability that it originated from the target. The main advantage of the algorithm is that it requires almost fixed computational resources per cycle. The problem is that the PDAF assumes that the target probability distributions and the probabilities of obtaining false measurements are known, which is not often the case. Also if the tracker is attempting to track a manoeuvring target surrounded by dense clutter, the PDAF will give the clutter more weighting than the target, and tracking will be lost. As with any filtering system, the better the input (the measurements in the case of trackers), the better the output. One method discussed by Rowe & Blake (1996) is to build a statistical model of the profile normal to help locate the target. The work in Chapter 5 and 7 of this thesis is aimed towards getting more robust features.

Spatio-temporal estimation based tracking is now widely achieved using Kalman filtering (Terzopoulos & Szeliski, 1992; Lowe, 1992; Reynard *et al.*, 1996; Yacoob & Davis, 1998b). The

Kalman filter is an optimal linear filter when the following two conditions hold: (1) The process model is represented by a linear differential equation corrupted by additive Gaussian noise; (2) The measurements are a linear function of the unknown states corrupted by additive Gaussian noise. Under these conditions, the prior of the state is Gaussian and can be represented by a mean vector (the state estimate) and a covariance matrix, which models the uncertainty. The relative strengths of the measurements are assimilated using a Kalman gain (Bar-Shalom & Fortmann, 1988, page 58).

A technique developed for the case where the underlying dynamics of the system are non-linear is the Extended Kalman filter (EKF) Bar-Shalom & Fortmann (1988, page 106). The EKF, applied in computer vision by Harris (1992), for example, uses a Taylor series expansion to linearise the non-linear model of the dynamics, and ultimately approximates the state distributions as a Gaussian.

A more general tool for dealing with non-Gaussian densities is the Condensation filter of Isard & Blake (1998) and Isard (1998). The Condensation (conditional density propagation) algorithm uses a set of samples to model the prior of the state variables. Then at each timestep, samples are randomly chosen, and are diffused forward according to the state noise model. Measurements are then taken and used to gauge the sample set. The idea of using random sampling was introduced to computer vision by Geman & Geman (1984) on a pixel by pixel basis, for static image restoration using Markov random fields (Ripley, 1988, page 79). Grenander (1983) and Grenander *et al.* (1991) applied a slightly different approach to find hands in noisy images. Iterative simulation based on factored sampling (Grenander *et al.*, 1991, page 39) was applied to a template of the hand, rather than the pixels themselves. Using a template curve was also applied to the problem of finding spirals in galaxies (Ripley & Sutherland, 1990). The idea of sampling has been generalised and termed Markov Chain Monte-Carlo (MCMC) (Gelfand & Smith, 1990; Ripley, 1992; Smith & Roberts, 1993; Grenander & Miller, 1994; Besag *et al.*, 1995). For a mathematical treatment of random fields and Monte Carlo methods see Winkler (1995). For an introductory text see Tanner (1991).

### **Incorporating prior information**

While snakes enforce constraints on smoothness and the amount of deformation, they can not, in their original form, be used to constrain the types of deformation valid for a particular class of object. This led to development of algorithms that enforce a-priori constraints on the type of deformation allowed in the tracking (Blake *et al.*, 1993), deformable templates (Amit *et al.*, 1991; Yuille & Hallinan, 1992) and trainable snakes (Cootes *et al.*, 1992; Baumberg & Hogg, 1994b).

We shall see in Chapters 4 and 5 that representing shape deformation using a low-dimensional shape-space, enables tracking of non-rigid motion.

The B-spline representation of image contours is widely used (Menet *et al.*, 1990; Blake & Isard, 1998) in tracking literature, enabling prior information, for example, smoothness constraints to be incorporated into the tracker.

A survey of deformable models in medical image analysis is given by McInerney & Terzopoulos (1996).

One technique for incorporating prior information is based on deformable templates (Lipson *et al.*, 1990a,b; Yuille & Hallinan, 1992; Zhong *et al.*, 1998). Deformable templates represent an object as a template made up of parameterised geometrical parts and their relative positions. They can be thought of as a combination of the geometrical information of the object, an imaging model relating the template of the expected image intensities and an algorithm to match the template to the image data. An example, is a template of an eye. This has been modelled using a circle to represent the iris and two parabolae to represent the extent of the white of the eye above and below the iris (Phillips & Smith, 1994). The model is also able to incorporate the expected intensity within each region of the template and the expected edges. The model is fitted by minimising an energy function that is made up of shape deformation and intensity constraints.

Several groups (Blake *et al.*, 1993, 1995; Heap & Hogg, 1998) favour incorporating prior shape information using a shape basis, for example, an affine basis.

Another alternative, the PDM, as we discussed, was introduced by Cootes & Taylor (1992). A limitation of the PDM approach, is that it makes the assumption that the training points in some space can be adequately modelled by a Gaussian distribution, or can be roughly contained in a hyperboloid. For an object undergoing a non-linear deformation, this assumption breaks down. This results in large areas of the model shape-space representing invalid shapes or means that the PDM does not cover all valid instances in the class. There have been several non-linear extensions to the PDM, based on non-linear PCA (Sozou *et al.*, 1995), using polar co-ordinates for rotating parts of the model (Heap & Hogg, 1996) and a mixture model (Cootes & Taylor, 1997). Hierarchical PDM's (Heap & Hogg, 1997) based on work by Bregler & Omohundro (1994), cluster the training set into subregions, performing a PCA of each, to give a set of locally-linear patches (hyperellipsoids), the union of which describe the constrained shape-space. The technique has been successively applied to a training set based on a hand (Heap & Hogg, 1998).

Another approach used for underconstrained shape recovery incorporates a-priori information using physically-based knowledge about the material properties of the object (Terzopoulos *et al.*, 1991; Bookstein, 1989; Cohen & Cohen, 1990; Pentland & Sclaroff, 1991; Nastar & Ayache,

1993; Nastar & Pentland, 1995; Nastar *et al.*, 1996; McInerney & Terzopoulos, 1996).

Sclaroff & Pentland (1994) describe a physically-based solution for recovering, tracking and recognizing solid models from 2-dimensional and 3-dimensional data. Their technique, which is based on modal analysis (Pentland & Sclaroff, 1991), is applied to MR images of the brain.

### **Correlation-based tracking**

Bascle *et al.* (1994) and Bascle & Deriche (1994) proposed using correlation to compare an image patch to a small template patch. In unpublished work, the algorithm was applied to echocardiographic images of a fetus. Difficulties arise when there is a change in the object's size, shape, brightness or contrast alters over the sequence. Some robustness can be obtained by normalising with respect to a mean intensity and mean intensity variation — so called normalised correlation. Even allowing for this, the approach is found to be sensitive to noise and situations where the motion of the object is not smooth or predictable. Simple auto-adaptive templates can adjust as a non-rigid object changes, but they tend to 'drift' off the region of interest without some powerful underlying model of allowable change (Intille & Bobick, 1995).

## **3.2 Computer vision in cardiology**

Having considered some of the extensive literature on visual tracking, we now turn our attention to specific work on cardiac image segmentation and tracking.

### **Active shape models**

Parker *et al.* (1994) and Cootes *et al.* (1995) apply a PDM to an apical and short-axis echocardiographic (transoesophageal) sequence. In transoesophageal echocardiography, an ultrasound probe is positioned in the esophagus. Imaging from inside the body overcomes many of the limitations of imaging through the lungs and ribs. This means that images are less noisy than those acquired using transthoracic (the transducer probe is outside the body) and hence is an advantage in image search, when the PDM is fitted to the data.

The PDM is based on labelling points around an object, in this case the left ventricular boundary. In the case of the apical view, 96 points were used to label the left ventricle in each of the image frames closest to end-systole and end-diastole in 33 patients. By selecting only two time points in the cycle, it is possible that the model will be incapable of accommodating all instances in the class of possible deformations of the left ventricular boundary, although this is less time

consuming than to do all frames. For reasons of stability, points along the left atrium were also labelled (Cootes *et al.*, 1995). The problem with labelling an object not of direct interest, is that the PDM may simply learn its deformation, instead of the deformation of the object of interest, in this case the left ventricle boundary.

Cootes *et al.* (1994) note that the technique could be applied to 3-dimensional and time sequence images, although it is not clear exactly how. The PDM is also applied to track a sequence of echocardiographic images. Using a PDM to statically segment the boundary is a little simplistic: on each time frame, the PDM is re-initialised to its mean shape, and no account is taken of the temporal nature of the data, and hence any motion dynamics. The work in Chapters 4 and 5 of this thesis on learning the dynamics of the heart aims to overcome this problem.

### Snakes

Chalana *et al.* (1996) propose using a snake to segment the endocardial and epicardium boundaries using an echocardiographic sequences of the left ventricle. The epicardium is initially found using a snake, which is then used as a starting point in finding the endocardium using the same snake with different parameter values.

The energy equation for the epicardium snake is composed of: an internal energy term to control the bending and an external term to cause the contour to be attracted to areas of high gradient. The external term also included a temporal constraint which enforces monotonic motion, i.e., in the systole phase, the heart is monotonically decreasing in size, and, in diastole, the heart is monotonically increasing in size. Boundaries are located using a Canny edge detector (Canny, 1986). The results are compared to manual tracings obtained from 'expert-observers'. Applying a standard edge detector to an echocardiographic image sequences, is shown in Chapter 5 of this thesis, to be inferior to applying specific detectors.

Sonka *et al.* (1995) propose a method to segment intravascular ultrasound images based on minimising a cost-function. A user-defined region-of-interest (an ellipse) is used as the starting point in finding the arterial wall and plaque borders. Boundary points are detected by searching in the direction of the contour normal. A graph-searching technique is then used to minimise a cost-function through these points. Experimental results of comparisons between their method and an observer are very good. There are limitations however; for example they require that the arteries they seek to segment are not badly deformed, and that they are approximately elliptical.

A Bayesian approach to active contours is proposed by Storvik (1994) who uses simulated annealing (Ripley, 1987, page 181) to minimise an energy function (or maximise a posterior probability). The method is applied to an ultrasound image of the left ventricle, although it is very slow

and the contours a little jagged.

Ayache *et al.* (1992) describe a system for the automated analysis of 2-dimensional echocardiographic images. The aim is to track the mitral valve and left auricle. Images are filtered in the polar domain and then converted into the Cartesian domain. The boundaries are tracked using snakes. They go on to develop a method for pointwise correspondence between the contours in successive frames, which uses points of high curvature for matching. This is based on the paradigm that high curvature points usually possess an anatomical meaning, and are therefore good landmarks to guide the matching process. Reliably finding high curvature points in a sequence, is clearly a problem in itself.

Snakes have also been used to track the motion of myocardial tag lines from tagged MR data (Amini *et al.*, 199; Park *et al.*, 1996b; Kerwin & Prince, 1998; Amini *et al.*, 1998a,b; Young, 1998; Haber *et al.*, 1998). McVeigh (1996) provides a review of motion tracking techniques applied in MR and MR tagging.

Texture based snakes have been used to segment the ultrasound images (Tained *et al.*, 1994; Boukerroui *et al.*, 1998).

### Markov Random Fields

Herlin *et al.* (1992a,b) present a Markov Random Field model (Graffigne, 1986; Clifford, 1989) to segment the auricle cavity from echocardiographic images based on an initial segmentation based on morphology (Herlin & Ayache, 1992). The energy function is modelled using visual properties Herlin *et al.* observe from echocardiographic images; the distribution of grey level values within the cavity approximately follows a normal distribution; pixels belonging to the structure's boundary have a high gradient norm; and cardiac cavities have smooth boundaries, and hence the segmentation boundary should be smooth. The energy function is minimised using the Iterated Conditional Mode (ICM) method described by Besag (1986, 1989). ICM is a deterministic hill climbing algorithm. Therefore, the algorithm can be trapped in local minima. An alternative, stochastic algorithm, would be to use simulated annealing (SA) (Geman & Geman, 1984; Ripley, 1987, page 181), which can escape from local minima by allowing the algorithm to jump down, as well as up.

Two drawbacks of the model, addressed by Herlin & Giraudon (1993) and Herlin *et al.* (1994), are that their method assumes that grey level values exhibit some global homogeneity and secondly the method does not use any temporal information. The results tend only to be satisfactory, probably because MRFs are limited by their inherent simplicity.

Friedland & Adam (1989) also use a MRF approach to segment echocardiographic images. An

iterative algorithm was employed to find the centre-of-gravity of the left ventricle. Once found, radii are cast from this centre, and estimates of the wall position made along these lines. These estimates, and prior information, are embedded within an energy function, which is then minimised using simulated annealing. A similar idea is used by Dias & Leitão (1996). In their method, the optimization of the posterior probability function (or energy function) is so complex, another algorithm is proposed simply to solve this problem. This type of radii-based approach to segmentation of the myocardial boundaries has been commonly used (Chu *et al.*, 1988; Wilson *et al.*, 1993; Sonka *et al.*, 1995). Other approaches of using MRFs to segment ultrasound images are Ashton & Parker (1995); Rueckert *et al.* (1997) and Boukerroui *et al.* (1998).

### Shape-based tracking

Shape-based approaches to tracking the endocardial wall have been used by McEachen II & Duncan (1997), following previous work (McEachen II *et al.*, 1994; Shi *et al.*, 1994; McEachen II *et al.*, 1995). The approach, which is applied to 2-dimensional MR and ultrasound images, assumes that distinguishable shape landmarks can be consistently followed through a sequence of images. The initial correspondences between frames are found by comparing the shape of contour segments and choosing the match that minimizes a metric based on differences in curvature. These matches are incorporated into a smoothness term in an optimization functional. The trajectories of the matched points are validated against gold-standard markers implanted in the left ventricular wall. This showed that, after extended periods, the trajectories of these matches diverged from the trajectories of the markers, and that a temporal constraint would be needed to correct for this. To the author's knowledge, this idea has not been implemented.

Clarysse *et al.* (1997) present a shape-based approach to track the motion of the endocardial surface of the left ventricle from 3-dimensional X-ray images of a dog's heart. Clarysse *et al.* make the point that there are no explicit point-to-point correspondences, and so propose to track using geometrical descriptors of the surface. First a curvature description of the surface is derived which allows the global shape dynamics to be estimated. This is then updated using regional curvature based tracking.

Mailloux *et al.* (1987) determined the apparent velocities of the heart on 2-dimensional echocardiographic sequences using optical flow. They also describe a number of methods for analysing the resulting tracked contour information, such as the velocity flow-field between frames. One of the more interesting aspects of their approach is to analyse the regional left ventricular function by partitioning the left ventricle short-axis view into 12 segments. The average tangential velocity (radial, tangential and magnitude) is then plotted for each segment, based on Schiller (1985). We shall see in Chapters 6 and 7 of this thesis, that we are able to quantify regional left ventricular

wall excursion.

Giachetti *et al.* (1995) also applied optic flow techniques to 2-dimensional echocardiographic image sequences. It was also used to estimate the 3-dimensional endocardial wall motion from high-resolution CT data of a canine heart (Gorce *et al.*, 1997). In further work, Giachetti (1998) builds a clinical system to track echocardiographic sequences. This is achieved using a combination of techniques, such as snakes, model-based constraints, optic flow, and the regionalised analysis of Mailloux *et al.* (1987). To get around the fact that optic flow is badly affected by noise in an image, the shape is heavily constrained. No temporal information is used, which can be important in echocardiography. The reason being that the position of the boundaries is often more evident on a moving sequence of images, than on a static image. We shall address this in Chapter 5 when we model the dynamics of the left ventricle. Giachetti (1998) only considers tracking the endocardium, and no attempt is made to track the myocardium. Myocardial thickening, although somewhat more difficult to assess because the epicardial border is not often clear, is more important than endocardial wall excursion. We tackle this problem in Chapter 7, and quantitatively assess regional myocardial thickening.

Using optic flow methods on echocardiographic images is questionable. In theory, echocardiographic images break the fundamental assumption of optic flow. Namely, that changes in the intensity in an image  $I(x, y, t)$  are caused only by motion  $(u, v)$  of the objects with fixed texture. The intensity in an echocardiographic image, is dependent on the angle of incidence of the ultrasound beam on the wall, which differs along the boundary (Rose *et al.*, 1995).

Sanchez-Ortiz *et al.* (1999b) applied anisotropic diffusion to the problem of segmenting cardiac cine MR images. The information in each image is 4-dimensional, made up of an intensity image, and three phase images detailing the three perpendicular velocity components. Spatial and temporal a-priori knowledge about the shape and dynamic of the heart is incorporated into the diffusion process.

### 3-dimensions

The work in this thesis considers 2-dimensional echocardiographic images. The following work deals with 3-dimensional images of the left ventricle.

Chen *et al.* (1994) model the deformation and motion of the left ventricle based on angiographic data, using a hierarchical coarse-to-fine decomposition. The global shape is modelled using a superquadric. The model is then updated using local deformations based on the residuals of the global model to the data. The results are less than impressive, and are more akin to something from computer graphics than medical imaging.

Modelling the global shape of the left ventricle with superquadrics and then refining the shape locally was also used by Bardinet *et al.* (1996).

Recovering the shape of the left ventricle was considered by Coppini *et al.* (1995). The basic shape of the ventricle is modelled using a thin elastic surface. Data points are treated as radial springs acting on this initial surface. Minimisation of an energy equation then leads to the final surface of the left ventricle. The data is based on an edge-map. This edge-map is derived in two phases. Initially a standard edge detection scheme is applied to the 2-dimensional echocardiographic images. Unsurprisingly (because of the noise in the image), this resulted in a noisy edge-map, with many unconnected lines. A neural network was then applied to this edge-map, to obtain a smaller edge-map, which was supposed to represent the left ventricle boundary. This map then formed the input into another neural network to recover the surface of the left ventricle. Among the weaknesses of the method are, the use of a standard edge detector to find the boundaries, and the use of positional information in the network. It is also unclear how the largely unconnected edge-map, can have much effect on the initial model, which probably explains the final shape being very elliptical.

Other approaches to recovering the 3-dimensional shape of the left ventricle are Azhari *et al.* (1987); Park *et al.* (1994, 1996a); Chen & Suter (1997); Rohling *et al.* (1997); Declerck *et al.* (1998); Sheehan *et al.* (1998).

Although the work in this thesis is directed to 2-dimensional echocardiography, the ideas generalise easily to 3-dimensions. This allows us the possibility of using the approaches described in the context of 3-dimensional echocardiography. Let us now turn to the modelling framework for tracking.

# CHAPTER 4

---

## Shape-spaces and Dynamical Tracking

---

This chapter details the active contour tracker framework that forms the basis for the ideas developed in this thesis. This framework was originally developed by Blake *et al.* (1993) and extended by Curwen (1993). Recent extensions were undertaken by Blake *et al.* (1995); Rowe (1996); Kaucic (1997); Wildenberg (1997); Isard (1998) and North (1998). The background material in this chapter is taken largely from Blake & Isard (1998).

The performance of an active contour is highly dependent on three key ingredients: firstly, an appropriate model of shape; secondly, modelling and learning the object's motion to guide a prediction; and, thirdly, using specific image feature detectors to enhance the measurement process, which in turn updates the contour's position.

Section 4.1 of the chapter begins by introducing the notation used to define curves in images. The three key ingredients for tracking are outlined briefly in Section 4.2. The main idea behind the first ingredient, modelling shape variation using 'shape-spaces', is described in Section 4.3. Shape-spaces can be taken off the 'shelf', or hand built. However, it is far more effective and less time-consuming to learn them probabilistically from a training set using principal component analysis. Section 4.4 explains this method. We shall see in Chapter 5 that building in finely tuned prior shape information is important for successful cardiac tracking.

The second and equally important ingredient is a model of the object's motion. This is discussed in Section 4.5, where we propose the model to be a second order autoregressive model. This is

used to predict the position of the contour on any given frame based on knowledge of the two previous frames. As with the model of shape variation, the motion model is learnt from a training set.

The third key ingredient to successful tracking is to use a measurement model to update the contour's predicted position. This is discussed in Section 4.6. We conclude by unifying the shape, motion and measurement models into a Kalman Filter framework suitable for tracking the left ventricle.

## 4.1 B-splines as tracking templates

Throughout this thesis, curves are represented by parametric B-spline curves (Bartels *et al.*, 1987) which have been used for example by Cipolla & Blake (1990) and Menet *et al.* (1990).

In this section, we briefly outline the basic ideas behind splines of relevance to this thesis. For a more detailed account the reader is directed to Bartels *et al.* (1987); Wahba (1990); Härdle (1990) and Green & Silverman (1994).

A 2-dimensional B-spline is a curve  $\mathbf{r}(s) = (x(s), y(s))$  in which  $s$  is a parameter that increases as the curve is traversed, and  $x$  and  $y$  are function of  $s$  called splines. A spline is made up of  $L$  piecewise polynomial segments or *spans* — quadratic (order 3) were used in this work — joined together at *breakpoints*.

A spline function is constructed as a weighted sum of  $N_B$  *basis-functions*  $B_n(s)$ . In our case, each basis function consists of 3 (the order) polynomials defined over a span of the  $s$ -axis. For simplicity, we assume that each span has unit length. The spans are joined at *knots*. For each basis function, a *control point*  $\mathbf{q}_n = (q_n^x, q_n^y)^\top$  is defined — we will assume that the spline has  $N_B$  control points — and the curve is a weighted sum of control points,

$$\mathbf{r}(s) = \sum_{n=0}^{N_B-1} B_n(s) \mathbf{q}_n \quad \text{for } 0 \leq s \leq L. \quad (4.1)$$

The component functions of  $\mathbf{r}(s)$  are then given by,

$$x(s) = \sum_{n=0}^{N_B-1} B_n(s) q_n^x, \quad (4.2)$$

and similarly for  $y(s)$ .

The configuration of the contour can be described entirely in terms of its control points' positions. A configuration column vector  $\mathbf{Q}$  is defined in a (spline) space  $S_Q$  of *control vectors*  $\mathbf{Q}$  containing

all the control point co-ordinates — the  $x$  co-ordinates ( $\mathbf{Q}^x$ ), then all the  $y$  co-ordinates ( $\mathbf{Q}^y$ ). That is,

$$\mathbf{Q} = \begin{pmatrix} \mathbf{Q}^x \\ \mathbf{Q}^y \end{pmatrix} \quad \text{where, } \mathbf{Q}^x = \begin{pmatrix} q_0^x \\ q_1^x \\ \vdots \\ q_{N_B-1}^x \end{pmatrix}, \quad (4.3)$$

and similarly for  $\mathbf{Q}^y$ .

The co-ordinate functions can then be written as,

$$x(s) = \mathbf{B}^\top(s)\mathbf{Q}^x,$$

and similarly for  $y(s)$ , so that, the curve  $r(s)$ , is given by,

$$r(s) = \mathbf{U}(s)\mathbf{Q}, \quad \text{for } 0 \leq s \leq L, \quad (4.4)$$

where,

$$\mathbf{U}(s) = \begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix} \otimes \mathbf{B}^\top(s) = \begin{pmatrix} \mathbf{B}^\top(s) & \mathbf{0} \\ \mathbf{0} & \mathbf{B}^\top(s) \end{pmatrix}, \quad (4.5)$$

a matrix of size  $2 \times 2N_Q$ , where  $N_Q = 2N_B$ , is the dimension of spline space. The operator  $\otimes$  represents the Kronecker product (see Appendix D.1.1) of two matrices.

When using parametric spline curves, it is useful to define a norm in order to measure how closely one curve approximates another. In this work, the  $L_2$ -norm,  $\|\mathbf{Q}\|$ , in *spline* space is used which is equivalent to the root mean square distance in the image plane:

$$\|\mathbf{Q}\|^2 = \frac{1}{L} \int_{s=0}^L \mathbf{r}^2(s) ds, \quad (4.6)$$

which can be written as,

$$\|\mathbf{Q}\|^2 = \mathbf{Q}^\top \mathbf{H} \mathbf{Q}, \quad (4.7)$$

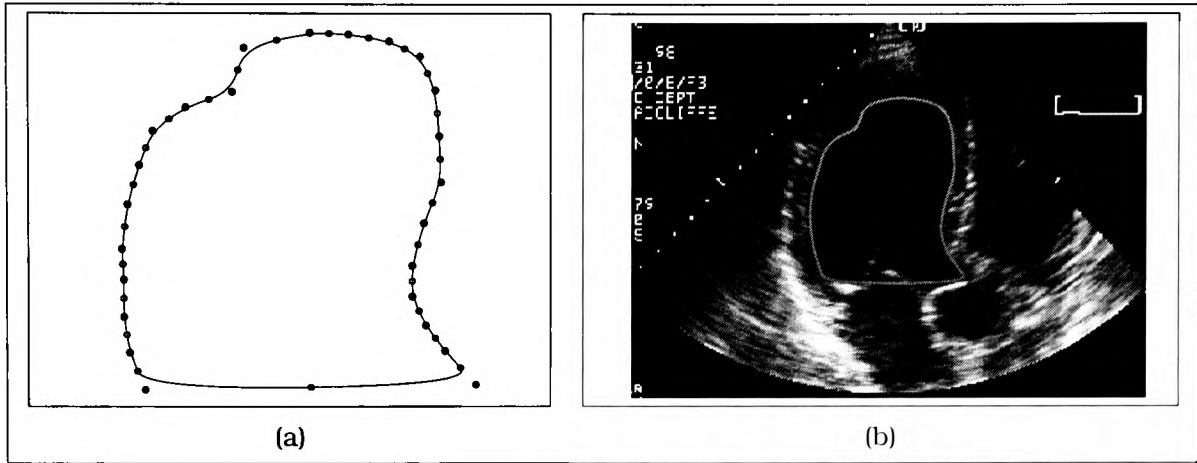


Figure 4.1: **B-Spline template.** A quadratic approximating B-spline with 24 control points used to model the endocardial boundary. Since an approximating B-spline is used, the control points (●) do not necessarily lie on the curve. The spline is made up of spans joined at knots (○). We shall build in knowledge of the shape deformation of the object into this template. (a) Schematic representation. (b) A B-spline template superimposed on an ultrasound image.

where,

$$\mathcal{H} = \begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix} \otimes \frac{1}{L} \int_{s=0}^L \mathbf{B}(s)\mathbf{B}(s)^\top ds. \quad (4.8)$$

Note that  $\mathcal{H}$  is a metric matrix converting between distances in spline space to image space.

We incorporate knowledge of the shape of the object by specifying a ‘typical’ configuration. This is usually a hand-drawn B-spline that is fitted by drawing around a single image frame of the endocardial wall that forms the blood tissue interface. This is referred to as the *template* and its control point configuration will be written as  $\mathbf{Q}_0$ . An example of a template used to model the endocardium boundary is shown in Figure 4.1.

## 4.2 Tracking strategy

As we shall see in the following sections, the basis for tracking is one of *prediction*, *measurement* and *update*. We shall use a B-spline contour to represent the boundary of the object. Our aim is for this contour to, as accurately as possible, track the shape and motion deformation of the object as it moves over time. We shall develop a:

- Model of the object’s deformation (or shape),
- Predictive learnt model of the motion of the object, and,
- Measurement model to update the prediction.

We firstly turn our attention to modelling the deformation of the object using a *shape-space*.

### 4.3 Shape-spaces

A 2-dimensional B-spline, with  $N_B$  control points, has dimension  $N_Q = 2N_B$ . Using this full space allows for arbitrary variations in control points over time. In our case, for example, this would allow maximum flexibility in deforming to the various left ventricular boundary shapes. However, tracking spaces which possess more degrees of freedom than are necessary are undesirable for two reasons.

- **Computational cost:** The computational cost of the tracking algorithm is  $\mathcal{O}(N_Q^3)$ , where  $N_Q$  is the dimension of the shape-space (Blake & Isard, 1998, pages 132 and 200)— the computational penalty for employing unnecessarily large spaces can be severe. Hence for computational sake we should chose to work in a considerably smaller dimensional space,  $N_X$ , which is a linear subspace of the full space. We use a technique known as subspace projection. For example, if 24 control points are needed to model the left ventricular wall boundary, the full space has dimension 48 — our aim is to work in lower dimensional space.
- **Stability:** A more important reason is that the additional non-essential degrees of freedom can lead to unstable tracking (Blake *et al.*, 1993). In cardiac tracking, this is particularly important because of the complexity of deformations.

The key is to exploit constraints imposed by the application. We explore this next.

#### Limiting the allowable shapes

Cardiac data has the nice property that *in an ultrasound image, we know what to expect, and consequently what to look for*. It is therefore a constrained environment. In very simple terms, a left ventricle can be described as being roughly elliptical (Friedland & Adam, 1989), so why even consider looking for a “square” shaped object ? This is the idea behind a shape-space — take advantage of the shape knowledge available, however coarse it is. This is precisely what Bayesian statisticians do when they combine data with a *prior* model, and seek to maximise the resulting posterior density function.

Shape-spaces restrict the shape of the object. In doing so they don't affect its ability to follow the boundary of the ‘normal’ object, but will prevent it from forming impossible shapes. This is important, because echocardiographic images are very noisy and prone to artifacts. Curwen (1993) and Blake *et al.* (1993) showed that allowing a tracker complete freedom to assume *any* shape can lead to unstable tracking. This is one of the major advantages of the (dynamic) active

contour paradigm over snakes as introduced by Kass *et al.* (1987) which allow too many degrees of freedom

A further advantage is that it is only necessary to estimate  $N_X$  parameters, instead of the  $2N_B$  free parameters of the curve's control points. For example, an affine shape-space has only 6 parameters, which, as we shall see, can all be defined from the initial template.

Let us now turn to a more formal definition of a shape-space.

### 4.3.1 Definition of a shape-space

A shape-space  $\mathcal{S}_X = \mathcal{L}(W, Q_0)$  is a linear mapping of a "shape-space vector"  $X \in \mathcal{R}^{N_X}$  to a spline vector  $Q \in \mathcal{R}^{N_Q}$ :

$$Q = Q_0 + WX, \quad (4.9)$$

where,  $W$  is a  $N_Q \times N_X$  "shape-matrix".  $Q_0$  is a constant offset, or initial template curve, of dimension  $N_Q \times 1$ . The elements of  $X$  ( $X$  has dimension  $N_X \times 1$ ) act as weights on the columns of  $W$ . Define the matrix  $W$  to have  $w_1, w_2, \dots, w_n$  as its columns.

Conversely, we can define a mapping from the spline space,  $\mathcal{S}_Q$ , to the shape-space,  $\mathcal{S}$ . Providing that  $W$  has full rank, a pseudo-inverse can be defined:

$$X = W^\dagger(Q - Q_0), \quad (4.10)$$

where,

$$W^\dagger = \mathcal{H}^{-1}W^\top \mathcal{U} = (W^\top \mathcal{H}W)^{-1}W^\top \mathcal{H}, \quad (4.11)$$

and,  $\mathcal{U}$  is metric matrix for B-spline parametric curves, where,

$$\mathcal{U} = I_2 \otimes \mathcal{B}. \quad (4.12)$$

A formal proof is given in Appendix E.1. We can view  $W^\dagger$  as an operator that minimises the projection error onto the shape-space. This notion is important, because when we fit a spline to some data, the pseudo inverse operator is used as a regulariser.

Turning back now to the notion of a shape, we shall be interested in the basis of a shape-space for the underlying vector space. The matrix  $W$  is comprised of columns which are the vectors of this basis.

The Euclidean similarity and affine spaces are two examples of vector spaces that can be used as a shape-space. The shape-space of Euclidean similarities has 4 degrees of freedom, allowing independent translation (horizontally and vertically), rotation and isotropic scaling. The 6 degree of freedom affine space allows independent translation (horizontally and vertically), independent scaling (horizontally and vertically) and two shearing motions.

Clearly, the Euclidean similarity space is not going to be able to model the complex non-rigid deformations of the left ventricle. This is because no shape changes are allowed in the Euclidean similarity space, yet the heart wall does change shape. More likely, is that the affine shape-space, with its two extra degrees of freedom, and more importantly independent scaling, that may be able to model the deformation. We will now consider constructing an affine shape-space.

### Affine shape-space

The shape-space matrix,  $W$ , defining an affine shape-space for a template  $Q_0 = (Q_0^x, Q_0^y)^T$ , is, given by,

$$W = \begin{pmatrix} 1 & 0 & Q_0^x & 0 & 0 & Q_0^y \\ 0 & 1 & 0 & Q_0^y & Q_0^x & 0 \end{pmatrix}. \quad (4.13)$$

One advantage of the affine space, as Equation 4.13 reveals, is that it can be directly constructed from the template. The 6 degrees of freedom of an affine space are illustrated in Figure 4.2.

A significant advantage of the affine space over other spaces, such as PCA (which we consider in Section 4.4), is that each component of the space is readily interpretable. For instance, the first basis vector horizontally translates the template, and the third horizontally scales the template. The space is rather appealing prompting its use in early applications of tracking (Blake *et al.*, 1993). Its usefulness stems from the fact that it can describe the shape variability of a planar object being imaged by an affine camera (Mundy & Zisserman, 1992). The affine camera is a good model of a real camera when the depth of the objects in the scene is small when compared to the distance from the objects to the camera. The main disadvantage of using the affine basis is that because it is very general, it is not finely tuned to the object's deformation, as demonstrated in Section 5.1.3

Despite this, and although in 3-dimensions the heart moves in a complex non-rigid manner (Bardinet *et al.*, 1996), several groups have modelled the 2-dimensional shape variability using an affine model (Ayache *et al.*, 1992; Giachetti, 1998). We shall see later in Section 5.1.3, that the performance of the affine model in tracking the endocardial boundary is not very good. This

is probably because the heart moves in a complex non-rigid manner, and an ultrasound image of the left ventricle is a 2-dimensional projection of a 3-dimensional object.

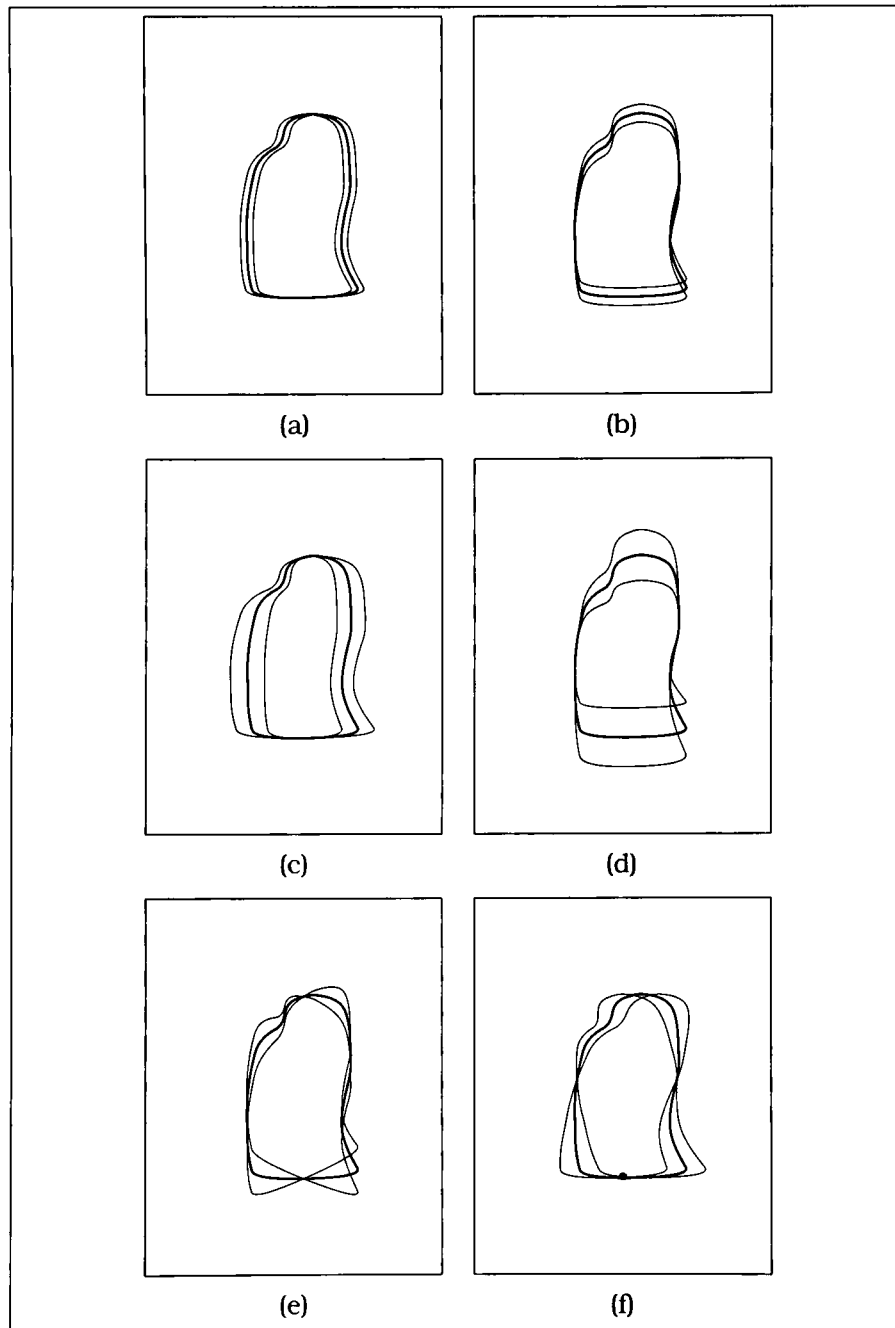


Figure 4.2: **Affine deformation.** A template (thick line)  $Q_0 = (Q_0^x, Q_0^y)$  undergoing an affine deformation (thin lines). Each basis vector independently deforms the template by an arbitrary amount. The 6 degrees of freedom are: (a) horizontal translation (b) vertical translation (c) horizontal scaling (d) vertical scaling (e) horizontal shearing and (f) vertical shearing.

### Building general shape-spaces via key-frames

As we have stated, the motion of the left ventricle is highly non-rigid. In the absence of any prior simple analytical description of the motion, it is equally effective to learn shape variation from

examples of 'normal' motion.

At a more simplistic level, a shape-space can be built by a linear combination of "key-frames", which it is assumed will span the full space of all possible deformations. As an example, facial pose could be modelled using a linear combination of a few basic facial expressions, such as, "neutral", "smile", "surprise" and "disgust" (Bascle & Blake, 1998). Alternatively, additional key-frames could be added to the affine space, generating a larger basis.

Although, like the affine space, construction of a shape-space using key-frames results in basis vectors corresponding directly to known deformations (they have clear interpretations), often the resulting spaces are unnecessarily large. This can cause two problems outlined at the start of Section 4.3.

At a more practical level, without any prior or anatomical knowledge, what and how many "key-frames" should be used as part of the basis? Using too few key-frames leads to a tracker that may fail because it is overly restrictive. Using too many results in a tracker that is unstable because it has too many degrees of freedom (Blake *et al.*, 1993). This direct tuning of a shape-space model,  $W$ , to the deformation is important. Principal component analysis (PCA) provides one remedy to the disadvantages of key-frames. We consider this in the next section.

#### 4.4 Probabilistic shape-spaces via Principal Component Analysis (PCA)

As we said above, we can use knowledge of the shape of the left ventricular boundary to impose a constraint on its deformation. This knowledge is turned into a prior, as in Bayesian statistics. It is this prior that we seek to learn as a shape-space.

Suppose that we were shown a training set of example contours,  $Q_1, Q_2, \dots, Q_M$ . Correspondingly,  $X_i$  is an example vector from the shape-space of interest  $S = \mathcal{L}(W, Q_0)$ . The training set in shape-space is given by,  $X_1, X_2, \dots, X_M$ . For example,  $Q_i$  may be a contour from a clinician's manual segmentation (as in Chapter 5), and  $X_i$  its shape-space representation. These examples should capture as much as possible of the variation in the space  $S_Q$ . Our goal is to determine a smaller shape-space,  $S' = \mathcal{L}(W', Q'_0) \subset S$ , that is a subspace of  $S$  that spans at least approximately, all of the shapes in the training set. These specific priors can be learnt using principal component analysis (PCA) (Jolliffe, 1986; Krzanowski, 1988, page 53).

Many authors have considered incorporating a-priori shape information (Bookstein, 1989; Shen & Hogg, 1994; Cootes & Taylor, 1995; Baumberg & Hogg, 1995; Heap & Hogg, 1996). Pentland (Pentland & Sclaroff, 1991) and Terzopoulos *et al.* (1991) considered using physical knowledge about the material properties of the object being modelled. Nastar & Ayache (1993) used modal

analysis for nonrigid tracking. Modal analysis changes the basis from the original modelling function to the eigenmodes of the deformation matrix. The most dominant deformations can then be characterised by the few largest eigenmodes, reducing the dimensionality of the object descriptor space substantially. Székely *et al.* (1996) modelled the shape variability of the corpus callosum using Fourier snakes. More recently, Heap & Hogg (1998) proposed a method for incorporating temporal shape discontinuities explicitly. Allowable shapes are represented as a union of learned bounded regions in a shape space. They term their approach 'wormholes in shape-space'. They note this overcomes the existing assumptions of tracking algorithms, which rely on the objects moving and deforming smoothly over time.

PCA was first introduced by Pearson (1901) and developed by Hotelling (1933). The central idea of PCA is to reduce the dimensionality of a data set in which there are a large number of interrelated variables while retaining as much as possible of the variation present in the data set. This reduction is achieved by transforming to a new set of variables, the principal components, which are uncorrelated, and which are ordered so that the first few retain most of the variation in all of the original variables.

This idea was introduced in the computer vision literature as a shape representation by Cootes & Taylor (1992), to construct a "Point Distribution Model" (PDM). Since then, a number of authors have used PCA in various contexts: echocardiography (Parker *et al.*, 1994), brain MR (Cootes *et al.*, 1994; Ruff *et al.*, 1997; Caunce & Taylor, 1998), modelling facial features (Bascle & Blake, 1998; Edwards & Taylor, 1998), modelling hands (Ahmad *et al.*, 1995; North & Blake, 1998) and modelling lips (Dalton *et al.*, 1995; Kaucic *et al.*, 1996; Kaucic, 1997).

#### **Principal Component Analysis (PCA) motivation**

The key idea behind PCA, is as follows. Suppose we have a  $p$ -dimensional set of data obtained from  $n$  individuals. We seek the 'best'  $r$  ( $\leq p$ )-dimensional representation of the data. The first principal axis is the line that maximises the spread of the  $n$  points when projected onto this line. Having obtained this, the second principal axis is the line in the  $(p - 1)$ -dimensional subspace, orthogonal to the first, that maximises the spread of the  $n$  points when projected onto this line. And so on.

The next few sections detail some of the relevant mathematical background to PCA. This material is, in part, taken from Blake & Isard (1998).

### Classical PCA

The classical approach to PCA allows the following version of the subspace approximation to be solved. Given a training sequence,  $\mathbf{X}_1, \mathbf{X}_2, \dots, \mathbf{X}_M$ , solve:

$$\min_{W', \mathbf{Q}'_0, \mathbf{X}'_1, \dots, \mathbf{X}'_{N'_X}} \left( \sum_{k=1}^M |\mathbf{Q}_k - \mathbf{Q}'_k|^2 \right), \quad (4.14)$$

where,

$$\mathbf{Q}'_k = \mathbf{Q}'_0 + W' \mathbf{X}'_k, \quad \text{and,} \quad \mathbf{Q}_k = \mathbf{Q}_0 W \mathbf{X}_k, \quad (4.15)$$

is the original data. The distance measure  $|\cdot|$  is the Euclidean norm, that is,  $|\mathbf{Q}|^2 \triangleq \mathbf{Q}^\top \mathbf{Q}$ . The solution to this problem (Blake & Isard, 1998, page 183) gives,  $\mathbf{Q}'_0 = \bar{\mathbf{Q}}$ , the mean of the training sequence. The columns of the matrix,  $\mathbf{W}'$ , are the first  $N'_X$  of the orthonormal eigenvectors of the covariance matrix,

$$\Sigma = \frac{1}{M} \sum_{k=1}^M (\mathbf{Q}_k - \bar{\mathbf{Q}})(\mathbf{Q}_k - \bar{\mathbf{Q}})^\top, \quad (4.16)$$

ordering the eigenvectors  $\mathbf{v}_i$  in descending order of their eigenvalues,  $\lambda_i$ .  $\Sigma$  is the covariance matrix, and necessarily has non-negative eigenvalues. In fact to be truly unbiased, the divisor is taken as  $\frac{1}{M-1}$ . The first  $N'_X$  eigenvectors form a basis for the subspace of dimension  $N'_X$  that explain as much as possible of the variance in the training set. Since the rank on the covariance is  $p = \text{Min}(M, N_X)$ , there will be at most  $\text{Min}(M, N_X)$  non-zero eigenvalues of  $\Sigma$ .

The proportion of the variance accounted for by a particular eigenvector,  $\mathbf{v}_k$  is given by,

$$\frac{\lambda_k}{\sum_{i=1}^p \lambda_i} \quad (4.17)$$

### Standardisation and the $L_2$ norm

PCA is scale dependent, since orthogonal projection is scale dependent. Statisticians have long argued whether the data should be standardised prior to PCA. Generally, standardisation is simply effected by considering each column (variate), and subtracting off its mean, and dividing by its standard deviation.

One problem with the classical PCA method outlined above is that it maximises the fit of one set of the control points of the data to the control points of the vectors in the  $\mathbf{W}$  matrix. However, in the same way as it was necessary to use the metric matrix  $\mathcal{H}$  when converting from control point space to state space, it is necessary to use the metric matrix when computing a PCA. In

practice this means that we use the  $L_2$  norm, rather than Euclidean, to compute the principal components.

The distance measure in the Euclidean norm is defined as  $|\mathbf{Q}|^2 \triangleq \mathbf{Q}^\top \mathbf{Q}$ . The  $L_2$  norm is defined by  $\|\mathbf{Q}\|^2 \triangleq \mathbf{Q}^\top \mathcal{H} \mathbf{Q}$ . The approximation problem for PCA using the  $L_2$  norm then becomes,

$$\min_{\mathbf{W}', \mathbf{Q}'_0, \mathbf{X}'_1, \dots, \mathbf{X}'_{N'_X}} \left( \sum_{k=1}^M \|\mathbf{Q}_k - \mathbf{Q}'_k\|^2 \right). \quad (4.18)$$

The solution to Equation 4.18, which can be derived from the solution to the classical problem, is that  $\mathbf{Q}'_0 = \bar{\mathbf{Q}}$  as before, and  $\mathbf{W}'$  is a matrix whose columns are the first  $N'_X$  eigenvectors of the matrix  $\mathcal{U}$

$$\Sigma = \frac{1}{M} \sum_{k=1}^M (\mathbf{Q}_k - \bar{\mathbf{Q}})(\mathbf{Q}_k - \bar{\mathbf{Q}})^\top \mathcal{H}, \quad (4.19)$$

where,  $\mathcal{U}$  is the symmetric square root of  $\mathcal{H}$ .

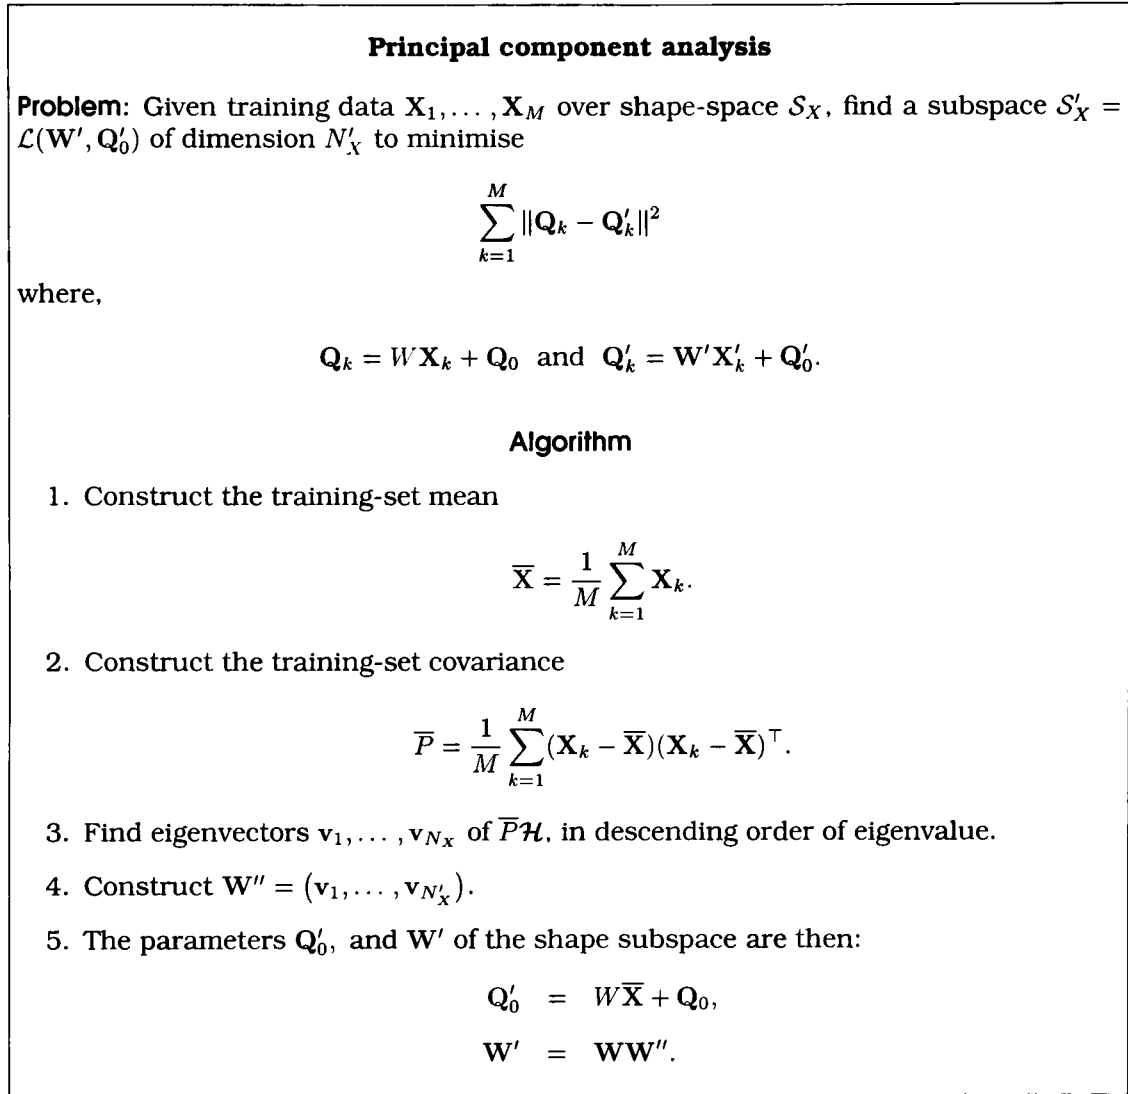
Essentially the PCA problem is, given a training sequence  $\mathbf{X}_1, \mathbf{X}_2, \dots, \mathbf{X}_M$ , find a template vector  $\mathbf{Q}'_0$  and a shape matrix  $\mathbf{W}' = (\mathbf{v}_1, \mathbf{v}_2, \dots, \mathbf{v}_n)$  such that the reconstructed sequence,  $\mathbf{X}'_1, \mathbf{X}'_2, \dots, \mathbf{X}'_M$  approximates most closely the training sequence in a least-squares sense. The resultant shape-space  $\mathcal{S}'_0 = \mathcal{L}(\mathbf{W}', \mathbf{Q}'_0)$  is then spanned by the basis vectors  $\mathbf{v}_i$  ( $= \mathbf{w}_i$ ) of the shape matrix  $\mathbf{W}$  plus an offset template  $\mathbf{Q}'_0$ . The full algorithm is described in Figure 4.3.

Having developed a model of the deformation of the object, we now consider obtaining a predictive model of the motion of the object. This was the second ingredient of our basis for tracking, given in Section 4.2.

## 4.5 Predictive Dynamical Model

We now turn our attention to establishing a framework for tracking curves in images. Szeliski & Terzopoulos (1991) raised the possibility of tracking curves within a Kalman Filter framework, and this was taken up by Blake *et al.* (1993). We adopt the approach of Blake & Isard (1994); Blake *et al.* (1995); Baumberg & Hogg (1995) based on *prediction, measurement and update*.

The central idea is based on recognising that an object moves coherently over time. This is important, because it means that we can use the position of the fitted contour on any frame as an initial estimate for its position in the next frame. We need then only 'look forward' from this initial estimate to find its new position. This repeated regularised fitting procedure is, however, rather crude — we can do better by using knowledge of the object's likely motion.

Figure 4.3: Algorithm for  $L_2$  PCA in shape-space

On any frame, the position of the contour on the next frame can be predicted — the *prediction* stage of the algorithm. We can then ‘look forward’ from this predicted position to find desirable image features — the *measurement* stage. We can then fit a new contour to these features weighting our belief in the prediction and measurements — the *update*. This circle of operations is illustrated in Figure 4.4.

This two-phase tracking approach is now fundamental to the tracking practice in North & Blake (1998); Bascle & Blake (1998); Isard & Blake (1998). In this work, the prediction is based on a second order auto-regressive process, which specifies the likely motion of the object. A measurement model specifies how the object is observed. The measurements are incorporated into an estimated target position using a Kalman Filter (Gelb, 1974; Harvey, 1989). The model assumes that the noise on each measurement is independent of the noise on the other measurements.

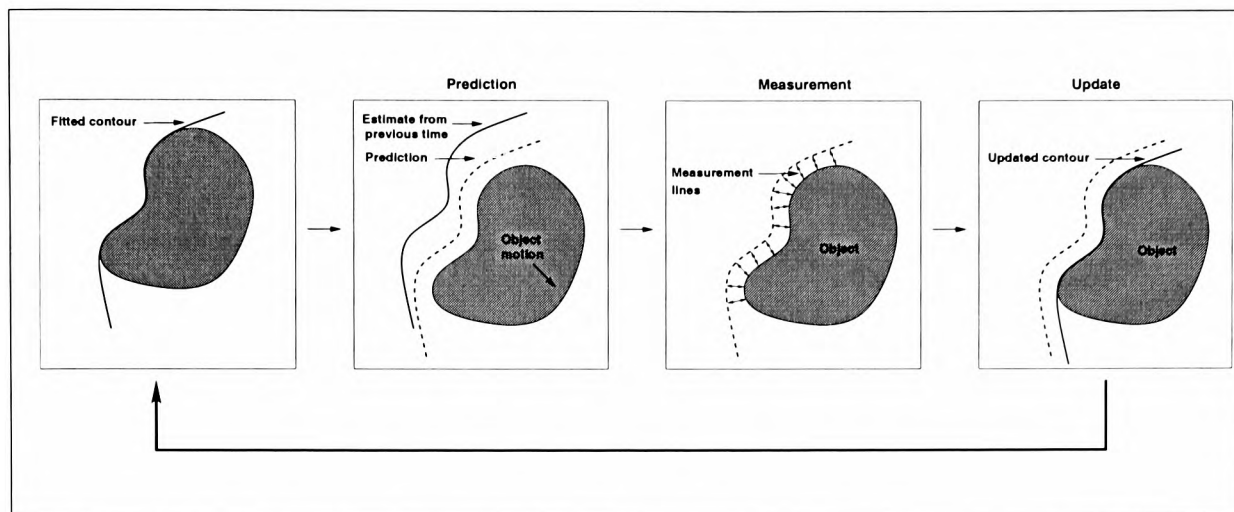


Figure 4.4: **Dynamical tracking.** The central idea behind the tracker is *prediction* and *measurement*. At any given frame, we have an estimate of the position of the contour. A prediction as to its position on the next frame is made. This prediction is then updated based on feature measurements taken from the image. A contour is fitted to these measurements, and becomes the new estimate of the position of the contour. This procedure is repeatedly applied to every frame of the sequence.

This approach to tracking differs from Cootes *et al.* (1994). In their approach they re-initialised the template to the mean position between successive frames, thereby making no use of the coherent motion of the object. Moreover, because the contour is then that much further from the object boundary, any boundary searches employed can cause the contour to be caught in local minima. Moreover, previous approaches to learning shape variability have only used static statistical models such as Grenander *et al.* (1991) and Cootes *et al.* (1993). In contrast, the learning here is dynamic, modelling temporal sequences.

We shall now turn our attention to the dynamical model used in this thesis.

Sections 4.5.1 and 4.6 contain a fair amount of mathematics. The uninterested reader can understand the basis for the tracker from Figures 4.4 , 4.5 and 4.6 and re-join the discussion in Section 4.7.



### 4.5.1 Dynamical Model

For ease of notation, we shall define a state vector,  $\mathcal{X}_t$ , as,

$$\mathcal{X}_t = \begin{pmatrix} \mathbf{X}_{t-1} \\ \mathbf{X}_t \end{pmatrix}. \quad (4.20)$$

$\mathbf{X}_t$  is the model space vector at time  $t$  defined in Equation 4.9 and considered in more detail below.

As we stated, the prediction is based on a second order model — i.e the state at time  $t$  is based on the state at time  $t-1$  and  $t-2$ . We use a second order model because it can model the object's acceleration, and the model is widely used in control theory (Astrom & Wittenmark, 1984). First order models are only able to model objects that have constant average velocity. If the object's velocity changed over time, tracking would inevitably be lost — the model would be out of its domain. There is another important anatomical reason, though. The left ventricle expands and contracts in an oscillatory manner every heart cycle. Blake & Isard (1998, pages 197 & 297) comment that first order processes are unable to model oscillatory signals adequately.

We shall only consider discrete-time models in this thesis — the image is only available at discrete times. The object's motion is modelled using a discrete-time model, in a stochastic dynamical model. This has the form,

$$\mathcal{X}_t - \bar{\mathcal{X}} = \mathbf{A}(\mathcal{X}_{t-1} - \bar{\mathcal{X}}) + \mathbf{B}\mathbf{w}_t, \quad (4.21)$$

where,  $\bar{\mathcal{X}}$  is the system mean,  $\mathbf{A}$  a parameter matrix that controls the deterministic part of the motion and  $\mathbf{B}$  is a matrix giving the stochastic part. The stochastic component of the model is a noise source  $\mathbf{w}_t$ , which has dimension  $N_X \times 1$ . We will assume that the random variables,  $\mathbf{w}_t$ , have a Gaussian distribution, with zero mean and unit variance, i.e.  $\mathbf{w}_t \sim N(\mathbf{0}, \mathbf{I})$ . It is conventional, although, not necessary to treat the  $\mathbf{w}_t$  as independent and identically distributed (i.i.d).

Since we have assumed that  $\mathcal{X}_t = (\mathbf{X}_t, \mathbf{X}_{t-1})^\top$ , then  $\bar{\mathcal{X}} = (\bar{\mathbf{X}}, \bar{\mathbf{X}})^\top$ . Without loss of generality, we can express the matrices  $\mathbf{A}$  and  $\mathbf{B}$  as,

$$\mathbf{A} = \begin{pmatrix} \mathbf{0} & \mathbf{I} \\ \mathbf{A}_2 & \mathbf{A}_1 \end{pmatrix} \quad \text{and} \quad \mathbf{B} = \begin{pmatrix} \mathbf{B}_0 \\ \mathbf{0} \end{pmatrix}.$$

The dynamical model then takes the form,

$$\mathbf{X}_t - \bar{\mathbf{X}} = \mathbf{A}_1(\mathbf{X}_{t-1} - \bar{\mathbf{X}}) + \mathbf{A}_2(\mathbf{X}_{t-2} - \bar{\mathbf{X}}) + \mathbf{B}_0\mathbf{w}_t. \quad (4.22)$$

Here,  $\mathbf{X}$  is the shape-space state vector of dimension  $N_X \times 1$ . The mean configuration,  $\bar{\mathbf{X}}$ , is taken as a constant in the model. The coefficient matrices  $\mathbf{A}_1$  and  $\mathbf{A}_2$  each have dimension  $N_X \times N_X$ . They represent the deterministic component of the model (Pentland & Horowitz, 1991).

The matrix  $\mathbf{B}_0$ , has dimension  $N_X \times N_X$ , and transforms these random variables into the states of the system. It is the noise that drives the system prediction. In fact,  $\mathbf{B}_0$  is not completely observable, only the covariance  $\mathbf{C} = \mathbf{B}_0\mathbf{B}_0^\top$  can be computed; this is because the probability distribution  $\mathbf{w}_t$  being a vector of i.i.d. normal variables is invariant to orthogonal transformations

and so  $B_0$  is under-determined. Using the covariance, we compute  $B_0$  as,

$$B_0 = \sqrt{C}, \quad (4.23)$$

where,  $\sqrt{\cdot}$  is the square root of a positive definite square matrix (Barnett, 1990).

The reader should note that our model is second order in relation to the shape-space vectors  $X$ , Equation 4.22, but first order in terms of the state vector,  $\mathcal{X}$ , Equation 4.21.

The deterministic part of Equation 4.22 is used as the prediction phase of the tracking algorithm. We need only now to specify the  $A_1$ ,  $A_2$  and  $B_0$ . These matrices characterise the temporal motion of the object.

It is important, in any system, to know the parameters of the process being modelled. In many applications, the dynamics are based on the underlying physics of the process. However, the heart moves in a non-rigid and complex manner. The dynamics can be specified in two ways.

#### Default Dynamics

In the absence of any information, we can choose plausible default dynamics that try to characterise intuitively the expected object behaviour. For example, our default dynamics may be a constant-velocity model or, more likely, a damped harmonic system with time-constants appropriate to the object's motion. Of course we shall be interested in default dynamics that provide oscillatory motion. If we are unable to specify approximate correct default dynamics, this is fine, but then we will be reliant on having strong image features to guide our tracking.

#### Learnt Dynamics

The alternative, and more robust approach, is to *learn* these matrices from a training set. By learning the matrices for the object's dynamics we allow the tracker to be *tuned* to the expected motion of the object. Of course, this is the 'chicken and egg' problem. We want to learn the object's motion from a tracked training set, but it is first necessarily to be able to track the object! One solution is, for each frame of a sequence, to manually delineate the boundary around of the object. This is very laborious, and impractical for long training sets. Fortunately, we can use an appropriate set of default dynamics to, as best as possible, track the object. The resulting contours, which best represent the object's position, then form the basis for our training set.

We now consider using a statistical framework to estimate the dynamics from this training set.

### 4.5.2 Learning the dynamical model

Suppose that we have tracked an object's motion for  $M$  time frames. We will assume that we obtain one contour at each time-frame. The resulting  $M$  contours,  $\mathbf{Q}_1, \mathbf{Q}_2, \dots, \mathbf{Q}_M$ , represent the object position at time-frames  $1, 2, \dots, M$ . Alternatively, in shape-space, we have  $\mathbf{X}_1, \mathbf{X}_2, \dots, \mathbf{X}_M$ . In mathematical terms, the problem now is to estimate the matrix parameters  $\mathbf{A}_1, \mathbf{A}_2, \bar{\mathbf{X}}$  and  $\mathbf{B}_0$  which model the motion observed in the training set of shapes,  $\mathbf{X}_1, \mathbf{X}_2, \dots, \mathbf{X}_M$ . We shall use maximum likelihood estimation (MLE). The assumption implicit in the estimation procedure described below, is that the shape-space vectors in the training set are "noise free" — they are the exact observations of the physical process.

Suppose that the random variables  $\mathbf{X}_1, \mathbf{X}_2, \dots, \mathbf{X}_M$  have a joint density function

$f(\mathbf{X}_1, \mathbf{X}_2, \dots, \mathbf{X}_M | \mathbf{A}_1, \mathbf{A}_2, \mathbf{B}_0)$ . Let  $\mathbf{B}_0^{-1} = \mathbf{V}$ . The probability density function takes the form,

$$f(\mathbf{X} | \mathbf{A}_1, \mathbf{A}_2, \mathbf{B}_0) = |2\pi|^{-\frac{1}{2}} |\mathbf{V}|^{\frac{1}{2}} \exp \left\{ -\frac{1}{2} (\mathbf{X}_t - \mathbf{A}_1 \mathbf{X}_{t-1} - \mathbf{A}_2 \mathbf{X}_{t-2})^\top \mathbf{V} (\mathbf{X}_{t+2} - \mathbf{A}_1 \mathbf{X}_{t-1} - \mathbf{A}_2 \mathbf{X}_{t-2}) \right\}. \quad (4.24)$$

The likelihood of  $\mathbf{A}_1, \mathbf{A}_2, \mathbf{B}_0$  as a function of  $\mathbf{X}_1, \mathbf{X}_2, \dots, \mathbf{X}_M$  is then,

$$L(\mathbf{X} | \mathbf{A}_1, \mathbf{A}_2, \mathbf{B}) = |2\pi|^{-\frac{M-2}{2}} |\mathbf{V}|^{\frac{M-2}{2}} \exp \left\{ -\frac{1}{2} \sum_{t=2}^M (\mathbf{X}_t - \mathbf{A}_1 \mathbf{X}_{t-1} - \mathbf{A}_2 \mathbf{X}_{t-2})^\top \mathbf{V} (\mathbf{X}_t - \mathbf{A}_1 \mathbf{X}_{t-1} - \mathbf{A}_2 \mathbf{X}_{t-2}) \right\}. \quad (4.25)$$

Maximising the likelihood with respect to  $\mathbf{A}_0, \mathbf{A}_1, \mathbf{B}$  leads to the following parameter estimators — see Appendix E.2.

$$\begin{aligned} \hat{\mathbf{A}}_1 &= (\mathbf{S}_{20} - \mathbf{S}_{21} \mathbf{S}_{11}^{-1} \mathbf{S}_{10}) (\mathbf{S}_{00} - \mathbf{S}_{01} \mathbf{S}_{11}^{-1} \mathbf{S}_{10})^{-1} \\ \hat{\mathbf{A}}_2 &= (\mathbf{S}_{20} \mathbf{S}_{01}^{-1} \mathbf{S}_{10} - \mathbf{S}_{20}) (\mathbf{S}_{10} - \mathbf{S}_{11} \mathbf{S}_{01}^{-1} \mathbf{S}_{10})^{-1} \\ \mathbf{V}^{-1} &= \frac{1}{M-2} \sum_{t=2}^M (\mathbf{X}_t - \mathbf{A}_1 \mathbf{X}_{t-1} - \mathbf{A}_2 \mathbf{X}_{t-2}) (\mathbf{X}_t - \mathbf{A}_1 \mathbf{X}_{t-1} - \mathbf{A}_2 \mathbf{X}_{t-2})^\top \end{aligned} \quad (4.26)$$

and,

$$\mathbf{S}_{ij} = \sum_{t=0}^{M-2} \mathbf{X}_{t+i} \mathbf{X}_{t+j}^\top \quad i, j = 0, 1, 2, \quad \mathbf{S}_{ij}^\top = \mathbf{S}_{ji}.$$

We then use the estimates of the parameter in our dynamical model, given in Equation 4.22.

The full training algorithm is illustrated in Figure 4.5. The trained tracker is now capable of *predicting* the position of the object of any frame, based on the two previous frames. Of course there will be inaccuracies between the predicted object position and actual (or observed) object

position. This discrepancy is resolved by *updating* the contour position based on *measurements* taken from the image.

We now turn to this second stage of the two-phase tracking algorithm — the measurement model. This is the third key ingredient of our basis for tracking, given in Section 4.2. In Chapter 5, and once again in Chapter 7, we shall see that it is important to use a specifically designed image feature detector as the basis for the measurement model.




Figure 4.5: **Training algorithm.** In the absence of any information, we chose a plausible default set of default dynamics for the tracker. Using the shape-space vectors from the sequence, we obtain an estimate of the parameters in the dynamical model, Equation 4.22, using Equation 4.26. The prediction part of the tracker now behaves according to the *learnt* dynamics — a trained tracker.

## 4.6 The measurement model

As we have seen, at each time step we predict the position of the contour on the next frame. This position is *updated* using “desirable” features taken from the image. These measurements are performed by casting rays along normals,  $\hat{\mathbf{n}}(s, t)$  to the predicted curve position, and measuring the positions,  $\nu(s, t)$ , of the object’s boundary along the ray — see Figure 4.6(a). These measurements are made relative to the predicted curve position, so that, the “innovation”  $\nu(s, t)$  is given by,

$$\nu(s, t) = [\mathbf{r}(s, t) - \bar{\mathbf{r}}(s, t)] \cdot \hat{\mathbf{n}}(s, t) + m(s, t), \quad (4.27)$$

where,  $m(s, t)$  is a spatial measurement noise (scalar) variable, assumed Gaussian with zero mean and variance. The variance is taken as constant both spatially and temporally. We further assume that the noise is independent of both the predicted and actual contour positions. Although they are not independent, we shall accept this assumption because they are constrained by the normal. We can think of the innovation as the difference between the actual measurements and the measurements that would be predicted at each normal based on the mean shape.

The innovations are defined only along normals to the curve as the tangential motion is unobservable locally — the well known aperture problem (Horn, 1986). The spatial measurements covariance,  $R(s)$ , depends on things like, the electrical noise involved in image formation, spatial transducer noise and the detection of erroneous features obtained via the feature detection




Figure 4.6: **Measurement process.** (a) Normals are cast from the predicted contour position (—). The length of the normal is based on the validation gate,  $\rho(s, t)$ . Measurements are taken along this normal to obtain the position of the object's boundary ( $x$ ). The contour is then updated using a Kalman Filter. (b) Measurements are taken at sampled positions of the left ventricular boundary contour. Image features are identified through image processing operations on the search line data. The innovations — the measurements relative to the predicted curve — are incorporated into the Kalman Filter.

process. However, in practice, it is often possible to approximate the measurement densities as Gaussian. Although, ultrasound images are characterised by speckle noise, in practice, this approximation works well.

Isard & Blake (1996) following work by Grenander *et al.* (1991) propose an alternative modelling scheme for non-Gaussian densities based on stochastic sampling.

The contour is then updated based on these measurements. To do this we need to relate the innovation,  $\nu(s, t)$  to the configuration vector,  $\hat{\mathbf{Q}}$ , or state vector,  $\hat{\mathbf{X}}$ . This is done using an observation matrix,  $\mathbf{H}$ , to be defined such that (see also Equation 4.20),

$$\nu(s, t) = \mathbf{H}(s, t)(\tilde{\mathcal{X}} - \hat{\mathcal{X}}) + v(s, t), \quad (4.28)$$

where,  $\tilde{\mathcal{X}}$  is the prediction of  $\mathcal{X}$ . From Equation 4.27 and using Equation 4.4,

$$\begin{aligned} [\mathbf{r}(s, t) - \tilde{\mathbf{r}}(s, t)] \cdot \hat{\mathbf{n}}(s, t) &= (\mathbf{U}(s)\tilde{\mathcal{X}} - \mathbf{U}(s)\hat{\mathcal{X}}) \cdot \hat{\mathbf{n}}(s, t) \\ &= (\mathbf{U}(s)[\mathbf{I} \ \mathbf{0}]\tilde{\mathcal{X}} - \mathbf{U}(s)[\mathbf{I} \ \mathbf{0}]\hat{\mathcal{X}}) \cdot \hat{\mathbf{n}}(s, t) \\ &= \hat{\mathbf{n}}(s, t)^\top \mathbf{U}(s)[\mathbf{I} \ \mathbf{0}](\tilde{\mathcal{X}} - \hat{\mathcal{X}}). \end{aligned}$$

Defining,

$$\mathbf{H}(s, t) = \hat{\mathbf{n}}(s, t)^\top \mathbf{U}(s)[\mathbf{I} \ \mathbf{0}] \quad (4.29)$$

leads to Equation 4.28.

In theory,  $\mathbf{H}(s, t)$  is a continuous function of  $s$ , although in practice the curve is not observed in its entirety but rather at sampled points,  $s_i$ , along the contour. The normals to the contour are referred to as *search lines* and are shown in Figure 4.6(b).

To estimate the final contour position, using a Kalman Filter (Gelb, 1974), the predicted contour is fused together with the measurements taken from the image. The distribution of the state estimates is Gaussian (this is because from Equation 4.21  $\mathbf{w}_t$  is assumed to be Gaussian). If we assume a measurement history  $\mathbf{Z}_1, \mathbf{Z}_2, \dots, \mathbf{Z}_t$ , then,

$$\hat{\mathcal{X}}(t_k) = E(\mathcal{X}(t_k) | \mathbf{Z}_1, \mathbf{Z}_2, \dots, \mathbf{Z}_t) \quad \text{and} \quad \mathcal{P}(t_k) = \text{Var}(\mathcal{X}(t_k) | \mathbf{Z}_1, \mathbf{Z}_2, \dots, \mathbf{Z}_t), \quad (4.30)$$

and so,

$$\mathcal{X}(t_k) \sim N(\hat{\mathcal{X}}(t_k), \mathcal{P}(t_k)).$$

This means that all the information about the left ventricular boundary position is carried by the 1st order (mean) and 2nd order (covariance) moments.

The prediction phase is based on the dynamical model described in Section 4.5.1,

$$\tilde{\mathcal{X}}(t_k) - \bar{\mathcal{X}} = \mathbf{A}(\mathcal{X}(t_{k-1}) - \bar{\mathcal{X}}), \quad (4.31)$$

$$\tilde{\mathcal{P}}(t_k) = \mathbf{A}\mathcal{P}(t_{k-1})\mathbf{A}^\top + \mathbf{B}\mathbf{B}^\top. \quad (4.32)$$

Measurements are then taken as described in Section 4.6 at a number of points along the spline perimeter. These measurements are then used to (possibly) influence the predicted contour position and variance. This can be done in two entirely equivalent ways; using an information weighting or a Kalman gain. The Kalman gain, a vector,  $\mathbf{K}(t_k)$ , is defined as

$$\mathbf{K}(t_k) = \tilde{\mathcal{P}}(t_k)\mathbf{H}^\top (\mathbf{H}\tilde{\mathcal{P}}(t_k)\mathbf{H}^\top + \sigma^2\mathbf{I})^{-1} \quad (4.33)$$

This is then used to produce estimates based on the new measurements:

$$\tilde{\mathcal{X}}(t_k) = \tilde{\mathcal{X}}(t_{k-1}) + \mathbf{K}(t_k)\mathbf{Z}(t_k), \quad (4.34)$$

$$\mathcal{P}(t_k) = (\mathbf{I} - \mathbf{K}(t_k)\mathbf{H})\tilde{\mathcal{P}}(t_k). \quad (4.35)$$

These equations are applied at the measurement points along the spline.

One advantage of using the Kalman filter is that each state estimate is accompanied by an estimate of its covariance. This enables a validation gate,  $\rho(s, t)$  to be defined which only allows



measurements to be included if they lie within a certain distance of the predicted position. It also allows for an adaptive search region, so that if image measurements fail along the contour, the search region across the spans will expand until observations are found again.

## 4.7 Conclusions

This chapter has provided an overview of a framework for a dynamical 2-dimensional tracker. This framework will serve as the basis for many of the ideas presented and developed further in this thesis. The key, as we shall see, to successful tracking of the endocardial boundary is based on three powerful ingredients:

- A prior (tuned) shape model
- Learning the object's motion
- Specific image feature detectors

The endocardial boundary is represented by a parametric B-spline with 24 control points (Figure 4.1), and measurements are taken at various points along the contour (Figure 4.6). These measurements together with the prior shape model and learnt dynamics allow efficient robust tracking of the endocardial boundary.

In the next chapter we shall address these 3 key ingredients. In particular, we shall see that an affine shape model is not capable of modelling the complex deformations of the heart wall. We shall also see that due to noise and clutter in the ventricular cavity and image it is difficult to reliably obtain image feature measurements using a standard feature detector.

# CHAPTER 5

---

## Tracking Endocardial Wall Excursion

---

In Chapter 4 we laid the foundations for successful model based dynamical tracking. In this chapter we apply the theory to the problem of tracking the left ventricular boundary of the heart, and in particular the blood-tissue boundary otherwise known as the *endocardium*. For the endocardial boundary, our aim will be to:

- Build an appropriate shape model,
- Learn its dynamics, and,
- Address the problem of what feature detector to use.

All the ideas described in this chapter are original contributions of the thesis. Earlier versions of this work are published in (Jacob *et al.*, 1998, 1999).

Discontinuities in the brightness between object contours and their surroundings often result in identifiable boundaries or edges. In the context of traditional computer vision, this has led to the development of edge detection algorithms (Canny, 1986; Deriche, 1987; Perona & Malik, 1990). Consequently, because the boundary features are step discontinuities, edge detection schemes have been widely used in tracking applications (Kass *et al.*, 1987; Herlin & Giraudon, 1993; Tained *et al.*, 1994).

Ultrasound images are formed by sending out high frequency sound into the body. As the beam hits the blood-tissue boundary, there is a discontinuity in the radio-frequency signal returned and hence the resulting intensity in the image. This discontinuity has been treated by many

researchers as an edge (Ayache *et al.*, 1992; Chalana *et al.*, 1996). However, when tracking the complex deformations of a moving heart, in an image which is characterised by speckle noise and clutter, edge detection schemes have proved to be inadequate (Friedland & Adam, 1989; Dias & Leitão, 1996).

The main reason for this is that speckle noise (Burckardt, 1978; Wells & Halliwell, 1981; Wagner *et al.*, 1983) and the spurious clutter inside the ventricle (Figure 5.1), result in spurious edges that are not due to the blood-tissue boundary of the endocardial wall. In addition, image intensities differ because the angle of incidence of the ultrasound beam and wall, differs along the wall (Figure 5.2). Stronger reflections result from those parts of the wall that are perpendicular to the beam. Whereas, those parts of the wall which are nearly parallel to the beam, such as the inferior wall and, in particular, the anterior wall (refer to Figure 6.1, page 80), are characterised by a lower contrast in intensity.

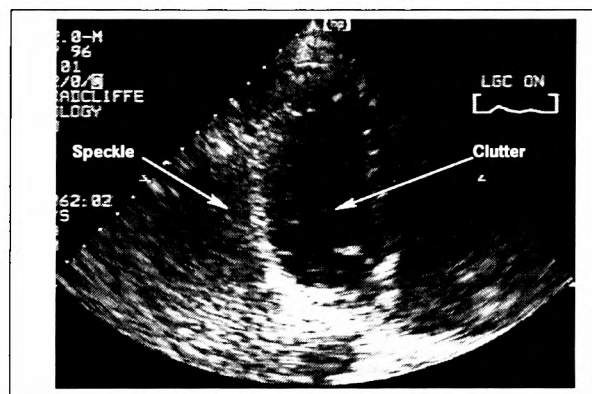


Figure 5.1: **Noise in echocardiographic images.** The noise in an echocardiographic image arises from two sources: spurious clutter inside the ventricle and speckle noise. This data was recorded on VHS video and then digitised.




Figure 5.2: **Intensity profiles.** (a) Tracker interface with six profiles normals. (b) Associated grey-level intensity profiles plotted against position along the profile. The endocardial boundary occurs at  $r = 0$ , while  $r < 0$  represents the intensity distribution of left ventricular chamber, and  $r > 0$  corresponds to the myocardium. Stronger reflections result from those parts of the wall that are perpendicular to the beam, such as profile 6. Reflections in the near field zone (profiles 3 and 4) are characterised by a lower contrast in intensity. Reflections from parts of the wall which are nearly parallel to the beam, such as profile 5 also show a reduced intensity. The clinician is able to adjust the gains to partially compensate for this.

As a consequence there is a lack of 'good' features. This problem is particular troublesome for the contour based tracker, where, in our approach, only one curve measurement can be integrated into the Kalman filter framework.

Methods exist for combining multiple observed features using a probabilistic data association filter (PDAF) Bar-Shalom & Fortmann (1988, page 163) by weighting each observation with the probability that it originated from the target. The problem is that the PDAF assumes that the target probability distributions and the probabilities of obtaining false measurements are known, which is not often the case. More fundamentally, the PDAF does not solve the problem that the endocardial border is not always visible. As discussed in Section 4.6, accurate and reliable feature measurements are a necessary ingredient in the dynamic contour tracker framework. Thus in addition to developing precise shape and motion models, there remains the difficult problem of how best to extract image features which can accurately delineate the endocardial boundary.

In this chapter, we explore a number of solutions for tracking in ultrasound. Section 5.1 deals with obtaining a shape-space via PCA. The time-varying shape-vector components of the tracking potentially offer clinical insight. This is illustrated in Section 5.2, although we go on to show in,

Chapter 6, that there are better ways of decomposing tracking into regionalised components.

To localise features along the endocardial boundary, we introduce intensity invariant feature detection in Section 5.3 and compare its performance against traditional intensity-based boundary.

Section 5.4 details the Chapter conclusions.

## 5.1 Prior Shape Model

In Sections 4.3.1 and 4.4 we saw that, in the general case, there are a number of ways to build an appropriate *prior* model of shape. In this section, using PCA (Section 4.4), we obtain a shape-space for a 2-dimensional scan of the left ventricle, and then go on to compare tracking performance using this shape-space with that of an affine shape-space Section 4.3.1.

We saw in Section 4.4 that we can construct a PCA shape-space using the control points from a number of manually segmented image frames. In the first part of this chapter, we will be mainly using a single echocardiographic sequence to illustrate the tracking approach and performance.

To obtain a shape matrix,  $W$ , for an echocardiographic data set, a B-spline with 14 control points was used to delineate the boundary of the left ventricle. The peak of the ECG R-wave was chosen as the starting point for the cardiac cycle. Note that the training sets were not normalised to take out global translation. This is beneficial for tracking purposes but makes the interpretation of the PCA in terms of physical motion more complicated. Table 5.1 summarises the results of PCA analysis. For this data set, four modes explained 95% of the variation.

The PCA analysis is intended to capture the main variability inherent in the data. As such, we would expect that each mode could represent part of the contraction/expansion stage of the heart cycle — the physical process. We can visually check to see how intuitive each PCA mode is by plotting the left ventricular template under the action of each deformation mode in turn. This is discussed in the next section.

Mode	Eigenvalue	Variance proportion	Cumulative proportion
1	4188.34	0.712	0.712
2	702.11	0.119	0.832
3	480.51	0.0817	0.913
4	255.65	0.0435	0.957
5	80.88	0.0138	0.971
6	55.77	0.00948	0.980

Table 5.1: **PCA analysis of shape.** The result of applying a Principal Component Analysis (PCA) to 4 manually segmented cardiac cycles. 4 modes of variation explain over 95% of the variability.



### 5.1.1 Deformation modes

Our goal is to determine a smaller shape-space,  $S' = \mathcal{L}(\mathbf{W}', \mathbf{Q}'_0) \subset S$ , that is a subspace of  $S$  that spans, at least approximately, all possible shapes in the training set. Mathematically, any shape,  $\mathbf{Q}_i$ , in the training set  $\mathbf{Q}_1, \mathbf{Q}_2, \dots, \mathbf{Q}_M$ , can be estimated as,

$$\hat{\mathbf{Q}}_i = \bar{\mathbf{Q}} + \sum_{i=1}^r \lambda_i \mathbf{w}_i, \quad (5.1)$$

where,  $r$  is the number of basis vectors retained,  $\lambda_i = \langle \mathbf{Q}, \mathbf{w}_i \rangle$ , is the  $i$ 'th eigenvalue, and  $\mathbf{w}_i$  is the  $i$ 'th associated eigenvector, or mode. The mean,  $\bar{\mathbf{Q}}$ , is added to the expansion because the data is mean corrected. In fact,  $\lambda_i$  is the sample variance of  $\mathbf{Q}_i$ .

To gain some insight into what each mode represents, we can take the mean shape,  $\bar{\mathbf{Q}}$ , or the initial template,  $\mathbf{Q}_0$ , and add/subtract a multiple of each mode, plotting the result. In fact, using, Equation 5.1, we can plot  $\bar{\mathbf{Q}} \pm m \mathbf{w}_i \sqrt{\lambda_i}$ , where,  $m$  is the number of standard deviations from the mean (we choose  $m = 3$ ). Figure 5.3 shows plots of the first 6 modes for this example. The thicker contour is the mean shape curve. The two thinner curves represent the mean shape  $\pm 3$  standard deviations. We can tentatively suggest the following. The first mode appears to be a translation mode. The second mode appears to be a scaling mode where the scaling applies to the bottom of the left ventricle next to the mitral valve. The third and fourth modes both appear to represent a combination of scaling and translation.

Attempting to relate both the number, and visually the nature of the modes of variation to any underlying physical process, is fraught with difficulty. PCA does not take account of this. It is merely a mathematical tool which partitions the variation in a data set, based on maximising the variance of a projection of the points onto a subspace of the full space. In this connection, it is worth quoting the following from Marriott (1974).

*It must be emphasised that no mathematical method is, or could be, designed to give physically meaningful results. If a mathematical expression of this sort has an obvious physical meaning it must be attributed to a lucky chance, or to the fact that the data have a strongly marked structure that shows up in the analysis. Even in the latter case, quite small sampling fluctuations can upset the interpretation; for example the first two principal components [in our case modes] may appear in reverse order, or may become confused altogether. Reification, then, requires, considerable skill and experience if it is to give a true picture of the physical meaning of the data.*

The implication from this is that when we try to relate our tracking results to heart function, just

looking at the PCA modes is not sufficient.

It is very common, both in medical and non-medical applications, to visualize the modes, as shown in Figure 5.3 — each mode in turn is added to and subtracted from a spline (2D contour) template (Cootes *et al.*, 1995; Ruff *et al.*, 1997; Heap & Hogg, 1998). In their face tracking paper, Edwards & Taylor (1998) build a point distribution model of a face. To visually show the effect of each mode, they independently add and subtract the first 4 modes to a mean face. Because, changes in facial expression are immediately recognisable, this gives a more intuitive feeling for what each mode represents. However, only their first mode was readily interpretable, being the dominant motion in the training set. Second and higher modes were mixtures of motions.

The problem with Figure 5.3 is that since we only view the 3 contours, it is not directly clear how the contours correspond to each other. For example, the dominant mode in Figure 5.3 appears to be a translational mode — this is actually not the case as Figure 5.4 shows. In this figure, flow vectors have been used to indicate the deformation for selected points (the start of each span) along the contour. This representation gives a better feel for the direction of each mode. It also reveals that although the most dominant mode is a small translation, there is a high degree of re-parameterisation (see Figure 5.5), which was unclear from Figure 5.3. Re-parameterisation contributes to measured variance, but does not represent a real shape-change, as shown in Figure 5.5.

Initially, and for the example discussed above, a B-spline using 14 control points was used. In the next, and later examples and analyses presented in this thesis, 24 control points were used. This is because it was not immediately obvious how many control points to use, and where they should be placed. It turns out that, although tracking is successful when using 'few' control points, a slight reparameterisation of the curve can occur — the position of the control points relative to their initial position can change, although the curve remains visually unchanged. As an example, suppose that we draw a circle using a B-spline which has 4 control points placed at equidistant angles. We could simply shift the control points along by 1 and visually obtain the same curve, but the curve has been re-parameterised, as shown in Figure 5.5.




Figure 5.3: **PCA deformation — example.** A principal component analysis was performed on 4 cardiac cycles of a ultrasonic image sequence using a B-spline with 14 control points. The results are shown in Table 5.1. To understand what each mode represents, we have plotted the mean shape (—) under the action of each mode in turn. The two other curves (---) represent the addition of  $\pm 3$  standard deviations to the mean shape. The modes are plotted in decreasing order of eigenvalue — *Top left*: The dominant mode to *Bottom right*: The least dominant. We can tentatively suggest the following. The first mode appears to be a translation mode. The second mode appears to be a scaling mode where the scaling applies to the bottom of the left ventricle next to the mitral valve. The third and fourth modes both appear to represent a combination of scaling, translation and rotation.

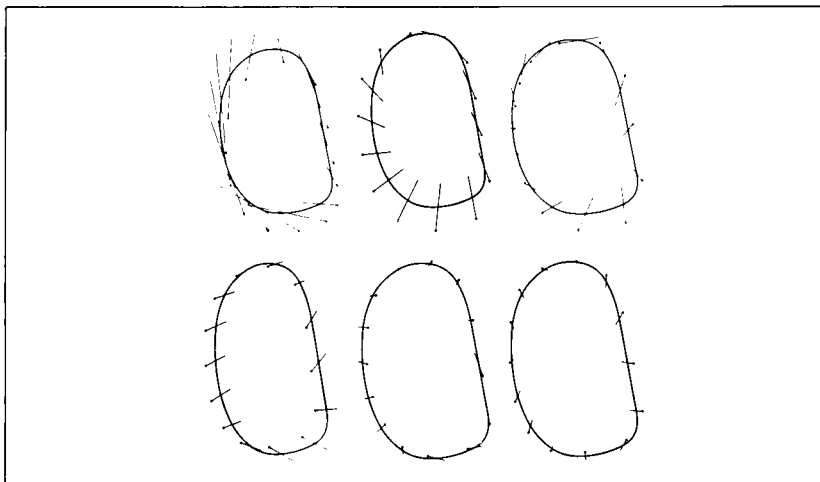


Figure 5.4: **PCA deformation — example.** The problem with Figure 5.3 is that since we only view the 3 contours, it is not directly clear which parts of the contours relate to each other. As an alternative, a flow vector is centred on the start of each span. The end points of each flow vector are located at  $\pm 3$  standard deviations from the mean shape in the direction of that particular PCA mode. To get a feel for the direction of the mode, (·) represents  $-3$  standard deviations. The modes are plotted in decreasing order of eigenvalue and match those in Figure 5.3. We can see that although the most dominant mode is a small translation, there is a high degree of re-parameterisation (see Figure 5.5), which was unclear from Figure 5.3.




Figure 5.5: **B-spline re-parametrisation.** A B-spline with 4 control points at equidistant angles is used to parametrise a circle. Suppose we shift the control points along by 1. We would visually obtain the same curve, but the curve has been re-parameterised.

The disadvantages of this re-parameterisation is two-fold. Firstly, to create a shape-space using PCA, a training set based on a number of manually segmented endocardial boundary outlines is used. If too few control points are used, then the first and dominant mode, is merely a re-parameterisational mode, and does not supply any information about the physical process going on.

What is effectively happening is that, if too few control points are used to manually segment a boundary, the control points then need to be moved into a slightly different anatomical position so that the curve can be fitted. Secondly, analysing the regional endocardial wall excursion over time is of interest here. This is done by considering the positions of the spline control points over time as an indicator of regional excursion. To do this, an *interpretational shape-space* is developed in Chapter 6. This makes the assumption that over time, relative to the initial template, the control points remain in roughly the same position. In the circle example above, the control points would be distorted relative to their initial position.

As a way of comparison, Figure 5.6, shows a second data set in which 24 control points were used. In this case our experience has shown that no re-parameterisation occurs, although there is no formal way to prove it that we are aware of.




Figure 5.6: **PCA deformation — example.** A principal component analysis was performed on an ultrasonic image sequence using a B-spline with 24 control points. In comparison to Figures 5.3 and 5.4, our experience has shown that no re-parameterisation of the contour takes place, although there is no formal way to prove it that we are aware of. (a) The most dominant six modes of variation. (b) Modes represented with flow vectors.

### 5.1.2 How many PCA modes should we use?

So far we have looked at generating a PCA shape basis, but the question arises as to how many modes shall we keep. Too few and we will not accurately capture enough of the shape deformation; too many, and that defeats the objective of reducing the dimension of the shape-space. Various principles have been put forward to introduce some objectivity into the process (Jolliffe, 1986, pages 92–114). Some authors (Cootes *et al.*, 1995; Bascle & Blake, 1998), choose the number of modes, such that the sum of the eigenvalues for those modes are greater than 95%. Picking 95% is rather subjective, and often 80% would suffice.

One approach, favoured by statisticians, is to use a scree diagram (Krzanowski, 1988, page 67)

and (Jolliffe, 1986, page 96), as shown in Figure 5.7. This plots the magnitude of the ordered eigenvalues against the dimension. Under the assumption that those dimensions corresponding to the flat portion of this graph represent undifferentiable ‘noise’ components of the data, then one should logically choose the number of dimensions as the elbow of the initial steep decline — or the elbow of the scree diagram. In the case of Figure 5.7 this would be 5 components.

Having obtained our shape basis we are now in a position to track the left ventricular boundary. However, before doing that, we will first mathematically compare this PCA basis with a standard affine basis.




Figure 5.7: **Scree diagram.** A scree diagram showing a plot of the magnitude of the ordered eigenvalues against dimension for the data in Table 5.1. The number of dimensions is typically chosen as the elbow of the initial steep decline — or the elbow of the scree diagram.

### 5.1.3 Can we assume an affine mode of deformation?

Recall that it is possible to define the  $\mathbf{W}$  matrix with varying degrees of freedom (dimensionality). A low-dimensional space, such as an affine space, is attractive as it is easier to compute and offers an intuitive interpretation (Section 4.3.1). Prior work on tracking hearts in 2-dimensional image sequences has assumed this model. On the other hand a higher dimensional space might be necessary for accurately characterising regional deformation and tracking.

Suppose that we project a PCA space in to an affine space and then back again. The residual,  $r$ , that is left, provides a measure of the degree to which the PCA space can be assumed affine. If the PCA space was affine, this residual would be 0, and if it was not, then it would be closer to 1. Formally, we define the residual,  $r$ , as,

$$r = \frac{\|\mathbf{w}_i - \mathbf{W}_A \mathbf{W}_A^\dagger \mathbf{w}_i\|_2}{\|\mathbf{w}_i\|_2}, \quad (5.2)$$

where,  $\mathbf{w}_i$  is an eigenvector of the PCA  $\mathbf{W}$  matrix,  $\mathbf{W}_A$  is an affine shape matrix,  $\mathbf{W}_A^\dagger$  is its corresponding pseudo-inverse and  $\|\cdot\|$  the  $L_2$  norm.

Table 5.2 summarises the residuals computed for the first 4 modes of Table 5.1. We can see that although modes 1, 2 and 4 are fairly close to affine components, only 12% of mode 3 can be explained by an affine deformation.

This is an important and interesting result because it tells us that the dynamics of the left ventricular boundary cannot be modelled well by an affine deformation. This result will be confirmed visually later in Section 5.1.4 when we show results from tracking the left ventricular boundary with an affine shape-space.

The residuals for the image sequence shown in Figure 5.6 are shown in Table 5.3. Again modes 3-6 are not affine.

Let us now turn to comparing the tracking results from an affine and PCA shape basis.

Eigenvector	$\frac{\ \mathbf{w}_i - \mathbf{W}_A \mathbf{W}_A^\dagger \mathbf{w}_i\ _2}{\ \mathbf{w}_i\ _2}$	$\left(\frac{\ \mathbf{w}_i - \mathbf{W}_A \mathbf{W}_A^\dagger \mathbf{w}_i\ _2}{\ \mathbf{w}_i\ _2}\right)^2$
$\mathbf{w}_1$	0.3411	0.1163
$\mathbf{w}_2$	0.2931	0.0859
$\mathbf{w}_3$	0.9392	0.8821
$\mathbf{w}_4$	0.4811	0.2315

Table 5.2: **PCA Residuals.** Projecting a PCA shape-space,  $\mathbf{W}$ , into and out of an affine space. The residual, as defined in Equation 5.2, as  $r = \frac{\|\mathbf{w}_i - \mathbf{W}_A \mathbf{W}_A^\dagger \mathbf{w}_i\|_2}{\|\mathbf{w}_i\|_2}$ , is shown. Only 11% of mode 1 and 8% of mode 2 can be characterised as not being affine components. Whereas, mode 3 is not affine. The data in 10 second image sequence in the top example, was recorded on VHS video and then digitised. There are approximately 22 frames per cardiac cycle. 14 control points are used.

Eigenvector	$\frac{\ \mathbf{w}_i - \mathbf{W}_A \mathbf{W}_A^\dagger \mathbf{w}_i\ _2}{\ \mathbf{w}_i\ _2}$	$\left(\frac{\ \mathbf{w}_i - \mathbf{W}_A \mathbf{W}_A^\dagger \mathbf{w}_i\ _2}{\ \mathbf{w}_i\ _2}\right)^2$
$\mathbf{w}_1$	0.4864	0.2366
$\mathbf{w}_2$	0.2911	0.08473
$\mathbf{w}_3$	0.8973	0.8051
$\mathbf{w}_4$	0.7838	0.61437
$\mathbf{w}_5$	0.8625	0.7439
$\mathbf{w}_6$	0.8114	0.6584

Table 5.3: **PCA Residuals.** Projecting a PCA shape-space,  $\mathbf{W}$ , into and out of an affine space. Equation 5.2, as  $r = \frac{\|\mathbf{w}_i - \mathbf{W}_A \mathbf{W}_A^\dagger \mathbf{w}_i\|_2}{\|\mathbf{w}_i\|_2}$ , is shown. Modes 3-6 are not affine. The data was digitally captured and stored on optical disk. 24 control points are used.

#### 5.1.4 Comparing shape models by visual inspection

An alternative way to compare how well different shape-models capture heart dynamics is to perform a visual inspection of tracking performance. Figure 5.8 shows typical frames from two different ultrasound sequences. The top pair represent data recorded on VHS video and then digitised. The bottom pair are digitally captured and stored on optical disk. Our aim was to compare an affine shape-space (Figure 5.8 (a)) and a PCA shape-space (Figure 5.8 (b)). It was concluded that heart dynamics are not well modelled by a (simple) affine model. Images based on an affine and PCA shape-space for the complete cardiac cycle in the second example are shown in Appendix F.1. This is not surprising. We saw in Table 5.2 in Section 5.1.3, that the modes of deformation (the columns of the  $W$  matrix) of the finely tuned PCA are not affine.

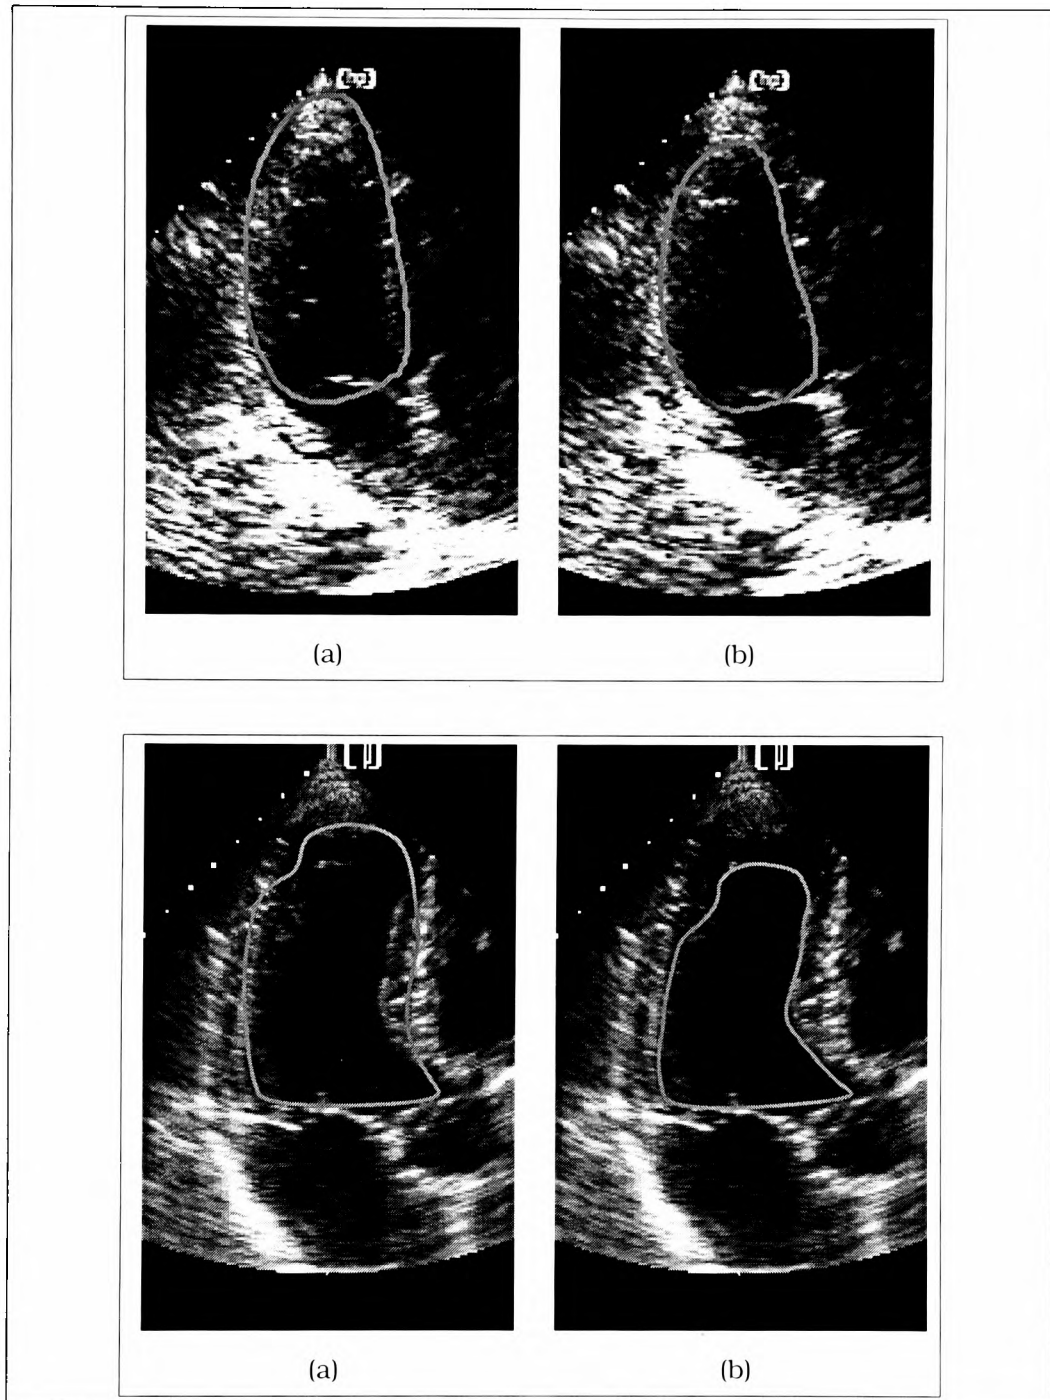


Figure 5.8: **Shape-space tracking.** Echocardiographic tracking using an (a) *affine* and (b) *PCA* shape space on two different hearts (top and bottom). The affine shape-space is unable to deal with the non-rigid deformations of the left ventricular boundary. In contrast, the finely tuned PCA shape-space, which is learnt from a training set of example contours, allows the contour to deform in accordance with heart motion. (Top Pair) The data in 10 second image sequence in the top example, was recorded on VHS video and then digitised. There are approximately 22 frames per cardiac cycle. 14 control points are used with a canny edge detector. The features chosen are the measurements with the strongest response to the filter. (Bottom Pair) The data in 3.5 second image sequence in the top example, was digitally captured and stored on optical disk. There are approximately 81 frames per cardiac cycle. 24 control points are used with a local phase detector (see Section 5.3.1). The features chosen are the measurements with the strongest response to the filter.

## 5.2 Component plots

After tracking, a time sequence of contours,  $\mathbf{Q}_1, \mathbf{Q}_2, \dots, \mathbf{Q}_M$  is obtained, which represent the estimated position of the object's boundary at times,  $1, 2, \dots, M$ . Recall from Equation 4.10 that,

$$\mathbf{Q} = \mathbf{Q}_0 + \mathbf{W}\mathbf{X} \quad \text{and} \quad \mathbf{X} = \mathbf{W}^\dagger(\mathbf{Q} - \mathbf{Q}_0). \quad (5.3)$$

Each component of  $\mathbf{X}$ ,  $X_i$ , represents a vector of time-varying weights which act upon the respective part of the shape-space basis,  $w_i$ . When their product is added to the initial configuration,  $\mathbf{Q}_0$ , a new configuration is generated which represents the object boundary at some (possibly different) point in time. Plotting the elements of  $\mathbf{X}$  against time provides insight into the influence of each shape-space deformation vector,  $w_i$ . For example, suppose that the dominant mode in a PCA basis represents the expansion/contraction of the left ventricle. Correspondingly, suppose that the magnitude of its shape-space vector component, is small relative to 'normal'. We could then tentatively suggest that the heart is not contracting and expanding properly. Figure 5.9 shows plots of the shape-space vector  $\mathbf{X}$  over time for the heart data of Figure 5.4. In particular, note the periodicity of the second mode. It is not surprising that the first mode is not very periodic, since as we saw above, it merely represents a re-parameterisation of the contour, and not any underlying physical deformation.

Figure 5.10 shows the components for the data in Figure 5.6 (where there is no-reparametrisation of the contour). In this case, all components exhibit periodicity.

Certainly, temporal plots of this kind are potentially of great clinical value for quantifying heart periodicity and asynchrony, although this has not been explored in this thesis (see the future work in Chapter 9).

A better way of interpreting the time sequence of contours turns out to be using an *interpretational shape-space* representation, analogous to the PCA shape-space,  $\mathbf{W}$ , used above for tracking, but designed to decompose the contours into a meaningful time varying shape-space vector. This representation is discussed in Chapter 6.




Figure 5.9: **Component plots.** PCA components 1 through 4 of the shape-space vector  $\mathbf{X}$  against time for the 10 second ultrasound sequence shown in Figure 5.4. The components represent a decomposition of the contours,  $\mathbf{Q}_1, \mathbf{Q}_2, \dots, \mathbf{Q}_M$ , into the time-varying weights,  $\mathbf{X}$ , which act upon the respective part of the shape-space,  $\mathbf{W}$ . Note that the first mode is not periodic. This is due to the re-parametrisation of the contour.




Figure 5.10: **Component plots.** PCA components 1 through 4 of the shape-space vector  $\mathbf{X}$  against time for the 3.5 second ultrasound sequence shown in Figure 5.6. The components represent a decomposition of the contours,  $\mathbf{Q}_1, \mathbf{Q}_2, \dots, \mathbf{Q}_M$ , into the time-varying weights,  $\mathbf{X}$ , which act upon the respective part of the shape-space,  $\mathbf{W}$ . Note the periodicity in the first mode compared to Figure 5.9. This is probably due to there being no re-parametrisation of the contour.

### 5.3 Improving feature detection

In this section we turn our attention to improving the feature measurement process. We develop an alternative to gradient-based feature detection using a combination of spatio-temporal noise reduction filtering (Evans & Nixon, 1996) and phase-congruency ideas (Kovesi, 1996). This has been subsequently developed further for use in echocardiographic images by Mulet-Parada &

Noble (1998). We show how this leads to a more stable and robust measurement process in the tracking algorithm. The TDLMS filter (Evans & Nixon, 1996) discussed in this section was implemented by Miguel Mulet-Parada.

Recall that the original *visual* image tracking algorithm uses a 1-dimensional gradient-based operator for detecting contour points. This approach produces many candidate responses on ultrasound images. This is due partly to the low signal-to-noise ratio and also to poor image contrast.

There are two disadvantages with using a gradient-based feature detector such as the Canny edge detector (Canny, 1986). Firstly, how should we choose the feature gradient threshold above which we accept features? Should this be a global or regional threshold? A regional threshold appears more sensible, because the intensities differ along the left ventricular wall, and hence the feature responses will differ. The intensities differ because the angle of incidence of the ultrasound beam on the wall, differs along the boundary (Rose *et al.*, 1995).

Moreover, during an echocardiographic study, the clinician will change the *gains* on the ultrasound machine to view the left ventricular boundary better. So no two patients will be imaged under the same acquisition conditions, and hence image quality differs. Secondly, the image intensity of the *same* boundary point over the cardiac cycle, is not constant. This is an especially challenging problem and depends on the orientation of the tissue with respect to the incident ultrasonic beam and the effect of occlusion artifacts.

For example, Figure 5.11 (a), shows the output of a gradient-based detector on an echocardiographic frame. A large threshold was set to eliminate noise responses on the posterior (left) wall of the cavity. As a result, low intensity features are missed on the anterior (right) wall.

This highlights that the threshold used in an intensity based gradient feature detector, even if regional, will need to be adapted to the stage of the cardiac cycle. In other words, we desire an intensity invariant feature detector. This would overcome any problem with the image intensity varying both by position and time.




Figure 5.11: **Feature detection.** Traditional computer vision edge detection has used intensity based feature detection. This fails in the case of echocardiographic data. Image intensities vary both by the position along the left ventricular wall, and the time in the cardiac cycle, as well as the gains used by the clinicians when recording the images. The result is that it is difficult to choose a global (and regional) intensity based feature gradient threshold. The contour is shown as a red dashed curve. For each normal, features above the noise threshold are shown as yellow dots with the strongest feature shown as a red cross. (a) Intensity based image feature detection with a gradient threshold that is the same for each normal. (b) Intensity-invariant phase based image feature detection with a local phase threshold that is the same for each normal.

We now show how to significantly improve the quality of feature detection by,

1. taking into consideration temporal continuity to reduce noise effects, and,
2. replacing the visual edge feature detector by,
  - (a) applying an acoustic boundary feature enhancement operator, and,
  - (b) modifying the localisation process to make it invariant to intensity.

To achieve (1) we apply the two-dimensional least mean square (TDLMS) filter proposed by Evans & Nixon (1996). Briefly, the TDLMS filter considers a given image in a sequence at time  $t$ . Images

close in time to this image, for example  $t-2, t-1, t+1, t+2$  are combined using a linear weighting scheme. The weights are chosen to minimise the least mean square error between the resulting output image and the corresponding image at time  $t$ .

Next, we enhance the boundaries using integrated backscatter (IBS) boundary enhancement (Vered *et al.*, 1987; Thomas III *et al.*, 1986, 1989; Peréz *et al.*, 1992a). There are various ways of obtaining integrated backscatter images, both in the frequency and time domain Rijsterborgh *et al.* (1993).

One method to calculate integrated backscattered images in the frequency domain is by firstly Fourier transforming the ultrasound signal to obtain a power spectrum. Calibration is accomplished by reference to the signal backscattered from a stainless steel reflector. The corrected power spectrum contains the full spectral information from the tissue. Finally, the corrected spectrum is frequency averaged to obtain the integrated backscatter measure.

IBS images show a marked reduction in both speckle and clutter. Accordingly, discrimination of the endocardial blood interfaces or boundaries is facilitated, allowing detection of the 'desirable image features' used in the tracking of these boundaries (Peréz *et al.*, 1992b). Recall that in visual image analysis, object boundaries are associated with step edges in intensity. Therefore boundaries can be found by localising extrema in the output of a gradient-based operator. This is the feature measurement operator used in the original visual image tracker. However, an ideal acoustic edge is defined as a discontinuity in acoustic impedance (an intensity ridge) or equivalently a discontinuity in acoustic energy or integrated backscatter (Peréz *et al.*, 1992b). The IBS model has been used to measure properties of myocardial tissue since changes in IBS relate to changes in acoustic impedance which in turn relate to changes in relaxation and elasticity (Lange *et al.*, 1995). One of the strengths of integrated backscatter is that it is a simple relative estimate of the energy backscattered from a small volume of tissue (Rose *et al.*, 1995). IBS boundary enhancement is also available in state-of-the-art commercial echocardiographic imaging systems<sup>1</sup> for single frame real-time edge detection (Vered *et al.*, 1987; Peréz *et al.*, 1992b).

Rijsterborgh *et al.* (1993) compared various ways of obtaining integrated backscatter images, both in the frequency and time domain, and concluded that there was no obvious difference between various methods to obtain integrated backscatter images. Subsequently, Moran *et al.* (1994) compared four backscatter parameters calculated in the time domain from the RF data to determine whether cardiac-cyclic variation of these parameters occurred, and if so, the magnitude. The parameters were; (1) the average integrated backscatter; (2) the average of the backscatter power from the log-compressed data; (3) the square of the average of the uncompressed RF amplitude; (4) the square of the average of the RF amplitude from the log-compressed

<sup>1</sup>Including the HP5500 SONOS ultrasound machine in the Medical Vision group and the JR hospital, Oxford.

data. In addition, video images derived from the same RF data were digitised and the magnitude of variation of the grey scale intensity over the region of interest was compared with each of the RF backscatter parameters. Moran showed significant correlation between the four backscatter parameters and integrated backscatter. But there was reduced correlation obtained between these parameters and parameters derived from the video data.

Figure 5.12 (a) shows one frame of an echogram together with (b) a spatio-temporal speckle-reduced and (c) spatio-temporal speckle-reduced/IBS enhanced versions of the same image. From Figure 5.12 (c) it is clear how speckle noise is significantly reduced and endocardial borders enhanced by this two step process.

To localise acoustic boundaries, an intensity invariant localisation process based on phase information is applied. This is discussed in the next section.




Figure 5.12: **Noise reduction.** (a) the original image; (b) the image after the TDLMS filter is applied; (c) TDLMS filter followed by IBS filtering.

### 5.3.1 Localisation based on local phase information

$\Sigma$  This section outlines how feature measurements are localised using an approach which is invariant to intensity amplitude (Venkatesh & Owens, 1990; Kovesi, 1996). This is a particularly attractive feature for echogram segmentation where finding global intensity thresholds that work well is problematic. The approach detailed in this section was implemented by Miguel Mulet-Parada, who applied the technique to 2-dimensional images Mulet-Parada & Noble (1999). Since the measurement process in the tracker is based on searching using 1-dimensional normals, Mulet-Parada re-implemented the approach to work on 1-dimensional data. This was then incorporated into the tracker. Further details of the approach are described in Mulet-Parada & Noble (1999); Mulet-Parada (1999).

After post-processing an image using TDLMS/IBS filtering the endocardial border resembles a large scale smooth step edge corrupted by small scale noise, Figure 5.12 (c). A step edge can be defined in terms of the phase values of its frequency components. Venkatesh & Owens (1990)

show that different 1D image features possess different local phase signatures, where local phase is defined in terms of the local Fourier expansion of the signal  $F(x)$  and its Hilbert transform  $H(x)$ , where,

$$\begin{aligned} F(x) &= \sum (a_n \sin(n\omega x + \phi)), \\ H(x) &= \sum (a_n \cos(n\omega x + \phi)). \end{aligned}$$

Here  $\omega$  is the signal frequency,  $a_n$  the frequency component amplitude,  $\phi$  a phase angle and  $n$  the  $n$ 'th Fourier component. The *local phase* of the signal  $E(x) = F(x) + jH(x)$  is then defined as,

$$\text{Arg}(E(x)) = \tan^{-1}(F(x)/H(x)).$$

In further work, Mulet-Parada & Noble (1999) consider avoiding the use of the non-linear arc-tan function, and also using a noise threshold. We tried this for our 1-dimensional profile normals, but the results were not improved.

In practice,  $E(x)$  can be estimated by computing an alternative form  $\hat{E}(x) = f(x) - jh(x)$  where  $f(x)$  and  $h(x)$  are the result of convolving the original signal with a quadrature filter pair formed by a symmetric (even) and an anti-symmetric (odd) filter. In their work, Venkatesh and Owens use Gabor wavelets as quadrature filters. They show that the output of these filters suffices to uniquely discriminate between different features, such as positive and negative edges, ridges and valleys, all of which possess different phase signatures. For example, a positive edge is uniquely defined as a zero-crossing of the output of the even filter and a maxima in the output of the odd filter. This gives a local phase signature of zero radians. Similarly, a negative edge gives a maximum in local phase ( $\pi$ ). Thus, in our application, consider taking 1D normals traced along the cavity boundary, starting from the interior of the ventricle outwards. A (positive going) edge would be located as a point of zero phase in the local phase representation.

Venkatesh and Owen's analysis applies to idealised image feature characterisation in the absence of image noise. Studying the properties of local phase measurements over scales allows us to adapt this scheme to localise edge features in a noisy profile (Kovesi, 1996). The wavelets we use are based on the Log-Gabor transfer function,  $G(\omega)$ , which is defined in the frequency domain by,

$$G(\omega) = \exp - \frac{(\log(\omega/\omega_0))^2}{2(\log(\kappa/\omega_0))^2}. \quad (5.4)$$

Here  $\kappa$  is related to the bandwidth of the filter and  $\omega_0$  the centre frequency of the filter. For the

two octave filters used in this work  $\kappa = 0.55$ . The advantage of using this representation is that the Log-Gabor even function has zero power for negative frequencies.

Consider taking a one-dimensional profile across an echogram of a left ventricle, Figure 5.13 (top). A scalogram can be derived from the output of a sequence of log-Gabor filters of increasing spread (scale), Figure 5.13 (bottom). In this figure, the absolute value of the local phase is plotted so that black corresponds to zero local phase, white corresponds to  $\pi$ . Observe that at small scales many features corresponding to these extreme values are found but that as the filter scale increases the number of features reduces to a few straight bands that correspond to the epicardial (1), to the endocardial (2,3) and to the pericardial (4) edges. It is important to stress that these features give a localised response over a wide range of scales. Thus, unlike other features, the centre of a step does not move (i.e. is stable) over a wide range of scales as long as it is 'significant' in the region of filter support. This contrasts with the poor localisation performance observed for a gradient-based operator at large scales. Obviously if the filter width is very large, two steps will fuse into a roof and the localisation of the phase signature will be lost. Nevertheless, provided the scale of the filter is reasonable, localisation is preserved. This is shown in Figure 5.13 (bottom) by the straight vertical lines corresponding to the endocardial borders (2,3).




Figure 5.13: **Scalogram**. A 1D profile across the left ventricle (top); local phase scalogram (bottom). Observe the stability of phase information over scale.

From the scalogram it is clear that noise and signal structure occupy distinct regions of scale-space. The two regions can be separated by a noise threshold set at the location where most small scale features due to noise disappear. This allows us to select a region in scale space where the main features are detected but where noise effects are insignificant. This noise floor

can be set experimentally with respect to the significance of the features that we want to detect. In our application, we assume that only one target is present along each normal and use the average of the local phase results over three scales above the noise threshold to localise a candidate boundary point. Edge points are then localised as minima in the averaged local phase measure under a pre-defined threshold (points near 0 degrees). In our implementation we apply 1D Log-Gabor filters (Kovesi, 1996) because these have no DC component and allow for large scales to be used without spreading into negative frequencies.

Figures 5.11 (a) and 5.11 (b) show results of the two feature extraction schemes for identical frames at different points in a cardiac cycle. Note that measurements have been deliberately ignored in the region of the mitral valve. This is to enable the tracker to follow the valve during heart contraction in Figure 5.11 (b). In the case of the gradient-based tracker it was necessary to use these (poor) measurements as they were (unfortunately) the only measurements that provided any information about right wall motion.

Frames in Figure 5.11 (a) show how a gradient-based finds very few features along the septum (right wall) and in regions of low image intensity. By contrast, using local phase measurements, a large number of features are detected along the contour, Figure 5.11 (b). These features give a larger measurement set on which the tracker can base a good estimate of the update to the prediction.



### 5.3.2 Comparing feature detection schemes

In this section we compare the performance of the tracking algorithm using (1) gradient-based feature detection on the original data (the original visual image tracker); (2) gradient-based feature detection on TDLMS/IBS enhanced images; and (3) local-phase based feature detection on TDLMS/IBS enhanced images.

A frame-by-frame visual comparison of tracking using the three approaches shows that the reliability of detection of boundaries was best using method (3). This is shown in Appendix F.2.

Quantifying the degree of improvement is difficult because ground truth is not available by which to compare the algorithms. However, since the state vectors can be estimated, we can look for consistency of state vector component trajectories over a number of cycles as a measure of algorithm robustness to measurement noise.

Figure 5.14 shows coefficient plots based on tracking the data in Figures 5.4 , 5.9 and 5.12. The original data (light curve) and the data pre-filtered by the spatio-temporal boundary en-

hancement method with gradient-based localisation (dark curve). Observe that the plots are less noisy for the algorithm which uses the spatio-temporal filtering approach. Figure 5.16 compares the shape-space component vectors for the two spatio-temporal tracking algorithms which use different localisation methods. Observe that the plots for the local-phase localisation are more consistent. This shows that the local-phase based localisation scheme gives a more stable tracking performance. In this case, we used an affine shape-space for comparison. The reason for this was that the affine components are directly interpretable.

It is possible to compare a clinician's manual segmentation with contours from the tracker. This idea was used by Chalana & Kim (1997) . But to do a systematic evaluation is very time consuming and was not felt to be of value at the point of discussion. Chapter 8 does address this when a clinical protocol is defined.




Figure 5.14: **Components plots with enhanced images.** Comparison of component plots using an affine shape-space for tracking on ultrasound data: original data (—) and enhanced data (—).




Figure 5.15: **Components plots with enhanced images.** Comparison of component plots using a PCA shape-space for tracking on ultrasound data: original data (—) and enhanced data (—).




Figure 5.16: **Components plots with enhanced feature detection.** Comparison of component plots using an affine shape-space for tracking on enhanced ultrasound data. Shown are the gradient-based localisation (—) and local-phase based localisation (—).

## 5.4 Conclusions

This chapter has evaluated the performance of a visual tracker on echocardiographic images. It was shown that the shape-deformation of a left ventricular boundary should be modelled using a PCA shape-space for good tracking. An affine basis was shown both empirically and mathematically to be incapable of providing the correct deformations. We also showed that simply applying a standard visual image boundary detector to obtain measurements was not sufficient. Instead, we applied intensity-invariant based feature detection and showed how it enhanced feature localisation and hence overall performance of the echocardiographic tracker was improved.

## CHAPTER 6

---

### Interpretational shape-spaces

---

In Chapter 4 we introduced the notion of a shape-space as a way of modelling shape deformation. We then took this idea and showed that tracking the endocardial boundary was dependent on several factors, not least the appropriate choice of shape-space. Up till now, we have talked about a shape-space designed to model shape deformation, but specifically used for tracking purposes. In this chapter, we shall extend the notion of a shape-space to cover tracking interpretation.

In the previous chapter, we suggested that plots of the shape-space vector components could potentially be used to look at heart function. As we shall see in this chapter, this is not the best way of doing this. We propose, instead to use and introduce the *interpretational shape-space*. An interpretational shape-space is a shape-space that is designed to transfer tracking results into a more clinically meaningful plot of the shape-space vector components. In our case, meaningful means more useful clinically, although, as we shall see, the theory can be applied to any tracking. The ideas in this chapter were developed in collaboration with the Departments of Cardiology and Nuclear Cardiology at the Oxford John Radcliffe hospital, Oxford, UK.

The central idea is based on the concept that the shape-space that we track in, need *not* be, and in fact is not, the same space that we interpret the tracking from.

Bascle & Blake (1998) applied a related idea in the context of automated re-animation, based on facial tracking.

In Section 6.1 we motivate the clinical basis for wall tracking as an assessment of global and

regional heart function. We discuss why regional wall motion is important. The limited amount of literature on this topic is *qualitative*. Section 6.2 presents first a 16-segment model of the heart, and second a qualitative standard scoring scheme currently used to assess wall motion and thickening. These can be combined to assess regional wall motion. In Section 6.3 we introduce and develop the idea of an interpretational shape-space to *quantitatively* assess wall motion.

We develop this idea for semi-automatic quantitative wall scoring in Section 6.4. Section 6.5 considers an alternative to interpretational shape-spaces based on 'key-frames' to build a shape basis that is directly clinically meaningful. We show that tracking in this basis is poor, because the basis allows for too much deformation. In Section 6.6 we discuss some colour-coded plots that can be used to assess wall motion, and visually compare their performance with our method. Section 6.7 summarises the key points of this chapter.



## 6.1 Heart Function

As we mentioned in Chapters 1 and 3, the assessment of left ventricular function is an essential component of the evaluation of any patient with known or suspected heart disease. As part of this assessment, the clinician will examine global systolic function and regional systolic function. We consider these below.

### 6.1.1 Global systolic function

The ejection fraction of the left ventricular cavity has been traditionally used as indicator of global heart abnormality. The ejection fraction represents the fraction (or percentage) of left ventricular diastolic volume (in 3-dimensions or 2-dimensional extrapolation) that is ejected in systole (refer to Table 6.1). A normal heart has an ejection fraction of between 0.4 and 0.8 (Feigenbaum, 1994, page 669). Much research has sought to automate calculation of the ejection fraction (Germano *et al.*, 1995; Giachetti, 1998).

The idea that automating the calculation of the ejection fraction will make a significant difference to the clinician is a myth — clinicians make a reasonable 'eyeball' estimate of it from normal two-dimensional echocardiography without making any measurements (Amico *et al.*, 1989; Mueller *et al.*, 1991; Wong *et al.*, 1991). In fact, whether an ejection fraction is 0.5 or 0.55 makes little practical difference — the only difference is if it is between 0.4 and 0.8 (Feigenbaum, 1994, page 669). A precise estimate of the ejection fraction is more applicable in a research rather than clinical setting. For example, a very accurate estimate may be needed to compare drug effects.

Presently, estimates of the volume of the left ventricular cavity are based on combining information from one or more single-plane views. A number of methods have been developed to make a volume estimation based on this information (Folland *et al.*, 1979; Kan *et al.*, 1981; Gordan *et al.*, 1983; Zile *et al.*, 1992). Methods for estimating the volume, for example the 'prolate ellipse' method outlined in Feigenbaum (1994, page 138), rely unrealistically on major assumptions about the geometry of the left ventricle. A more common technique is to use Simpson's rule (Rogers *et al.*, 1979), which calculates the volume by adding together the volumes of a series of parallel slices. Difficulties with this approach are that, it is not always possible to obtain multiple slices, due to occluding factors such as the ribs, the slices must be parallel, and also densely sampled.

The technique that is recommended by The American Society of Echocardiography, is that the volume should be calculated from the apical two-chamber and four-chamber views — this is called the biplane method (Schiller *et al.*, 1989). The method is based on orthogonal views of the chamber. In the case, where only one apical view is available — as in our case — the single plane method, can be used (Schiller *et al.*, 1989). The calculation of the volume ( $V$ ) is then as follows:

$$V = \frac{0.85 \cdot A^2}{L}, \quad (6.1)$$

where,  $A$  is the area of the left ventricular cavity, and  $L$  is the length of the ventricular cavity, as shown in Figure 6.1.

Figure 6.2 shows application of the single plane method to the data in Figures 5.6 and 5.10. For this example, the ejection fraction is 0.63, which is in the range of a normal heart (Feigenbaum, 1994, page 669).

Clearly, volume calculations, as in Equation 6.1, rely on unrealistic geometrical assumptions. Moreover, because the equation relies on the square of the area, errors in the calculation of the area are squared in the calculation of the volume. An alternative global measurement that can be used is the cavity area of a projected 2-dimensional slice. The potential problem of area calculations is that they can be affected by foreshortening of the chambers. This is because, an ultrasound image is a 2-dimensional projection of a 3-dimensional object. This is also true of volume estimates.

Measurement	Definition	Normal range
LV ejection fraction (EF)	$EF = \frac{LV \text{ diastolic volume} - LV \text{ systolic volume}}{LV \text{ diastolic volume}}$	0.4–0.8
Fractional area change (FAC)	$FAC = \frac{\text{End-diastolic area} - \text{End-systolic area}}{\text{End-diastolic area}}$	0.36–0.64

Table 6.1: **Global measurements, definitions and normal ranges.** Some global measurements that can be obtained from two-dimensional echocardiography. The left ventricular diameters and fractional shortening provide an assessment of the size and systolic function of the base of the left ventricle. The normal ranges are as given in Feigenbaum (1994, page 669).




Figure 6.1: **Single plane volume calculation.** The volume of the left ventricular cavity can be measured using the single plane area length method, as shown by Schiller *et al.* (1989). The volume ( $V$ ) is given by  $V = \frac{0.85A^2}{L}$ .  $A$  is the area of the left ventricular cavity, and  $L$  is the length of the ventricular cavity.

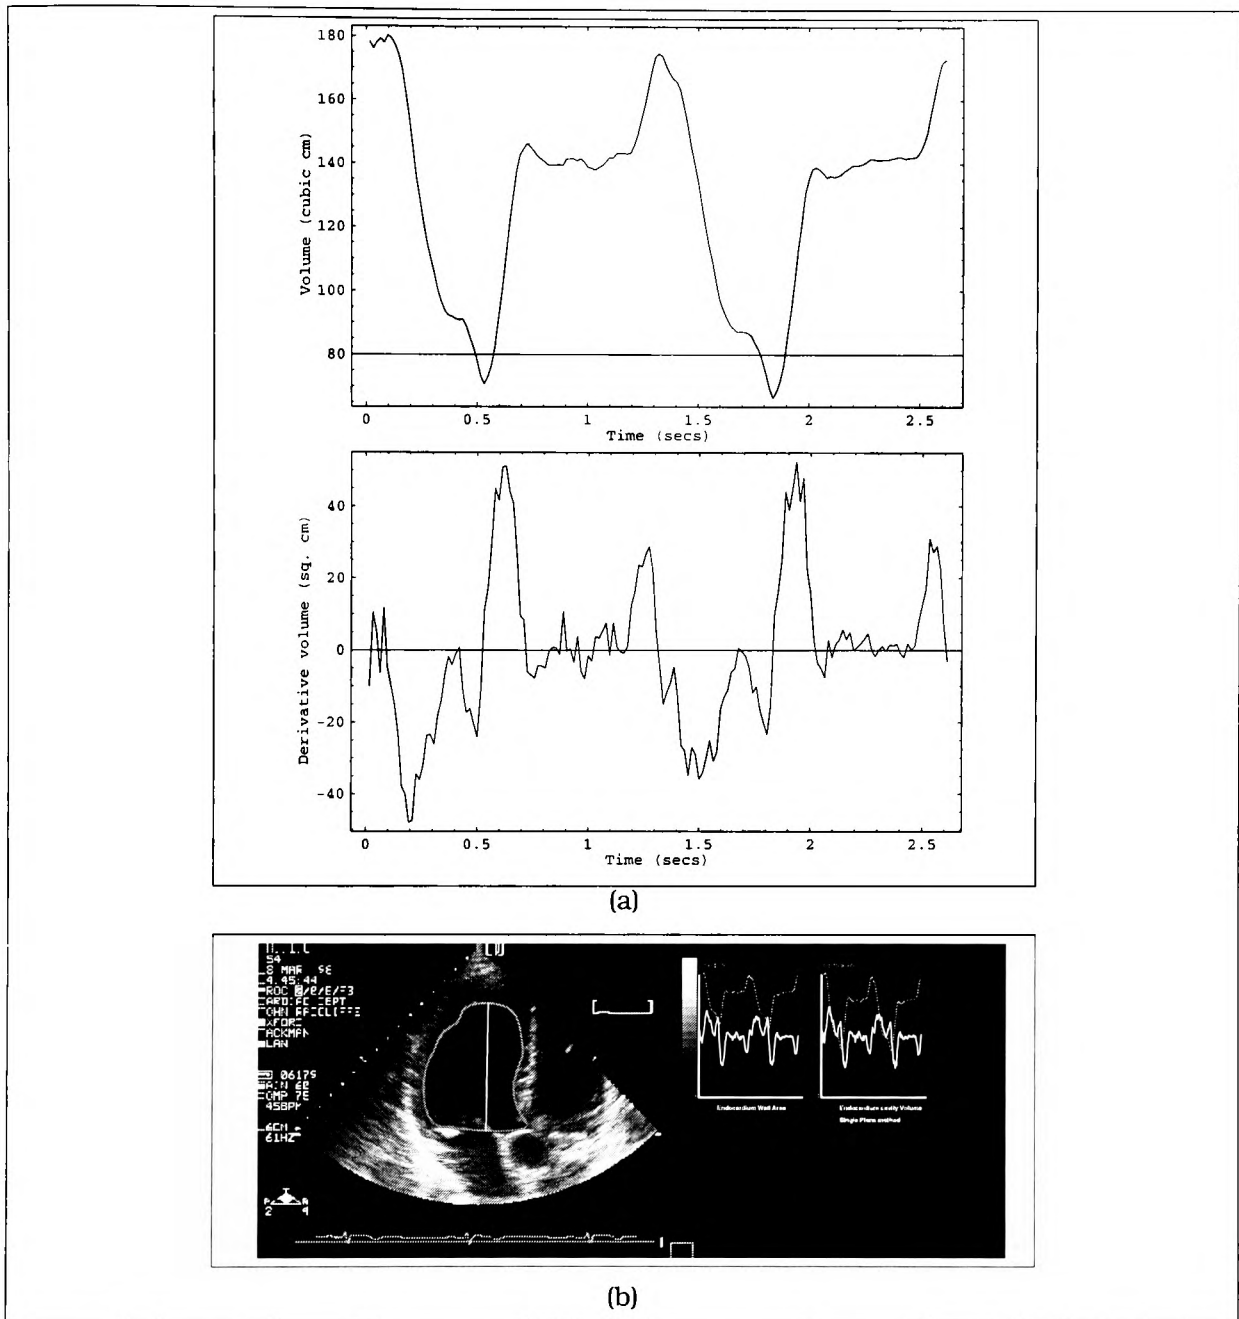


Figure 6.2: **Single plane volume calculation example.** The volume and its derivative of the left ventricular cavity measured using the single plane area length method, as shown by Schiller *et al.* (1989). In this example, the calculated ejection fraction, which represents the fraction of left ventricular diastolic volume that is ejected in systole, was 0.63, which is in the range of a normal heart (Feigenbaum, 1994, page 669). (a) Volume ( $\text{cm}^3$ ) plotted against time for two cycles. (b) Tracker interface.

### Left ventricular area

Traditionally, the measures most used by clinicians to assess wall motion have been:

1. The area of the left ventricular cavity.
2. The first derivative of the area with respect to time. This is used as an index of the rate of left ventricular cavity shrinkage and the rate of cavity expansion during systole and diastole,

respectively.

3. The fractional area change (refer to Table 6.1).

Peréz *et al.* (1992c,b) present a clinical tool for on-line quantification of left ventricular size and function. Their work developed an *“automatic boundary detection approach that depicts and tracks the interface between endocardium and blood in real-time and quantitates cardiac chambers area and function with instantaneous graphics display of the data”*.

In their paper, images are shown of the tracked endocardial borders, the area, its derivative, and fractional area change. The problem is that *all* these measures refer to assessment of global heart function, and are not directly applicable to regional heart function.

For example, Figure 6.3 shows a plot of the area and its first derivative, of the left ventricular cavity for the data in Figure 5.6. Note that for plotting purposes, the derivatives have been scaled by a factor of 5. The calculated fractional area change is 0.45. The area plot and fractional area change are both in the range of normality (Feigenbaum, 1994, page 669).



we will consider in the next chapter.

Abnormal endocardial wall motion of the ischemic segment is one of the main ways clinicians detect ischemic muscle action. Kerber & Abboud (1973) note that when the muscle becomes ischemic, its motion is altered almost immediately.

An early technique for evaluating left ventricular function (excursion and thickening) used manual measurements based on (1-dimensional) M-mode echocardiography (with its high sampling rate) (Kerber & Abboud, 1973; Jacobs *et al.*, 1973) and associated measurements. Abnormalities in endocardial wall motion (Kerber *et al.*, 1975, 1976), as well as systolic wall thickening (Sasayama *et al.*, 1976; Corya *et al.*, 1977; Kerber *et al.*, 1979) have been related to coronary perfusion.

However, the one-dimensional nature of the M-mode echocardiography limits its usefulness. Although, its measurements are still used, they have limitations. An obvious limitation can occur with patients with coronary artery disease, where the apex may be dilated and akinetic, but the M-mode measurements totally normal — a severe limitation in diagnosing patients with regional wall abnormalities.

The alternative is to move to two-dimensional echocardiography. This is now the principal (Mann *et al.*, 1986) technique used for analysing wall motion and thickening around the entire left ventricle. It is more accurate than M-mode, because the M-mode beam does not always cut across walls perpendicularly.

Two-dimensional echocardiography provides the technique to image the endocardial boundary, but the regional wall function needs to be analysed in some way. We consider this in the next section.

## 6.2 The 16-segment model of the heart

Regional function can be assessed in many ways, most commonly by observing wall motion.

One way to qualitatively assess wall motion in real-time is to indicate which segments are moving abnormally. The obvious problem with a total qualitative assessment is that it is subjective, and that the clinician is more likely to notice the most striking abnormalities and may overlook subtle effects taking place which might be of equal importance.

A more systematic way to assess regional wall motion is to divide the ventricular boundary into segments, corresponding to significant blood supplies to the heart surface. Various groups have used a different number of segments, ranging from 9 to 20 (Stamm *et al.*, 1983; Previtali *et al.*,

1991, 1993; Mazeika *et al.*, 1992; Hecht *et al.*, 1993; Beleslin *et al.*, 1994; Senior & Lahiri, 1995).

The generally accepted model adopted by The American Society of Echocardiography (ASE) is a 16-segment approach (Bourdillon *et al.*, 1989; Schiller *et al.*, 1989).




Figure 6.4: **16-segment anatomical model of the heart.** The 16-segment anatomical model of the heart proposed by The American Society of Echocardiography. Depicted clockwise from top left are parasternal long and short axis, and 2 and 4 chamber views. Figure based on Hoffmann *et al.* (1993).

With this anatomical model of the left ventricle now 'defined', it remains to devise a scoring scheme to assess wall motion and thickening. As with the choice of heart model, many groups use different scoring schemes.

One of the earliest schemes by Limacher *et al.* (1983) graded walls on the “*subjective impression of inward motion of the endocardial echo towards the centre of the left ventricular cavity and the degree of thickening of the myocardium*”. Limacher goes on to describe an average score based on each segmental score.

A scoring scheme proposed by the American Society of Echocardiography has been adopted by a number of groups (Krahwinkel *et al.*, 1997; Cokkinos *et al.*, 1998) and is shown in Table 6.2. A normal segment is given the value 1. Abnormal segments have higher numbers. Occasionally, 5 is awarded to an aneurysmal region. The clinician’s score is based on the myocardial thickening of each segment. If however, the myocardium thickening is not clearly visible, but the endocardial border is, then the clinician’s score is based on the observed endocardial excursion of each segment.

Score	Grading	Characterized by
1	Normal	A uniform increase in wall excursion and thickening
2	Hypokinetic	A reduced (< 5 mm) inward systolic wall motion
3	Akinetic	An absence of inward motion and thickening
4	Dyskinetic	Systolic thinning and outward systolic wall motion

Table 6.2: **Echocardiography scoring system.** The scoring system used in echocardiography. A normal segment is given the value 1 with abnormal segments having higher numbers. The clinician’s score is based on the myocardial thickening of each segment. If however, the myocardium thickening is not clearly visible, but the endocardial border is, then the clinician’s score is based on the endocardial excursion of each segment.

Although the introduction of the 16-segment model forces the clinician to assess each segment separately, the scoring scheme is still subjective. This means that clinical reporting of echocardiography examinations is operator-dependent and largely qualitative. Moreover, the index is based on a (discrete) ordinal scale — a segment can only be classified as normal or abnormal. This has the effect that the clinician must be convinced a segment is abnormal before they award a score greater than one. The net result is that subtle wall motion abnormalities may be lost within the discrete scoring system.

In the next sections, we shall develop an approach to assess wall motion quantitatively, and devise an automatic quantitative scoring scheme. Fundamental to the approach is the way in which we decompose the tracking results using an interpretational shape-space. We consider this next.



### 6.3 Interpretational shape-spaces



Recall from Section 4.3.1, that we defined a tracking shape-space by (Equation 4.10),

$$\mathbf{Q} = \mathbf{Q}_0 + \mathbf{W}_{PCA} \mathbf{X}_{PCA}. \quad (6.2)$$

For clarity, we have introduced the subscript 'PCA' to indicate that the tracking is based on using a PCA shape-space. We recover the shape-space vector using the pseudo inverse,  $\mathbf{W}_{PCA}^\dagger$  of the shape-space we track in,  $\mathbf{W}_{PCA}$ :

$$\mathbf{X}_{PCA} = \mathbf{W}_{PCA}^\dagger (\mathbf{Q} - \mathbf{Q}_0). \quad (6.3)$$

Figure 6.5 shows an example of PCA based component plots. Note that it is not clear how to relate the plots to any underlying physical deformation process, and hence clinical cause.

Each contour can be expressed as a weighted sum of the shape-space vectors, where each weight is determined by the significant PCA component – this does not necessarily correlate with clinical significance. So for example, the first deformation mode in Figure 6.5 (a) appears to be a scaling of inferior part of the left ventricular boundary (refer to Figure 6.4). The corresponding shape-space vector is plotted as 'Component 1' in Figure 6.5 (b).

All we can say is that the expansion of that part of the left ventricular wall is periodic. Moreover, because this first shape-space deformation affects more than just a regionalised part of the left ventricular wall — the anterior wall is also weighted — we probably can't even say that. A similar conclusion to this can be drawn from the plots of the shape-vector in Figure 5.9 that correspond to the deformation modes in Figure 5.4.




Figure 6.5: **Shape-space vector components.** (a) Shape-space deformation modes and, (b) the corresponding shape-space vector components. Each shape-vector component acts upon its corresponding shape-space deformation vector to obtain the contour at time  $t$ .

We can overcome this problem in a few ways:

- We can assume an affine model, which ensures that each shape-vector component is readily interpretable. But, as we saw in Section 5.1.3, an affine model is not capable of modelling the complex non-rigid motion of the left ventricular boundary, and does not generate a clinically meaningful description.
- We can define a new shape-space, analogous to the affine space, where each component is readily interpretable, but the space, as whole, has the ability to model the heart deformation. For example, we can take an approach using 'key-frames' as in Section 4.3.1, but as we shall see in Section 6.5, this allows for too much deformation.
- We can use one shape-space for tracking, and then decompose the results using another interpretational shape-space.

We choose the last approach, and use a PCA shape-space for tracking, which has the advantage of being finely tuned to the object's deformation and computationally convenient. We then decompose the tracking results using another shape-space, called an *interpretational shape-space*.

To illustrate the approach, suppose as illustrated in Figure 6.3, we tracked an object using 12 control points in a shape-space of dimension 6. Then each of the 6 column vectors in the shape-space,  $\mathbf{W}$ , will be of dimension 24 (2 times the number of control points). Correspondingly, each of the 24 column vectors in  $\mathbf{W}^\dagger$ , will be of dimension 6. Suppose we now define the interpretational space,  $\mathbf{W}^\dagger$ , so that the first row vector is  $(1/24, 1/24, \dots, 1/24)$ , but all the remaining 23 row vectors are  $(0, 0, \dots, 0)$ . Then the resulting component magnitudes based on Equation 6.3 will all be zero except for the first, which will be the average of the control points positions.

In the apical 2-chamber view, the left ventricle is split into 6 segments, as shown in Figure 6.4. We can use this regional model to interpret the tracking. To achieve this, we can then choose an interpretational shape-space so that the pseudo-inverse,  $\mathbf{W}^{\dagger*}$ , is defined as,

$$\mathbf{W}_{CLIN}^\dagger = \begin{pmatrix} \frac{1}{4} & \frac{1}{4} & \frac{1}{4} & \frac{1}{4} & 0 & 0 & 0 & 0 & \dots & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & \frac{1}{4} & \frac{1}{4} & \frac{1}{4} & \frac{1}{4} & \dots & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & \ddots & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & \dots & 0 & 0 & 0 & 0 & \frac{1}{4} & \frac{1}{4} & \frac{1}{4} & \frac{1}{4} \end{pmatrix}, \quad (6.4)$$

and,

$$\mathbf{X}_{CLIN} = \mathbf{W}_{CLIN}^\dagger (\mathbf{Q} - \mathbf{Q}_0). \quad (6.5)$$

Suppose we track an object with 12 control points in a shape-space of dimension 6. Then, the shape-space equation is given by,

$$\begin{pmatrix} q_1^x \\ \vdots \\ q_{12}^x \\ q_1^y \\ \vdots \\ q_{12}^y \end{pmatrix} = \begin{pmatrix} q_{0,1}^x \\ \vdots \\ q_{0,12}^x \\ q_{0,1}^y \\ \vdots \\ q_{0,12}^y \end{pmatrix} + (w_1 \ w_2 \ w_3 \ w_4 \ w_5 \ w_6) \begin{pmatrix} x_1 \\ x_2 \\ x_3 \\ x_4 \\ x_5 \\ x_6 \end{pmatrix}$$

Alternatively in matrix notation,

$$\begin{aligned} \mathbf{Q}_{24 \times 1} &= \mathbf{Q}_0 + \mathbf{W}_{24 \times 6} \mathbf{X}_{6 \times 1}, \text{ and,} \\ \mathbf{X}_{6 \times 1} &= \mathbf{W}^\dagger_{6 \times 24} (\mathbf{Q} - \mathbf{Q}_0)_{24 \times 1}. \end{aligned}$$

Now choose,

$$\mathbf{W}^\dagger = \begin{pmatrix} 1/24 & 1/24 & \dots & 1/24 \\ 0 & 0 & \dots & 0 \\ \vdots & \vdots & \vdots & \vdots \\ 0 & 0 & \dots & 0 \end{pmatrix}.$$

Then,

$$x_1 = \frac{1}{24} \sum (\mathbf{Q} - \mathbf{Q}_0) \quad \text{and,} \quad x_2, \dots, x_6 = 0.$$

Figure 6.6: **Interpretational pseudo-inverses.** Tracking an object using 12 control points in a shape-space of dimension 6. Each column vectors in the shape-space,  $\mathbf{W}$ , is of dimension 24. Each of the 24 column vectors in  $\mathbf{W}^\dagger$ , has dimension 6. Suppose we choose the interpretational shape-space so that its first row vector was  $(1/24, 1/24, \dots, 1/24)$ , but all the remaining (23) row vectors were  $(0, 0, \dots, 0)$ . Then, the shape-space components will all 0 except for the first, which will just be the average of the control points position.

For clarity we have added the subscript 'CLIN' to indicate a 'clinical' interpretation.  $\mathbf{W}^\dagger$  has dimension,  $N_X \times N_q$ ; in our case, the left ventricular boundary is modelled using a B-spline of 24 control points (Figure 4.1), and a shape-space of dimension 6. Hence,  $\mathbf{W}^\dagger$  has dimension,  $6 \times 24$ .

In fact, we divide the template of the left ventricular boundary into the 16-segment model outlined in Section 6.2. In the model, the apical 2-chamber view has 6 segments. For our analysis, we use 4 control points in each segment, as shown in Figure 6.7. We choose to weight the control points in each segment equally, as  $1/\text{number of control points}$ .




Figure 6.7: **Interpretational space.** The apical 2-chamber view of the left ventricular boundary is divided into 6 segments, in accordance with the 16-segment outlines in Section 6.2. We model this boundary with a B-spline of 24 control points, with 4 control points in each segment. The interpretational shape-space, given in Equations 6.4 and 6.5 then decomposes the tracking results into the local average motion of each segment. This ability to quantitatively analyse regional wall function is clinically very useful. Once tracking has begun, it is not guaranteed that the six tracking segments will precisely match the anatomical regions, but in practice they are very close. This is similar to a clinician scoring the six segments by visual inspection.

Therefore, Equations 6.4 and 6.5 simply decompose the tracking into 6 regional time-varying shape-vectors components. The components represent the 'average' local motion of each segment. The advantage of this approach is that now the shape-space vector components are meaningful — they represent the average regional function of the left ventricular boundary.

A problem with Equation 6.5 is that wall excursion is measured from a fixed external frame of reference. This means that it ignores the heart's translational (and rotational) motion. One way of overcoming this is to use a floating of internal frame of reference that attempts to control for or correct for the translation (and rotation). In general, as the translation (and rotational) motion increases, fixed reference systems become less reliable and floating systems more reliable (Wiske *et al.*, 1990).

To address this, we subtract out the centroid, so that,

$$\mathbf{X}_{CLIN} = \mathbf{W}_{CLIN}^{\dagger} ((\mathbf{Q} - \bar{\mathbf{Q}}) - (\mathbf{Q}_0 - \bar{\mathbf{Q}}_0)) \quad (6.6)$$

Although, the floating system we have used may have advantages in monitoring the normal cardiac contraction pattern when cardiac translation is significant, it can have disadvantages when assessing regional wall motion in patients with abnormalities (Parisi *et al.*, 1981; Schnittger *et al.*, 1984). Moreover, floating systems introduce unavoidable variability of their own (Katz *et al.*, 1996). A discussion of these two systems and ways of measuring regional wall motion can be found in Gelberg *et al.* (1979); Sheehan *et al.* (1986); Wiske *et al.* (1990), and Katz *et al.* (1996).



To show the value of this decomposition, we can compare the components plots based on a normal PCA shape-space pseudo-inverse,  $\mathbf{W}_{PCA}^\dagger$ , with that from using the new interpretational space,  $\mathbf{W}_{CLIN}^\dagger$  — Figure 6.8 (the figure is plotted so that the component magnitude axis is fixed). Figure 6.8 (a) represents the shape-vector components based on the pseudo-inverse of the tracking shape-space. In comparison, Figure 6.8 (b) represents the shape-vector components based on the new interpretational shape-space. We can now clearly relate the plots of the shape-vector components in the interpretational space to the underlying local motion of the left ventricular boundary. Although Figure 6.8 (a) shows periodicity in principally the first three of the components. It is very difficult to relate these plots to the underlying motion of the left ventricular boundary. In comparison, Figure 6.8 (b) shows that all six of the components of the new space are periodic. We can see that the basal anterior, mid anterior, mid inferior and basal inferior segments all move normally. The observed smaller movement in the basal inferior is in accordance with normal heart function. However, we can see that the apical inferior and apical anterior segments, although moving periodically, have a reduced endocardial wall excursions. This observation is exactly in accordance with the diagnosis that this patient has a myocardial infarct in the apical region. This is shown by the way in which we decompose the tracking results using our interpretational shape-space.

Of course it is not necessarily true that any ischemic region would only fall in one segment of our model alone. In fact, in the above example in Figure 6.8, the ischemic region covers roughly half of the apical inferior segment and half of the apical anterior segment. We need some way of trying to highlight this. We consider this next.




Figure 6.8: **Comparison of tracking decompositions.** Shape-vector component plots based on (a) a PCA shape-space and (b) an interpretational shape-space. We can see that decomposing the tracking results using an interpretational shape-space leads to more clinically meaningful component plots. Note that the component magnitude scale has been fixed for each sub-plot.

### 6.3.1 Segmental variation in wall excursion

Recall from Section 6.3 that one of the ways in which we can view the interpretational shape-space is as a local average of the motion in a segment. The pseudo-inverse,  $W^\dagger$ , averages the difference between the curve at any time with the initial template curve. To understand whether the ischemia is homogeneous throughout a segment or whether it covers only part of the segment, i.e. heterogeneous, we need to determine a measure of the variation in each segment. One way of doing this is to calculate the standard deviation of control point movement in each segment. Strictly speaking, calculating the standard deviation of control points is not the same as the calculating the standard deviation of the points along the contour, although in practice, there is little difference. For each segment, we therefore obtain two summary measures – the mean and standard deviation — of the endocardial wall excursion.

We can then generate a 95% (2 standard deviations) confidence interval (error plot) for the interpretational shape-space vector components for each of the six segments. An example is shown in Figure 6.9. The starting point for this dataset is diastole, with the heart entering systole after around a tenth of a second. We can see that the variation (standard deviation) in the segments is approximately the same for the basal anterior, mid inferior and basal inferior segments. In comparison, the apical anterior and apical inferior segments show a very noticeable increase in

variation during the systolic stage of the heart cycle. This implies that not all of the apical inferior and apical anterior segments are ischemic. More importantly it shows that this difference is more evident during systole, when the heart is contracting. This observation is again in accordance with the literature.

We now turn our attention to proposing a novel method for an automatic scheme to 'score' each segment based on the interpretational space. Table 6.2 depicts the standard scoring scheme presently used throughout the literature.

We will perform a more thorough clinical evaluation in Chapter 8.




Figure 6.9: **Interpretational shape-space variation.** 95% confidence intervals for the 6 interpretational shape-space vector components. The starting point for this dataset is diastole. The heart enters systole after around a tenth of a second. The variation in the segments movement is approximately the same for the basal anterior, mid inferior and basal inferior segments. In comparison, the apical anterior and apical inferior segments show a very noticeable increase in variation during the systolic stage of the heart cycle. This observation shows that not all of the apical inferior and apical anterior segments are ischemic. Moreover, this is more evident when the heart is contracting. (a) Segmental wall excursion is plotted (—), along with 95% confidence intervals (---). (b) Tracker interface.

## 6.4 Segmental scoring of wall excursion

Recall from Section 6.2, that a clinician scores the myocardial thickening of each segment by an ordinal scale (Table 6.2).

In this section, a quantitative scoring scheme is developed based on the endocardial wall excursion. To do this, we shall use the interpretational shape-space vectors components, developed earlier in this chapter.

The endocardial wall excursion is modelled by calculating the maximum displacement of the contour from its initial position. The excursions are then normalised with respect to the maximal movement in the 6 segments (recall we are using the apical 2-chamber view). Only the relative displacements of the contour are important. The reference is arbitrary.

Then the *maximal endocardial wall excursion* is calculated as the maximum (signed) excursion of the contour minus the minimum (signed) excursion of the contour. Excursion is taken with reference to the initial template. This is calculated for the whole sequence, not by cycle. This is illustrated in Figure 6.10.

As an example, Figure 6.11 (a) shows the maximal endocardial wall excursion (in cm) for the data in Figure 6.8 (b). The relative magnitude of these values is consistent with the diagnosis of a myocardial infarct in the apical inferior and apical anterior. Moreover, the reduced motion of the basal inferior section is also recognisable.




Figure 6.10: **Schematic maximal endocardial wall excursion.** A schematic plot of a shape-space vector in interpretation space. We use the maximal endocardial wall excursion as a measure of wall excursion. Maximal endocardial wall excursion is calculated as the maximum (signed) excursion of the contour minus the minimum (signed) excursion of the contour. Excursion is taken with reference to the initial template.




Figure 6.11: **Maximal endocardial wall excursion scores for the data in Figure 6.9.** Maximal endocardial wall excursion (in cm) is given by, the maximum excursion relative to the initial template minus the minimum excursion relative to the initial template minus. (a) Maximal endocardial wall excursion by segment. We can normalise these values by the segment with the largest maximal endocardial wall excursion (in cm). (b) Normalised maximal endocardial wall excursion by segment.

### 6.4.1 Scoring scheme normalisation

The maximal endocardial wall scores, discussed in the previous section are normalised by the segment with the largest maximal endocardial wall excursion. That is we simply divide the value of each segment's maximal endocardial wall excursion by the largest of the six segments — Figure 6.11 (b). A similar approach is used in MUGA scanning, where pixel counts are normalised to the pixel with the highest count.

Clearly, there is one extreme case when this method will not work well. This is when every segment is ischemic, then every maximal endocardial wall excursion score will be small, and hence the normalised score will appear totally normal. We could check this by having a minimal allowable maximum and minimum endocardial wall excursion.

We can see from Figure 6.11 (b) that the normalised scores correctly enables the reduced motion in the apical interior and anterior segments to be identified. It also highlights the reduced, although totally normal, motion of the basal inferior segment.



## 6.5 Tracking in a clinical space

The original motivation for defining an interpretation space arose because we wanted to decompose the tracking results into a shape-space vector that was clinically meaningful. The problem we faced was that the vectors in the PCA shape-space were not readily interpretable. In comparison, a shape-space, such as affine, does not suffer from this problem. In Section 6.3, we suggested three methods for overcoming this problem. We adopted the last approach, based on

tracking in one shape-space, but interpreting the tracking results in another shape-space — the so called interpretational shape-space.

In this section, we look at the second suggestion made to overcome the problem. This involves creating a shape-space that can be used for tracking and that is clinically meaningful.

The basis is created using the idea of key-frames, outlined in Section 4.3.1. This would have both the advantage of being tuned to the object's deformation, but also each basis vector would be directly interpretable. The shape-space vector components would then have direct clinical meaning.

In this section, we shall explore this idea further, and show that although the idea seems appealing, it allows for too many degrees of freedom in a tracker in practice. The tracker becomes underconstrained, with the result being poor tracking.

We need to answer the question of what should be in the basis? Clearly, both horizontal and vertical translation, will be needed, if only to allow for either the clinician moving his hand (and hence transducer) or the patient moving. During a cardiac cycle, the left ventricular boundary expands and contracts. So intuitively, the basis should contain scaling, so that the whole boundary and also each segment can expand and contract. We shall use the 16-segment model, outlined in Section 6.2, so that for the apical 2-chamber view, 6 basis vectors are need to control the regional scaling. The full dimension of the basis is therefore 8.

The key-frames for the shape-space basis are generated from the template,  $Q_0$ . The horizontal and vertical translation are, as in the affine case, simply  $(1, 0)^T$  and  $(0, 1)^T$ . The identity vector,  $\mathbf{1}$ , and the null vector,  $\mathbf{0}$ , are of length 24 (the number of control points).

The 6 region scaling vectors are  $(Q_1^x, Q_1^y)^T, (Q_2^x, Q_2^y)^T, \dots, (Q_6^x, Q_6^y)^T$ . So the shape-space is given by,

$$\mathbf{W} = \begin{pmatrix} \mathbf{1} & \mathbf{0} & Q_1^x & Q_2^x & Q_3^x & Q_4^x & Q_5^x & Q_6^x \\ \mathbf{0} & \mathbf{1} & Q_1^y & Q_2^y & Q_3^y & Q_4^y & Q_5^y & Q_6^y \\ \underbrace{w_1 \quad w_2}_{\text{Translation}} & & \underbrace{w_3 \quad w_4 \quad w_5 \quad w_6 \quad w_7 \quad w_8}_{\text{Scaling of 6 segments}} & & & & & \end{pmatrix} \quad (6.7)$$

There are (at least) three potential ways to obtain the basis vectors that will scale each segment, based on an initial template  $Q_0$ .

1. The first two techniques expand the initial template,  $Q_0$ , to a new template,  $Q'_0$ , and generate basis vectors from the difference between the configurations. The methods rely on manually 'pulling' the control points of the initial configuration outwards by an arbitrary amount and

then recording the new configuration.

In both cases, the translational basis vectors are obtained in the same way. We simply horizontally and vertically translate the initial template,  $\mathbf{Q}_0$ , and use the difference between the two configurations to obtain the basis vectors.

- (a) Expand (evenly) the initial template,  $\mathbf{Q}_0$ , to a new template,  $\mathbf{Q}'_0$ . Then obtain the vector,  $\mathbf{Q}''_0 = (\mathbf{Q}'_0 - \mathbf{Q}_0)$ , that represents the difference between the two templates. We now generate 6 vectors,  $(\mathbf{w}_3, \mathbf{w}_4, \dots, \mathbf{w}_8)$ , from  $\mathbf{Q}''_0$ , that uniquely represent the scaling of the 6 segments of the ventricle. Recall that there are 4 control points in each segment, so, for example,  $\mathbf{w}_3 = (\mathbf{Q}''_{0,1,x}, \mathbf{Q}''_{0,2,x}, \mathbf{Q}''_{0,3,x}, \mathbf{Q}''_{0,4,x}, 0, \dots, 0, \mathbf{Q}''_{0,1,y}, \mathbf{Q}''_{0,2,y}, \mathbf{Q}''_{0,3,y}, \mathbf{Q}''_{0,4,y}, 0, \dots, 0)^\top$ , and similarly for  $\mathbf{w}_4, \dots, \mathbf{w}_8$ . These vectors are then normalised so that,  $\mathbf{w}_i^\top \mathcal{H} \mathbf{w}_i = 1$ .
- (b) An alternative method to above would be to directly generate the individual basis vectors. For each segment in turn, we manually 'pull' outwards by an arbitrary amount, the 4 control points that represent the respective segment and record the new configuration. We then normalise, as above, the difference between this configuration and the initial configuration of the template.

2. The previous approach dealt with manually altering the initial configuration control points by hand. Another option would be to orthogonally project each point on the contour outwards, and then re-fit the spline, and hence obtain the control points of the new configuration. This could be done for either of the two approaches above.

For the methods above, there is a problem of the continuity across the segment boundaries — i.e. moving any control point affects three spans of the spline, and so moving control points on the boundary between two segments affects both segments.

An 8-dimensional shape-space basis was created using the first technique in the above default experiment. Figure 6.12 shows the contours on 5 consecutive frames (28–32) from an ultrasound sequence. Tracking based on this 'hand-made' shape-space is shown in the top row, whereas the bottom row shows tracking based on a PCA shape-space. The 'hand-made' shape-space allows for too much deformation in the contour, and the tracker becomes unstable. This observation was particularly true in contraction, although in expansion the tracker is able to recover its correct configuration.



We can conclude from this, that the better approach is to track the endocardium in the PCA shape-space, and then interpret the tracking in an interpretational shape-space.




Figure 6.12: **Clinical shape-spaces.** An 8-dimensional clinical shape-space was created using the key-frames idea outlined in Section 4.3.1. The basis vectors were as follows; horizontal and vertical translation and 6 segmental scaling vectors. Tracking using this basis was visually compared to a PCA shape-space. Shown are 5 consecutive frames (28–32) during systole from an ultrasound sequence. (a) ‘Hand-made’ clinical shape-space. (b) PCA shape-space. The ‘hand-made’ clinical shape-space allows or too much deformation in the contour, and the tracker becomes unstable. In comparison, as expected, the PCA shape-space, being finely tuned to the object’s deformation, tracks well. The interpretational shape-space, described in Section 6.3, decomposes the tracking into clinically meaningful shape-space components.

## 6.6 Clinical plots

As we saw in Section 6.1.1, the traditional measure most used by clinicians to assess wall motion — the area — fails to diagnose a patient with regional ischemia. We saw in Section 6.3 that our interpretational shape-space picked out the reduced excursion in the apical anterior and inferior segments. In Section 6.3.1, using the variation in each segments motion, we went on to show that this abnormality was more apparent in systole than diastole.

In this section, we will compare and contrast two pictorial ways of assessing (regional) endocardial wall excursion, against the interpretational shape-space idea developed in this chapter.

### 6.6.1 Colour kinesis for wall excursion

Lang *et al.* (1996) introduced the idea known as colour kinesis (Mor-Avi *et al.*, 1997) as a way of visualising wall motion, as well as the phase of the wall motion (described below). Godoy *et al.* (1997) and Vitarelli *et al.* (1998) give an up-to-date state-of-the-art review of colour kinesis.

Colour kinesis facilitates the qualitative diagnosis of wall motion abnormalities by allowing the clinician to visualize the extent and synchrony of endocardial contraction and expansion. Colour kinesis keeps track of the estimated motion of the blood-tissue border (endocardium), and colour codes the moving endocardial border to clearly show, on every frame, and in real time, sequential stages of wall motion throughout the cardiac cycle (Lang *et al.*, 1996; Vignon *et al.*, 1998). To do this, colour kinesis uses the endocardial borders obtained from Acoustic Quantification (AQ) (Mor-Avi *et al.*, 1995). Therefore, colour kinesis provides a tool for looking at the magnitude and timing of endocardial motion.

An example is given in Figure 6.13. Here we have plotted the previous 20 contours together. Each contour is colour coded from the most recent, in green, to the oldest, in red. The wall motion in systole is shown in Figure 6.13 (a), and during diastole, in Figure 6.13 (b). The figure demonstrates the reduced motion of the apical inferior and anterior segments and offers the potential for looking at the phase of wall motion. Colour kinesis, as shown in Figure 6.13, is now integrated as a tool in the Hewlett-Packard 5500 SONOS ultrasound machine.

An alternative way of plotting the same frames as in Figure 6.13 is shown in Figure 6.14. This difference here is that each contour has been 'opened out' starting from the mitral valve, going clockwise around the contour — i.e. the basal anterior segment through to the basal inferior segment. Again we have used the contour on the 20 previous frames. For each of these 20 contours, the difference between the contour on the frame  $t$  with that on the frame  $(t - 1)$  is plotted in a colour. These are then coded from the most recent, in blue, to the oldest, in red. Again we can see the reduced motion of the apical inferior and anterior segments.

Koch *et al.* (1996) used a similar approach. In their work, end-systolic colour kinesis images were segmented, and pixel counts were used to calculate incremental regional fractional area change. These were then displayed as stacked histograms. Koch *et al.* (1996) uses the technique in the context of dobutamine stress echocardiography (Cokkinos *et al.*, 1998). The method highlights a reduced segmental area in patients with regional wall abnormalities, following introduction of the pharmacological drug.




Figure 6.13: **Colour kinesis contours.** Colour kinesis facilitates the qualitative diagnosis of wall motion abnormalities. Here we have chosen to plot the previous 20 contours together on frames 25 and 45. The present contour position is coded in blue, with a natural scale in-between. The previous 20 (every other) contours are then coded from the most recent, in green, to the oldest, in red. Wall motion is visualized during endocardial (a) contraction and (b) expansion. Note the reduced wall motion in the apical inferior and anterior segments. This is consistent with the actual clinical diagnosis of a myocardial infarct.




Figure 6.14: **Colour kinesis chart.** Colour kinesis facilitates the qualitative diagnosis of wall motion abnormalities. Here we have chosen to plot the previous 20 contours together on frames 25 and 45. The present contour position is coded in blue. The previous 20 contours are then coded from the most recent, in green, to the oldest, in red, with a natural scale in-between. The kinesis chart begins at the basal anterior segment, and then goes clockwise around the boundary. Wall motion is visualized during endocardial (a) contraction and (b) expansion. Note the reduced wall motion in the apical inferior and anterior segments. This is in line with the patient's myocardial infarct.

Both of these colour kinesis plots, correctly highlight the reduced motion of the apical inferior and anterior segments, although are qualitative. In comparison, the interpretational shape-space components plotted in Figure 6.9, and the scoring scheme shown in Figure 6.11, are quantitative.

### 6.6.2 Velocity coding of wall motion

Colour kinesis provides a way of qualitatively characterising the phase of endocardial wall motion. In this section, we consider coding the instantaneous velocity of the wall motion, rather than the motion of the wall.

To do this, at each time frame, we calculate the velocity of the contour between the present time-frame, and the previous time-frame. We calculate the distance between each point on the present contour to the same point on the previous contour. We then colour code this distance, so that low velocities are coded in blue, and faster velocities in red.

Since we are more interested in highlighting reduced wall motion, we use a non-linear curve to map the contours displacements to assigned colours. This curve was experimentally chosen to be,  $1 - e^{-d/2}$ , where,  $d$ , is the distance (in  $\text{cm}^2$ ) between the same point on the two curves.

Figure 6.15 shows the instantaneous velocity of the wall motion in systole (frames 17 to 25) and diastole (frames 37 to 45) (every other frame is shown). We can see from Figure 6.15 (a) that during systole that the contour is coloured by shades of reddish orange, yellow and green, which indicates normal contractability. However, in each of the 5 frames shown, the apical inferior and anterior segments are represented by blue, colder colours. This highlights these segments reduced motion. We can also see that the cyan/green colours in the basal inferior segment indicates a reduced (but normal) motion. Systole ends at frame 25 — hence the whole contour (on frame 25) is a shade of blue.

In general, in diastole, Figure 6.15 (b), shows that the whole contour is coded by bright red, oranges and yellows, which indicated high velocities. However, the apical inferior and anterior segments are still represented by blue, colder colours. The heart is 'relaxed' at frame 45 — hence the whole contour (on frame 45) is a shade of blue.




Figure 6.15: **Colour velocity.** At each time frame, we calculate the velocity of the contour between the present frame, and the previous frame. This velocity is then mapped (non-linearly) to a colour — low velocities are coded in blue, and faster velocities in red, with a natural scale in-between. (a) 5 (every other) frames in systole. (b) 5 (every other) frames in diastole. Note the ‘colder’ colours in the apical inferior and anterior segments. This is in line with the patient’s myocardial infarct.

## 6.7 Conclusions

This chapter has defined a framework for interpreting a tracked echocardiographic sequence. We have seen that decomposing a tracker based on the interpretational shape-space provides a clinically meaningful quantitative way of assessing regional heart function. The approach also allows the clinician to assess the homogeneity of the ischemic region.

We proposed using this interpretational shape-space to define a normalised quantitative scoring scheme based on maximal endocardial wall motion.

We conducted a comparison of established qualitative ways of assessing wall motion, with our

approach.

The next chapter extends some of the ideas to myocardial thickening quantification. Chapter 8 contains a more rigorous evaluation of the methods in this chapter applied to clinical data.

# CHAPTER 7

---

## Myocardial Wall Thickening

---

The two most frequently used measures to characterize the dynamics of the heart are: endocardial wall motion (excursion) and myocardial wall thickening. We considered tracking wall excursion in Chapter 5. In this chapter, we develop a framework for determining myocardial wall thickening. This is a key contribution of this thesis. To our knowledge, this is the first fully automatic algorithm capable of quantifying wall-thickness. Normal myocardium demonstrates systolic thickening and inward movement of the endocardium. Thus, absence of systolic thickening/endocardial inward motion (akinesis), or systolic thinning/endocardial expansion (dyskinesis), is clearly abnormal when seen during an echocardiographic examination.

As we saw in Chapter 6, one way of detecting myocardial ischemia was by noting the abnormal motion of the ischemic wall segment. In that chapter, we showed that we could quantitatively measure regional endocardial wall motion. We then went on to develop a semi-automatic quantitative scoring scheme.

Our approach to tracking and quantifying myocardial wall thickening will be similar to that used in endocardial wall excursion.

One problem with solely using endocardial motion as a criterion for ischemia is that the movement of any given ventricular segment is directly influenced by the adjacent segments. This can cause two effects. Firstly, that a dyskinetic segment may affect an adjacent normal segment, resulting in the normal segment appearing to be hypokinetic. Conversely, a vigorously contract-

ing normal segment may affect an adjacent ischemic segment, masking any abnormality (Falseti *et al.*, 1981).

As an alternative Kerber *et al.* (1975, 1976) and Torry *et al.* (1991) propose assessing heart function using systolic thickening. This provides a more sensitive predictor of myocardial ischemia than changes in epicardial segment length (Heikkilä *et al.*, 1972; Lieberman *et al.*, 1981). Normal myocardial muscle increases in thickness with systolic contraction. Reduced systolic thickening, or indeed in some cases systolic thinning (the thickness of the left ventricular wall is greater in diastole than systole), is an indication of acute ischemia or infarction. In this chapter, we consider determining myocardial wall thickening.

The myocardium is the muscle tissue that forms the middle layer of the heart, and is the region in between the epicardium and endocardium, shown in Figure 7.1. (Note: In this image the myocardium is clearly visible. This is not often the case). Put another way, the endocardium and epicardium are the inner and outer boundaries of the myocardium. A fuller explanation of the structure of the heart wall is given in Appendix B.1.5. During contraction, as the heart propels blood from the left ventricle, the thickness of the myocardium increases. The position of the epicardium remains fairly unchanged over the cycle — only its relative position to the endocardium alters.

In Section 7.1 we discuss the strategy we have developed to track myocardial thickening. The approach we have taken is to track the endocardium and epicardium and then infer knowledge about myocardial thickening from the relative positions of these boundaries. The tracking of the epicardium turns out to be much harder than tracking the endocardium. We outline the novel method we have developed to track the epicardium in Section 7.2. As with modelling the endocardium, a key ingredient for successfully tracking the epicardium, is having a prior model of its deformation (good system model). We discuss this in Section 7.3. We show that simply modelling the epicardium using a shape-space, is not good enough. We need to constrain the fitted position of the epicardium with reference to the endocardium — Section 7.4.

The second key ingredient in tracking is the measurement model. In Section 7.5, we show that a local-phase ridge-feature detector is not good enough, and we show, in Section 7.6, that localisation based on a wavelet based ridge-feature detector works well. To add robustness to the feature detection, we go on to describe a method for obtaining feature measurements based on assimilating the information in successive profile normals.

In Section 7.7 we look at quantifying wall thickening and introduce a novel segmental scoring scheme for this in Section 7.8. In Section 7.9 we compare using wall excursion and wall thickening to assess regional heart function. In Section 7.10, we extend the ideas of colour kinesis,

introduced in Section 6.6.1, to wall thickening.




Figure 7.1: **The myocardium.** The heart is enclosed by a double-layered *pericardial sac*. The inner layer of this loose-fitting sac is the *serous pericardium*. The *epicardium* is formed by the serous pericardium turning back on itself at the base of the heart. This thin serous membrane forms the thin, outer layer of the heart wall. (a) and (b) Diagrammatic representation of the layers of the heart wall and of the pericardium. (c) Echocardiographic image showing myocardial thickening during systole.

## 7.1 Myocardial tracking strategy

Our goal is to determine myocardial thickening. It is more convenient to track the inner and outer boundaries of this muscle — the endocardium and epicardium — and then infer knowledge about myocardial thickening from the relative positions of these boundaries.

Two potential ways of doing this are as follows:

1. Track the myocardial thickening using a 'ribbon snake' (Zeng *et al.*, 1998). As with traditional snakes (Kass *et al.*, 1987; Irwins & Porrill, 1995; Székely *et al.*, 1996), the optimal contour can be found by minimising an energy function. The energy function consists of internal and external forces, where the internal forces serve as smoothness constraints and the external forces guide the active contour towards image features. These forces then place constraints on the configuration of the contour.

In the multiple active contour model of Chalana *et al.* (1996), an added temporal constraint is placed on the contour. Chalana makes the assumption of monotonic motion of the heart in the diastolic and systolic phases of the cardiac cycle, i.e. in the systole phase, it is monotonically decreasing in size, and in the diastolic phase it is monotonically increasing in size.

A ribbon snake makes a similar assumption, but the multiple active contour model, places a joint constraint on the contours. In our application, the two contours would represent endocardial and epicardial boundaries. For example, a constraint could be placed on the joint shape deformation of the boundaries. Zeng *et al.* (1998) used a similar idea to seg-

ment the cortex. In their work, the cortex was modelled by two bounding surfaces and the homogeneity in between.

2. Alternatively, track the endocardium and epicardium separately, and infer knowledge about myocardial thickening using the distance between the tracked contours. We know that the endocardium and epicardium represent the boundaries of the myocardium. Therefore, this places a physical constraint on their relative position — the distance between the boundaries must be positive. We will return to this below. Our experience, shown below, shows that tracking the boundaries separately, does not satisfy the constraint. However, we can overcome this by modelling the joint shape deformation of the two boundaries using a shape-space.

We have adopted this approach.

The framework we have developed to track the epicardium is similar to that used for the endocardium, discussed in Chapter 5. The difference is that we have not explicitly predicted the position of the epicardial boundary. Instead, we will rely on the estimated position of the endocardial boundary as a starting point in finding the epicardial boundary.

We rely on the two other key ingredients for tracking stated in Section 4.2. Specifically, we developed:

- a model of the shape deformation of the epicardium, and,
- a measurement model to find the epicardial boundary.

Clearly, both of these ingredients are important. As with the endocardium, we model the shape deformation of the epicardium first, and then develop the measurement model. One reason for this is that when we fit a spline to a set of measurements, we place a constraint on the fitted spline so that it lies in the shape-space of the object. Therefore, in this fitting process, having the appropriate model of shape will automatically down-weight poorer observations, which do not lie in the main body of the shape-space.

We consider tracking the epicardium next.

## 7.2 Tracking epicardial boundary motion

Three possible ways of tracking the epicardium are shown in Figure 7.2. We will use the method illustrated in Figure 7.2 (b).

We first track the endocardium. Then, using the knowledge of the position of the endocardium, 'look outwards' to find a set of features that represent the epicardium. The position of the

epicardium is then estimated by fitting a B-spline contour to these measurements.

The advantage of this approach is that the endocardium is often easier to find in 2-dimensional images produced during an echocardiographic examination. The reason for this is that this boundary represents a blood-vessel interface, and ultrasound reflections from this discontinuity are often stronger than that from the epicardium/pericardium. Therefore, tracking is based on the more robust (of the two) set of feature measurements.

A similar approach was used by Chalana *et al.* (1996). In their work, they firstly detected the epicardium for the entire sequence. The epicardium was then used as an initial curve in finding the endocardium. Since the endocardium must lie inside the epicardium, finding the epicardium first provides a constrained enclosed region of interest in the image in which to look for the endocardium. The advantage of Chalana *et al.*'s approach over ours, is that searching for the endocardium is then constrained. However, the problem with this method is that finding the positions of the epicardium is difficult, both for a computer algorithm and a clinician. It would be too unreliable to use this as the starting point. The endocardium, which forms the blood-tissue boundary, is often easier to find.

A potential disadvantage of our approach is that we model the shape of the epicardium, but not its motion. Moreover, although we know that the endocardium and epicardium are related, we do not model their joint motion. We discuss the effectiveness of this strategy next.




Figure 7.2: **Myocardial wall tracking — tracking coupled contours.** The myocardium is the muscle tissue that forms the region in-between the epicardium and endocardium. As the thickness of the myocardium alters during the cardiac cycle, the relative position of the epicardium and endocardium changes. In Chapter 5, we considered tracking endocardial wall excursion. Three possible ways to track the epicardial boundary are: (a) Model the shape, motion and measurements of the epicardium, without using knowledge of the endocardium. (b) Model the shape and measurements of the epicardium. The estimated position of the endocardium is then used as a starting point in finding the endocardium. (c) Jointly model the shape, motion and measurements of the endocardium and epicardium.

### 7.3 Epicardial shape-space estimation

We saw in Section 4.4, that we can construct a PCA shape-space using the control points from a number of manually segmented image frames. We applied this idea in Section 5.1, to obtain

a shape-space that is capable of modelling the shape deformation of the endocardial boundary. An example of the modes of deformations of a shape-space was given in Figures 5.4 and 5.6.

Now consider building a shape-space for the deformation of the epicardium for the dataset of Figure 5.6. This was done in precisely the same way as for the endocardium, by manual segmentation of the boundary and subsequent PCA on the control points. Table 7.1 summarises the results of the PCA analysis. Figure 7.3 shows a scree plot of the magnitude of the ordered eigenvalues against the dimension. We chose to use six modes, which explained 97.5% of the variation, based on the elbow of the scree diagram, shown in Figure 7.3. As we saw in Section 5.1.1, we can visualize the modes of deformation of the epicardium, as shown in Figure 7.4. As with the endocardium (see Table 5.3, page 62), the PCA residuals shown in Table 7.2 show that the deformation of the epicardium is not affine.

Although the shape-space created in this section looks appealing, it has shortcomings. We turn to this next.

Mode	Eigenvalue	Variability %	Cumulative %
1	850.67	0.453	0.453
2	519.032	0.277	0.730
3	351.45	0.187	0.917
4	66.47	0.0354	0.953
5	22.35	0.0119	0.965
6	19.07	0.0102	0.975

Table 7.1: **PCA analysis of the deformation of the epicardium.** The result of applying PCA to the epicardial boundary. 4 modes of variation explain over 95% of the variability.




Figure 7.3: **Scree diagram.** A scree diagram showing a plot of the magnitude of the ordered eigenvalues against dimension for the data in Table 7.1. The number of dimensions is typically chosen as the foot of the initial steep decline — or the elbow of the scree diagram.




Figure 7.4: **Modes of deformation of the epicardium.** A principal component analysis was performed on the epicardial boundary using a B-spline with 14 control points. The results are shown in Table 7.1 and Figure 7.3.

Eigenvector	$\frac{\ \mathbf{w}_i - \mathbf{W}_A \mathbf{W}_A^\dagger \mathbf{w}_i\ _2}{\ \mathbf{w}_i\ _2}$	$\left( \frac{\ \mathbf{w}_i - \mathbf{W}_A \mathbf{W}_A^\dagger \mathbf{w}_i\ _2}{\ \mathbf{w}_i\ _2} \right)^2$
$\mathbf{w}_1$	0.4766	0.2271
$\mathbf{w}_2$	0.2869	0.08232
$\mathbf{w}_3$	0.9016	0.8130
$\mathbf{w}_4$	0.7747	0.6002
$\mathbf{w}_5$	0.8708	0.7584
$\mathbf{w}_6$	0.8153	0.6648

Table 7.2: **PCA Residuals.** Projecting a PCA shape-space,  $\mathbf{W}$ , into and out of an affine space. Equation 5.2, as  $r = \frac{\|\mathbf{w}_i - \mathbf{W}_A \mathbf{W}_A^\dagger \mathbf{w}_i\|_2}{\|\mathbf{w}_i\|_2}$ , is shown. Modes 3–6 are not affine. The data was digitally captured and stored on optical disk. 24 control points are used.

## 7.4 Fitting spline templates and choice of epicardium shape-space

As we have said, to find the epicardium, we could look outwards from the estimated position of the endocardium. Image features are then identified through image processing operations on the search line data (contour normals).

As with any measurements taken from an image, these measurements will be 'noisy' making fitting problematic. One way to overcome this is to use regularised fitting. In general, regularisation tries to avoid over-fitting to a training set, by penalizing the fit with a measure of 'smoothness' of the fitted function. In the context of fitting a contour to these measurements, a regulariser can be used to impose prior constraints on the likely shape of the contour. The regulariser we use is based on the shape-space of the object. This is precisely what we did in Chapter 5 when we updated the predicted position of the endocardium contour, with image measurements.

The solution to this curve fitting problem is given in Figure 7.5, and is taken from Blake & Isard (1998, page 133). For example, the regulariser,  $\alpha\|\mathbf{X} - \bar{\mathbf{X}}\|^2$ , is obtained by setting  $\bar{\mathbf{S}} = \alpha\mathcal{H}$ . In this case,  $\alpha$  controls the trade off between bias and smoothness. A large value of  $\alpha$ , gives greater tolerance to noisy data, but biases towards the mean shape. A low value of  $\alpha$ , gives greater fidelity to the data, but greater sensitivity to noisy data. In fitting a contour to measurements made of the epicardium, we empirically chose a low value of  $\alpha = 0.03$ .

Figure 7.6 shows typical frames from an ultrasound sequence. The tracked position of the endocardium is shown in red, and is based on the tracker framework of Chapters 4 and 5. The epicardium is then estimated using image features along the search lines, normal to the endocardium contour. A contour is then fit to these measurement ( $\alpha = 0.03$ ). The estimated position of the epicardium is shown in yellow.

Although, tracking the epicardium in this way appears appealing, Figure 7.6 begins to reveal a problem with the method. The approach we take for estimating the epicardium, guarantees that the measurements that represent the epicardium will lie outside the measurements that represent the endocardium. However, despite this, there is no guarantee that the contour we fit to estimate the epicardium, will lie outside the contour previously fit to represent the endocardium.

This is because, the epicardium shape-space does not contain any knowledge about the endocardium. This problem is illustrated in the top pair of images shown in Figure 7.8. We can see that because there is no constraint on the relative distance between the endocardium and epicardium, the contour that represents the epicardium can cross-over the contour that represents the endocardium. In reality this distance should be positive.

An alternative solution is to model the distance between the epicardium and endocardium. This provides a relationship between the two boundaries of the myocardium. In the probabilistic framework of shape-spaces already developed, this translates into using a shape-space to model the shape-deformation of the difference between the epicardium and endocardium.

Suppose that the manually segmented contours that represent the endocardium and epicardium are given by,  $\mathbf{Q}_{En,1}, \mathbf{Q}_{En,2}, \dots, \mathbf{Q}_{En,M}$ , and  $\mathbf{Q}_{Ep,1}, \mathbf{Q}_{Ep,2}, \dots, \mathbf{Q}_{Ep,M}$ , respectively. Then, as in Section 5.1, we can construct a PCA shape-space using the control points of the contours, except now, the ‘difference shape-space’,  $\mathbf{W}_{Diff}$ , is based on the PCA of the difference,  $\mathbf{Q}_{Diff,i}$ , in control points,  $\mathbf{Q}_{Diff,i} = \mathbf{Q}_{Ep,i} - \mathbf{Q}_{En,i}$ . The shape-space equivalent of  $\mathbf{W}_{Diff}$  is  $\mathbf{X}_{Diff}$ . The advantage of using  $\mathbf{W}_{Diff}$  is that now there is a natural constraint on the distance between the estimated position of the epicardium and endocardium, which is embedded into the shape-space.

**Curve fitting problem**

Given an initial shape estimate  $\bar{\mathbf{r}}(s)$  (or  $\bar{\mathbf{X}}$  in shape-space) with normals  $\bar{\mathbf{n}}(s)$ , and a regularisation weight matrix  $\bar{\mathbf{S}}$ , solve:

$$\min_{\mathbf{X}} T \quad \text{where} \quad T = (\mathbf{X} - \bar{\mathbf{X}})^\top \bar{\mathbf{S}} (\mathbf{X} - \bar{\mathbf{X}}) + \sum_{i=1}^N \frac{1}{\sigma_i^2} (\nu_i - \mathbf{h}(s_i)^\top [\mathbf{X} - \bar{\mathbf{X}}])^2.$$

**Algorithm**

1. Choose samples  $s_i, i = 1, \dots, N$ , s.t.  $s_1 = 0, s_{i+1} = s_i + h, s_N = L$ .
2. For each  $i$ , apply some image-processing filter along a suitable line (e.g. curve normal) passing through  $\bar{\mathbf{r}}(s_i)$ , to establish the position of  $\mathbf{r}_f(s_i)$ .
3. Initialise
 
$$\mathbf{Z}_0 = \mathbf{0}, \quad \mathbf{S}_0 = \mathbf{0}.$$
4. Iterate, for  $i = 1, \dots, N$ :
 
$$\begin{aligned} \nu_i &= (\mathbf{r}_f(s_i) - \bar{\mathbf{r}}(s_i)) \cdot \bar{\mathbf{n}}(s_i), \\ \mathbf{h}(s_i)^\top &= \bar{\mathbf{n}}(s_i)^\top \mathbf{U}(s_i) \mathbf{W}, \\ \mathbf{S}_i &= \mathbf{S}_{i-1} + \frac{1}{\sigma_i^2} \mathbf{h}(s_i) \mathbf{h}(s_i)^\top \text{ and,} \\ \mathbf{Z}_i &= \mathbf{Z}_{i-1} + \frac{1}{\sigma_i^2} \mathbf{h}(s_i) \nu_i. \end{aligned}$$
5. The aggregated observation vector is,
 
$$\mathbf{Z} = \mathbf{Z}_N,$$
 with associated statistical information,
 
$$\mathbf{S} = \mathbf{S}_N.$$
6. Finally, the best fitting curve is given in shape-space by:
 
$$\hat{\mathbf{X}} = \bar{\mathbf{X}} + (\bar{\mathbf{S}} + \mathbf{S})^{-1} \mathbf{Z}.$$

Figure 7.5: **Fitting splines to data.** Recursive algorithm for curve fitting. Taken from Blake & Isard (1998, page 127). For example, the regulariser,  $\alpha \|\mathbf{X} - \bar{\mathbf{X}}\|^2$ , is obtained by setting  $\bar{\mathbf{S}} = \alpha \mathbf{H}$ . In this case,  $\alpha$  controls the trade off between bias and smoothness.

Measurements of the position of the epicardium are then obtained, relative to the endocardium. Using the difference shape-space,  $\mathbf{W}_{\text{Diff}}$ , a contour can be fit to these measurements. The resulting contour,  $\mathbf{Q}_{\text{Diff}}$ , can be thought of as the extra contribution that is added to the configuration of the endocardium contour, to estimate the configuration of the epicardium contour. The full algorithm for estimating the epicardium is described in Figure 7.7.

Figure 7.8 compares tracking results based on modelling the shape deformation of the epicardium (top) with modelling the shape deformation of the difference between the boundaries. We can see that simply modelling the shape deformation of the epicardium, without knowledge




Figure 7.6: **Epicardium estimation based on a shape-space for the deformation of the epicardium.** The endocardium (red) is tracked based on the framework discussed in Chapters 4 and 5. The epicardium is then estimated using image features along the search lines, normal to the endocardium contour. A contour is then fitted to these measurement, according to Figure 7.5.

### Epicardium estimation

Given an estimate of the endocardium  $Q$ , ( $X$  in shape-space), with normals,  $\bar{n}(s)$ , estimate the position of the epicardium.

#### Algorithm

1. Obtain a shape-space,  $W_{\text{Diff}}$ , of the difference between the manually segmented endocardium and epicardium contours,  $Q_{\text{En},1}, Q_{\text{En},2}, \dots, Q_{\text{En},M}$ , and  $Q_{\text{Ep},1}, Q_{\text{Ep},2}, \dots, Q_{\text{Ep},M}$ , respectively.
2. Search normally to the estimated position of the endocardium, to find image measurements that represent the epicardium.
3. Using these measurements, obtain a best fitting curve,  $\hat{X}_{\text{Diff}}$  (using the algorithm in Figure 7.5), in the difference shape-space,  $W_{\text{Diff}}$ . Call this contour  $\hat{Q}_{\text{Diff}}$ .
4. The estimated epicardium position,  $\hat{Q}_{\text{Ep}}$ , is then given by  $\hat{Q}_{\text{Ep}} = \hat{Q}_{\text{En}} + \hat{Q}_{\text{Diff}}$

Figure 7.7: **Algorithm for epicardium estimation.**

of the endocardium, results in the epicardium contour crossing over the endocardium contour. By comparison, modelling the shape deformation of the difference between these boundaries imposes a constraint on the distance between the boundaries, and there is a gap between the contours.



Next, we shall briefly consider the choice of search region. This is the region in which we expect the epicardium (and hence image features) to lie in.




Figure 7.8: **Choice of shape-space for epicardium estimation.** Image measurements are performed by casting rays along normals, to the estimated position of the endocardium, and then measuring the positions, of the epicardium along the ray. The position of the epicardium is then estimated by fitting a contour to these measurements. These measurements, as any other will be noisy. A regulariser is used to impose prior constraints on the likely shape of the contour.

For each image frame, the top image represents a shape-space based on the epicardium. The bottom image represents a shape-space based on the difference between the epicardium and endocardium boundaries. Simply modelling the shape deformation of the epicardium does not embed any knowledge about the relative distance between the endocardium and epicardium — this distance must be positive. This is shown by the epicardium contour crossing over the endocardium contour.

In comparison, modelling the shape deformation of the difference between these boundaries imposes a constraint on the distance between the boundaries, and there is a gap between the contours.

### 7.4.1 Choice of search scale

Since we do not use a Kalman filter for the estimation of the epicardium, there is no state covariance upon which to base a search scale.

Initially we tried using the search scale for the endocardium,  $\rho(s, t)$  for the epicardium. Specifically, the length of the  $i$ 'th epicardium search scale, was arbitrarily taken as  $2\rho(s, t)$ . This is shown in Figure 7.9.

The problem with this, as Figure 7.11 (a) illustrates, is that then the search scale is essentially based on our confidence in the predicted position of the endocardium. This means that the search scale is too small along parts of the contour, and tracking is lost. We found that a better choice was to use an empirically chosen constant search scale of 30 pixels (1.14cm in this example) as shown in Figure 7.10; tracking results are shown in Figure 7.11 (b). In fact, it would be better to link the search scale for the epicardium to the stage of the cardiac cycle, so that the search scale would be longer in systole than in diastole. This was not pursued in this thesis because the ECG information is not easily obtainable.

Having obtained a model of the deformation of the epicardium and choice of search region, we now consider a measurement model to localise the features of the epicardium. Our initial approach was to localise these features using local phase, as outlined in Section 5.3.1. We turn to this next.




Figure 7.9: **Epicardium search scale based on endocardium search scale.** To estimate the position of the epicardium, normals are cast from the estimated position of the endocardium. (a) Measurements are taken along a normal to the predicted position of the endocardium. The search scale for the  $i$ 'th endocardium normal is  $\rho(s, t)$ . These measurements are used to estimate (and update) the position of the endocardium. Using the estimate of the endocardium, we search normally to this boundary to find the epicardium. The search scale for the  $i$ 'th epicardium normal is arbitrarily taken as  $2\rho(s, t)$ . (b) Image measurement are taken along normals (green) to the predicted position of the endocardium (red). This updated position then forms the starting point in estimating the epicardium. Image measurement are taken along normals (blue) to the estimated position of the endocardium.




Figure 7.10: **Measurement process for epicardium estimation.** (a) Normals are cast from the estimated position of the endocardium. The length of the normals is chosen to be constant (see Section 7.4.1). Measurements are taken along this normal to obtain the position of the epicardium. The estimate of the epicardium is obtained by fitting a contour to these measurements. (b) Measurements are taken at sampled positions of the left ventricular boundary contour. Image features are identified through image processing operations on the search line data.




Figure 7.11: **Epicardium search scale. An example.** (a) The search scale for the epicardium is based on the search scale for the endocardium. The epicardium contour lies inside the visually observed position of the epicardium. The reason for this is that the search scale is too small along parts of the contour, and the epicardium does not lie in the search line. More importantly, the search scale is based on our confidence in the predicted position of the endocardium. (b) Using a constant search scale of 30 pixels.

## 7.5 Measurement model to localise epicardium features

In this section, attention is turned to the feature measurement process. A ridge detector is developed for the detection of the epicardium. The initial approach was to find the epicardium/pericardium boundary using a local phase step detector, as per the endocardium.

This was unsatisfactory as Figure 7.12 shows. The reason for this is explained in Figure 7.13, where one-dimensional intensity profiles are shown starting from the estimated endocardium (based on tracking). The epicardium boundary is observed to be an intensity ridge.

A local phase ridge detector was then tried, but this was unsatisfactory because the detector was too sensitive to noise. As an alternative, a ridge feature detector was investigated using wavelets.

This is considered next.




Figure 7.12: **Epicardium estimation using a local phase step detection.** Typical frames from an ultrasound sequence. The position of the endocardium is tracked (red). Search lines are cast, normal to this boundary, and a local phase step detector was applied. The largest feature response along the search line is shown (cyan). We can see that this feature detector is not good enough. Note: The measurements along the mitral valve region are fixed to be the same for the endocardium and epicardium estimation.




Figure 7.13: **Intensity profiles along the myocardium.** Intensity profiles (60 pixels) from the tracked endocardium, through the myocardium, past the epicardium. The epicardium boundary is observed to be an intensity ridge. (a) Tracker interface. (b) Associated intensity profiles.



## 7.6 Localisation based on a wavelet ridge detector

This section outlines how feature measurements are localised along normals using an approach based on a wavelet ridge detector. We go on to discuss a novel way of estimating measurements on the  $i$ 'th normal by weighting the wavelet information on the  $i$ 'th normal with wavelet information from the previous  $(i - 1)$  normals. The wavelet approach detailed in this section was implemented by Christian Behrenbruch. Further details of the approach are described in Behrenbruch *et al.* (1999). We now turn to the choice of wavelet.

We saw in Figure 7.13 (b) that one-dimensional profiles across the epicardium boundary can be viewed as an intensity ridge. We assume that the maximum of this intensity profile represents the epicardial feature of interest. However, noise and artifacts in an ultrasound image make this feature detection challenging.

Our problem is to find a wavelet approximation (Daubechies, 1988) to the intensity profile, which ignores the noise without smoothing regions of interest and losing accuracy. This is made harder by the feature being a ridge, and not an edge. Fundamental to the use of wavelets is the application of scale-space and multi-resolution decomposition. Using wavelets we can reconstruct a profile using resolutions optimised for specific signal components (Chambolle *et al.*, 1998).

### 7.6.1 Choice of wavelet

In this section, we will consider two issues involved when choosing a wavelet, namely, the wavelet type and computation cost.

The choice of wavelet for this application is the Coifman wavelet series (Daubechies, 1994) (Figure 7.14). These orthogonal, compactly supported wavelets feature the highest number of vanishing moments for a given filter length. In addition, several packets within the Coifman set have a spatial distribution that is particularly suitable for evaluating ridge-like structures (such as the epicardium) at low resolutions. Each of the profiles (such as those shown in Figure 7.13) are initially dyadically-decomposed, and then reconstructed using a best-basis algorithm (Coifman & Wickerhauser, 1992). Each packet within this optimised set is then compared with the total best-basis reconstruction to determine which decompositions contribute most strongly to the original profile signal. These particular low-resolution decompositions are used in the next stage of the analysis to restrict the reconstruction of the filtered profile to that of the ridge-characteristics of successive profiles. The first profile normal is chosen to be at the start of the basal anterior segment, with the profile normals incremented anti-clockwise from this point, with the last being at the end of the basal inferior segment.

Figure 7.15 shows features detected by the three wavelets with different filter lengths (subscripted). For comparison, we also show results from the Daubechies and Beylkin wavelets. The figure indicates that the Coifman wavelet with a filter length of 30 (C30) appears to be best. However, computationally, it is slow to compute. We, therefore, choose the Coifman wavelet with a filter length of 6, i.e. C6 in the Figure.




Figure 7.14: **A Coifman wavelet packet.** A Coifman wavelet packet. Note how the ridge shape of the wavelet packet matches the ridge feature we seek.

## 7.6.2 Assimilating information using wavelets

As we saw in the previous section, the Coifman wavelet, is able to localise features along the epicardium. In this section, we develop an approach to assimilate information about the features. We do this by spatial fusion of the wavelet packets. To do this, as we step through each of the search lines along the boundary, we aggregate information about the successive profiles.

Although the epicardium is usually quite visible across an ultrasound sequence, automated detection is made difficult by the presence of noise and artifacts that effectively spread the range of possible ridge points. Simple smoothing (even using a morphologically effective tool such as a wavelet) does not fully enable accurate detection of the ridge, as shown in Figure 7.15, as it only smooths in 1-dimension. The adjustment of the reconstruction is developed iteratively across the image space, i.e. across the normal profiles of a single frame. The effective result of this adjustment, is that peak or ridge-like features are negatively weighted if (1) they deviate outside of the estimated image space for the myocardium (given by the tracker software), or, (2) if local maxima deviate too greatly from an allowable curvature of the boundary. This curvature is important because it can change the profile geometry, and hence shape of the features.

Once the decomposed profile has been adjusted to reflect ridge localisation in the previous profiles, the low-resolution representation is then adjointly convolved with the low and high-pass quadrature filters for reconstruction. The reconstructed profile has now been smoothed (due to the use of a limited number of packets for reconstruction), shifted and normalised (due to the low-resolution weighting function); then a ridge detection that finds the maximum of the reconstructed profile is applied to this. The net result is that the ridge detection is much more consistent with the visual location of the epicardium, as Figure 7.16.

Localisation of the epicardium features is improved for all three wavelets. The figure also shows (frame 36) that the curvature information is important. In the basal anterior segment, localisation is poorer than on other parts of the epicardial boundary. The reason for this is that the position of the first profile normal is around the mitral valve. The profile normals are incremented anti-clockwise from this point. The change in curvature of the boundary along this part of the boundary on this frame deviates too much from the allowable curvature. However, we can see that the assimilated wavelet information is quickly updated, and localisation improved. An example of reconstructed profiles are shown in Figure 7.17. Some typical image features detected using the Coifman wavelet are given in Figure 7.18. Figure 7.19 shows a plot of the estimated myocardium thickness at various stages in the cycle, based on this.

Having obtained models of the shape and measurements of the epicardial boundary, we are now able to use the tracked myocardium to quantify wall thickening. We turn to this next.

A more elegant solution to this problem, which gives the same answer is to use shape-spaces. This is discussed in Appendix E.3. Figure 7.20 shows a plot of the myocardial thickening for the data in Figure 7.19 (two cycles). We can see that the basal inferior, mid anterior, mid inferior and basal inferior segments all move normally. The observed smaller thickening in the basal anterior is in accordance with normal heart function.

However, we can see the reduced change in wall thickness in the apical anterior and apical inferior segments. This observation is exactly in accordance with the diagnosis that this patient has a myocardial infarct in the apical region.




Figure 7.20: **Myocardial wall thickening.** Having tracked the epicardium and endocardium boundaries, we calculate regional wall thickening using the 16-segment model introduced Section 6.2. This ability to quantitatively analyse regional wall thickening is clinically very useful. The data is as in Figure 7.19.

### 7.7.1 Segment variation in wall thickening

It is also useful to look at the variation in segmental wall thickening in each of the six segments, as shown in Figure 7.21.

The starting point for this dataset is diastole, with the heart entering systole after around a tenth of a second. The basal anterior and basal inferior segments show a lot of variation in wall thickness. The lack of thickening and variation in the apical anterior segment is consistent with the actual clinical diagnosis of a myocardial infarct in the apical region. We can see that there is more variation in the apical inferior than the apical anterior segment. This implies that not all of the apical inferior segment is ischemic.

## 7.8 Segmental scoring of wall thickening

The scoring scheme we used is based on regional percentage of wall thickening (called fractional thickening), %Th, defined as (Torry *et al.*, 1991):

$$\%Th = \frac{Th_{ES} - Th_{ED}}{Th_{ED}} \times 100 \quad (7.1)$$

$$= \left( \frac{Th_{ES}}{Th_{ED}} - 1 \right) \times 100. \quad (7.2)$$

Where,  $Th_{ES}$  is the thickness of the myocardial segment at end-systole, and  $Th_{ED}$  is the thickness of the myocardial segment at end-diastole (in cm).

If the myocardium is 'normal', then  $Th_{ES}$  will be greater than  $Th_{ED}$ , and hence, %Th will be positive. For example, if %Th = 100, it implies that myocardial thickness was double in systole than diastole. If %Th is negative, then systolic wall thinning is occurring. Goldberg (1984) notes that wall thickness at end systole has a value approximately 100% greater than the mean end diastolic value with a range 60–120% (Goldberg *et al.*, 1980). Chapter 8 contains a more rigorous evaluation of the methods in this chapter applied to clinical data.

Regional percentage of wall thickening scores for the data in Figure 6.11 are shown in Figure 7.22 (a). For reference, the maximal endocardial wall excursion scores, developed in Section 6.4.1, for the same data set, are shown in Figure 7.22 (b). The regional percentage of wall thickening scores, confirm the reduced motion in the apical interior and anterior segments. The scores also highlight the reduced, although totally normal, motion of the basal inferior segment.

quantifying myocardial thickening is more important. A similar conclusion was reached by Lieberman *et al.* (1981).

## 7.10 Colour kinesis for wall thickening

Recall that colour kinesis keeps track of the estimated blood tissue border and their transitions, and then colour codes the moving endocardial border to clearly show, on every frame, and in real time, sequential stages of wall motion throughout the cardiac cycle. An example was given in Figure 6.13. As we saw there, an equivalent way to show this information was to 'open out' the contour starting from the mitral valve, going clockwise around the contour, as shown in Figure 6.14.

Clearly a similar idea can be applied to wall thickening as follows.

For every point along the myocardium, calculate the change in the thickness between frame  $t$  with that on the frame  $(t - 1)$ . We then plot these cumulative distances, which are colour-coded so that the most recent is blue, to the oldest, which is red. An example is given in Figure 7.23. Here we have chosen to use the previous 20 frames.




Figure 7.23: **Colour kinesis chart for myocardial thickening.** Colour kinesis facilitates the qualitative diagnosis of wall motion abnormalities. We have extended the idea to wall thickening. For every point along the myocardium, the change in the thickening between frames  $t$  and  $(t - 1)$  is calculated. The cumulative distances are then plotted. Here we have chosen to plot the previous 20 contours together on frames 45 and 105. The present contour position is coded in blue. The previous 20 contours are then coded from the most recent, in green, to the oldest, in red, with a natural scale in-between. The kinesis chart begins at the basal anterior segment, and then goes clockwise around the boundary. Wall thickening is visualized (a) during contraction (b) end of contraction (c) during expansion and (d) end of expansion. Note the reduced wall motion in the apical inferior and anterior segments (3rd and 4th segments in kinesis chart). This is in line with the patient's myocardial infarct.

## 7.11 Conclusions

This chapter has provided a framework for determining myocardial thickening. The solution we adopted was to track the endocardium and then use this as a starting point in finding the epicardium. It was shown that modelling the shape-deformation of the epicardium using a shape-space did not constrain the position of the epicardium. Instead, we modelled the difference between the epicardium and endocardium.

To find the epicardium, we proposed using a wavelet-based ridge detector. We went on to show that this could be improved by assimilating wavelet information along the contour.

Using the tracked myocardium, we proposed a quantitative scoring scheme based on the change

in myocardial thickness during the cycle.

We also extended colour kinesis for myocardial wall thickening.

In the next chapter we apply the ideas developed so far in this thesis to dobutamine stress echocardiography data.

## CHAPTER 8

---

### Stress Echocardiography

---

In Chapters 4 to 7 a set of algorithms were presented that showed that we could track and quantify regional endocardial wall excursion and myocardial wall thickening.

In this chapter, the performance of these algorithms is preliminary assessed in a retrospective clinical test using stress echocardiography. A clinical review of stress echocardiography is provided by Marwick (1996) and is also reviewed in Appendix G.

Stress echocardiography provides both diagnostic and prognostic information for patients with known or suspected coronary artery disease. It comes in many guises including post-treadmill imaging, imaging at the time of exercise with supine or upright bicycle ergometry as well as various forms of pharmacological and pacing stress. The clinical study detailed in this chapter uses a pharmacologic drug, dobutamine, which mimics the process of exercise. This technique is called dobutamine stress echocardiography.

On each patient, the general idea is to visually (and hence qualitatively) compare regional wall motion and wall thickening before, during, and after, any of the stress protocols. By doing this, it is hoped that the clinician will be able to distinguish between the normal hyperdynamic response of the left ventricle to stress, the presence of prior infarction or the development of myocardial ischemia.

It would seem sensible that given the patient is subjected to a known, precise, and in many cases quantitative stress level, that the clinician would likewise use quantitative methods to assess the

change in assessment between the levels of stress. Unfortunately, this is not the case. The clinician's response is qualitative, and simply based on a visual comparison. This is because, until now, automatic quantification of regional left ventricular function has not been possible,

This anomaly affords us with an ideal opportunity to apply the ideas developed so far in this thesis, since we are able to quantitatively assess wall motion and wall thickening. From a technical perspective, this application is interesting as it allows us to focus on assessing intra-patient rather than inter-patient variability.

Before the details of the clinical tests are described and subsequent analysis, some background to stress echocardiography is provided in Section 8.1. The interested reader is directed to a fuller account of stress echocardiography in Appendix G.

The John Radcliffe Hospital dobutamine stress echocardiography protocol is detailed in Section 8.2 and our analysis is presented in Section 8.3. In Section 8.4 we validate the results from a segmentation algorithm. We propose a number of measures for a regional comparison of the algorithm against ground truth, which is taken as a clinician's segmentation. Section 8.5 summarises the chapter conclusions.

## 8.1 Background to stress echocardiography

Under resting conditions, patients with coronary artery disease may have normal left ventricular function. If, at the time of the examination, no permanent myocardial damage has occurred, and if the ventricle is not ischemic, the study will not reflect any underlying coronary artery disease.

There has been considerable interest in combining standard echocardiography with stress intervention to induce ischemia (Limacher *et al.*, 1983; Marwick *et al.*, 1992; Aldrich & Reichek, 1993). This is because (exercise) stress increases cardiac oxygen consumption, and whereas normal myocardial muscle increases in thickness with systolic contraction, ischemic myocardial muscle shows reduced systolic thickening (Pozen *et al.*, 1981).

Early studies were based on exercise electrocardiograms (ECG), but the usefulness of this method in terms of sensitivity and specificity is questionable (Epstein, 1976). Currently, there are a number of methods available to the clinician to induce ischemia — Table 8.1 summarises these. The most common methods are treadmill exercise and pharmacological stress.

With treadmill exercise (Armstrong, 1997), the patient is placed in the supine position, and four standard views are imaged. These views create a baseline for reference. Then the patient exercises on the treadmill. The speed and inclination of the treadmill is increased according to a defined protocol (Bruce & Hornstein, 1969; McHenry *et al.*, 1972), so that in practice the

patient remains on the treadmill for around 8–10 minutes. When the patient indicates that they can no longer exercise or a target heart rate is reached, they are immediately returned to the examination couch and re-imaged post exercise using as close to as possible the same four standard views.

Although treadmill exercise may be considered the most vigorous form of stress, there are some key limitations; for example, many patients are not suitable for, or unable, to exercise and imaging only takes place after exercise. Other limitations are listed in Appendix G.1.1.

Many of the limitations can be overcome by controlling the increases in oxygen demand with a pharmacologic drug agent such as dobutamine (Ryan, 1991; Secknus & Marwick, 1997). Dobutamine is intravenously injected into a patient and increases oxygen demand, and hence the heart is required to work harder. The technique was first described by Tuttle & Mills (1975). It was first reported as a cardiac stress agent by Mason *et al.* (1984). A thorough review of all dobutamine stress echocardiography studies is given by Geleijnse *et al.* (1997).

During a dobutamine infusion, both the patient's heart rate and blood pressure increases, mimicking exercise (Carlson *et al.*, 1988). This increase in heart function is observed echocardiographically by all segments thickening and moving inward to a greater extent than in the corresponding rest image. However, in the presence of significant coronary artery disease, a wall motion abnormality develops in the region of the myocardium supplied by the critically stenosed vessel (Pozen *et al.*, 1981).

Dobutamine stress echocardiography is used by the cardiologists at the John Radcliffe Hospital, Oxford. Their protocol is detailed next.

## **8.2 The John Radcliffe Hospital dobutamine stress echocardiography protocol**

In this section we detail the standard infusion and imaging protocols employed at the John Radcliffe Hospital. This is the standard protocol used by many hospitals (Krahwinkel *et al.*, 1997; Rallidis *et al.*, 1997; Cokkinos *et al.*, 1998).

### **Drug protocol**

The standard protocol for dobutamine stress echocardiography involves injecting a patient with increasing doses of 5, 10, 20, 30 and 40Mg/Kg/min of dobutamine. Each stage may last for as much as 5 minutes, allowing time for the heart rate to stabilise at a 'within dose' level. At peak dose, if the patient has not reached the target heart rate — the chronotropic response is inad-

---

Exercise
Post Treadmill
Supine bicycle
Upright bicycle
Pharmacologic stress
Adrenergic stimulation
Dobutamine
Isuprel
Arbutamine
Epinephrine
Vasodilation
Dypridamole
Adenosine
Pacing
Esophageal
Atrial
Ventricular
Handgrip
Coldpress

---

Table 8.1: **Stress Echocardiography.** Methods available to induce ischemia. Based on Feigenbaum (1994).

equate — 0.5Mg of atropine (i.e glaucoma) may be given every minute up to 2Mg to achieve this (Rallidis *et al.*, 1997; Owen *et al.*, 1998). If patients are unable to receive atropine, a dobutamine dose of 50Mg/Kg/min may be given.

### Imaging Protocol

The patient is imaged while lying on the examination couch in a left decubitus position, as shown in Figure 8.1. ECG (12 lead) recordings are made at each drug dose and in recovery. The patient's heart rate and blood pressure is also recorded at each dose. During the examination, ventricular wall motion is visualised in the standard parasternal long and short axis, and the apical two and four chamber views.

Images <sup>1</sup> are recorded onto an optical disk at baseline and peak infusion. The images acquired during dobutamine stress echocardiography are of poorer quality than during standard echocardiography. This is because logistically the room was too small for the patient to be placed on their side, whilst at the same time monitoring the patient's heart rate, blood pressure and ECG monitored. Therefore, the stress echocardiographic images we shall analyse in this chapter are significantly worse than standard echocardiographic images. Further, the majority of the analysis uses the apical 2-chamber view, which tends to be the hardest to obtain clear images from. We now turn to the analysis of the dobutamine stress echocardiographic images.

<sup>1</sup>The number of cycles capable of being stored digitally by the ultrasound machine is dependent on the patient's heart rate, i.e. the faster the heart rate, the more cycles that can be recorded.




Figure 8.1: **Left decubitus position.** The patient is imaged in the left decubitus position. Figure courtesy of Hewlett Packard.

### 8.3 Stress echocardiography analysis

At the John Radcliffe Hospital, dobutamine stress echocardiography is evaluated by two clinicians. Wall function is assessed at varying levels of dobutamine, starting from a baseline level of no dosage. To assess wall abnormalities, the left ventricle is divided into the 16-segment model of the heart (Figure 6.4). This model is presently used by a number of groups (Krahwinkel *et al.*, 1997; Marwick, 1997; Rallidis *et al.*, 1997; Cokkinos *et al.*, 1998).

Regional myocardial contractile function is graded using a standard scoring scheme (Table 8.2). A normal segment is given the value 1. Abnormal segments have higher numbers. Occasionally, 5 is awarded to an aneurysmal region. The clinician's score is based on the myocardial thickening of each segment. If however, the myocardium thickening is not clearly visible, the clinician's score is based on the observed endocardial excursion of each segment. Appendix H shows the actual scoring sheet used at the John Radcliffe Hospital.

There are a number of outcomes to a dobutamine stress echocardiography procedure as outlined in Table 8.3.

**Normal test result:** A normal stress echocardiogram is defined by a uniform increase in wall motion and systolic wall thickening, with a reduction in end-systolic cavity area.

**Positive test result:** A positive test result is generally defined as the appearance of a new or worsening wall motion abnormality (Sawada *et al.*, 1991; Afridi *et al.*, 1994; Cokkinos *et al.*, 1998).

**A biphasic response:** In patients with rest wall abnormalities, initial improvement of wall motion and thickening at low dose followed by worsening at higher dose.

**Other results:** Tardokinesia (delayed excursions) and relative failure to augment wall thickening.

Score	Grading	Characterized by
1	Normal	A uniform increase in wall excursion and thickening
2	Hypokinetic	A reduced (< 5 mm) inward systolic wall motion
3	Akinetic	An absence of inward motion and thickening
4	Dyskinetic	Systolic thinning and outward systolic wall motion

Table 8.2: **Echocardiography scoring system.** The scoring system used in echocardiography. A normal segment is given the value 1 with abnormal segments having higher numbers. The clinician's score is based on the myocardial thickening of each segment. If however, the myocardium thickening is not clearly visible, but the endocardial border is, then the clinician's score is based on the endocardial excursion of each segment. This figure was as per Figure 6.2.

Test outcome	Wall motion		Wall thickening	
	Low Dose	High Dose	Low Dose	High Dose
Normal	increase	increase	increase	increase
Positive	increase/decrease	increase/decrease	increase/decrease	increase/decrease
Biphasic	increase	decrease	increase	decrease

Table 8.3: **Dobutamine stress echocardiography outcomes.** The main outcomes of a dobutamine stress echocardiography test.

Progress in research into a quantitative even (partially automated) scoring system for dobutamine stress echocardiography has been slow. Little mention is made in the literature about quantitative analysis. Even if it is, the authors have tended to dismiss the possibility (Marcovitz & Armstrong, 1992). Mazeika *et al.* (1992) comments, without reference, that "*Quantitative analysis of regional function was not performed in this study because of the known confounding effects of cardiac translation and rotation*". Comments like these are often supplemented with "*The qualitative approach adopted by us and other workers resulted in good interobserver agreement*".

A possible reason for the lack of interest in a quantitative approach has been the poor quality of images obtainable from an ultrasound machine. With the advent of improved ultrasound machine in the last five years, in particular the capabilities to save data digitally, endocardial (if maybe at times not epicardial) border definition has been significantly improved. Of course, whatever the capabilities of the machine, no machine can ever combat a patient with a poor acoustic window.

Even with the newer technology, recent published work on dobutamine stress testing (Cokkinos *et al.*, 1998; Krahwinkel *et al.*, 1997) still uses subjective scoring, and no attempt, or even mention, was paid to quantifying the procedure, or at the very least noting possible failings of a qualitative approach.

The most obvious reason for this is that no one has been able to robustly and reliably track

the endocardium and epicardium boundaries from echocardiographic images. We have seen in Chapters 5 and 7 that we can track these borders.

What little work there is on quantitating stress echocardiography, has dealt with trying to partially automate area and volume measurements (Pérez *et al.*, 1992c; Vitarelli *et al.*, 1997). The clinician then uses these (global) measurements to assign scores to heart function. This still leaves the subjective wall scoring to the clinician. More importantly, area and volume measurements are global measures that can only look at global systolic function and hence *only* global heart disease. Whereas, coronary artery disease, tends to be regionalised. In contrast, we will use our endocardial/myocardial border detection to *automatically* calculate regional scores based on endocardial wall excursion and myocardial wall thickening, as discussed in Chapters 6 and 7.

We now turn to the results of the John Radcliffe Hospital study.

### 8.3.1 John Radcliffe Hospital study results

Between October 1998 and May 1999, 11 dobutamine stress tests were performed at the John Radcliffe Hospital. Although, images were stored for each patient, several of the datasets were unusable. The reasons for this were: the data was incomplete; the imaging view recorded at rest, was not the same at peak stress; the images recorded were in quad-screen format with each quadrant showing a ciné loop of a single cycle played synchronously. Hence, we had four data sets to work with.

We analysed data from four patients at rest and at peak stress. The data is given in Table 8.4. Echocardiographic images from the eight sequences are shown in Figure 8.2. For each dataset, the endocardial and epicardial borders were tracked using the approach detailed in Chapters 5 and 7. The tracker was blind to the observers scores, and the observer was blind to the computer score.

The maximal endocardial wall excursion and myocardial wall thickening were calculated for each patient at rest and at peak stress, and the ratio of peak to rest obtained. The results for these four patients are shown in Figure 8.3 and discussed in more detail in the next section. Where an 'x' or ⊗ is placed next to the measure of wall excursion or wall thickening, this indicates, respectively, that the endocardial or epicardial wall segment was poorly or very poorly visualised. Hence, care should be taken when assessing the quantitative measures from these segments. The ratio of the peak to rest score should be greater than one for normal myocardial muscle, and less than one for ischemic myocardial muscle.

For comparison, the score awarded by two observers is given in Figure 8.4. The clinician's score

Patient	Sex	View	Image Size (Pixels)	Imaging Depth (cm)		Pixel size (cm)		Frequency (Hz)	
				Rest	Peak	Rest	Peak	Rest	Peak
1	F	4C	720 × 512	15	15	0.02646	0.02646	25	25
2	M	2C	720 × 512	15	15	0.02646	0.02646	64	60
3	F	2C	720 × 512	14	14	0.02457	0.02457	25	25
4	F	2C	720 × 512	14	13	0.02457	0.02268	67	67

Table 8.4: **Dobutamine stress echocardiography clinical test data.** The four patients we analysed from the John Radcliffe Hospital dobutamine stress echocardiography centre. 4C = 4 chamber, 2C = 2 chamber.

was based on the myocardial thickening of each segment. If however, the myocardium thickening was not clearly visible, but the endocardial border was, then the clinician's score was based on the observed endocardial excursion of each segment.

In the next section, we compare the tracker and clinician's scorings of each patient.

### 8.3.2 Comparison of tracker and clinician scores

Refer to the colour-coded results in Figure 8.3.

For patients 1, 2, and 4, the clinician graded all segments as normal, before and after stress. For patient 3, two segments were graded as abnormal at peak stress, and one segment was not graded. We will compare these gradings with those from the tracker using the ratio of the peak to rest score, for wall excursion and wall thickening.

#### Patient 1

The image quality at rest and at peak stress was the highest of the four patients. The image quality was poorest in the apical regions, because of near-field clutter in the image. The increase in wall excursion was correctly identified in two of the six segments. Part of the reason for four segments not being identified was the poor tracking of the septal wall at rest and the poor image quality in the apical septal segment. The tracker correctly identified the increase in myocardial thickening between rest and peak stress in all segments.

#### Patient 2

The images at rest and at peak were of average quality. The epicardial border was not clear along the anterior wall at rest and at peak. The increase in wall excursion was correctly identified in four of the six segments. It was incorrectly identified in the apical anterior segment because of the near-field clutter in the image. The increase in wall thickening was correctly identified in

two of the six segments. This was because the tracking of the endocardial border at rest was satisfactory.

### **Patient 3**

The image quality at rest was average, but at peak stress visual delineation of the anterior wall was very difficult. The increase in wall excursion was correctly identified in three of the six segments. The reason for incorrectly identifying three segments was the poor image quality in the anterior wall and apical inferior segment at peak stress. Correspondingly, the increase in wall thickening was not identified in these segments.

### **Patient 4**

The image quality at rest and at peak stress was average. The image quality was poorest in the apical inferior segment. The increase in wall excursion was correctly identified in all segments except the apical inferior segment, where image quality was poor. The increase in wall thickening was correctly identified in the inferior wall, but not in the anterior wall. The inconsistency in results between wall excursion and wall thickening was because the delineation of the epicardial border was not clear along the anterior wall at rest and at peak.

A summary of the comparison of the results is provided in Table 8.5. Here, we have assumed that the clinician's gradings apply for wall excursion and wall thickening. We can see that the tracker correctly classified around half of the segments for both wall excursion and wall thickening. One segment was not graded by the clinician.

Where the image quality was not poor, the tracker generally performed well, as in some of the segment in patients 1, 2 and 4. Where the image quality was poor, as in patient 3, and in some segments in patients 2 and 4, the tracker did not perform well.

Therefore, we can conclude that the image quality had a major effect on the results of the tracker. Some of other reasons for these results and the limitations of the study are discussed in the next section.




Figure 8.2: **Dobutamine stress echocardiography clinical test images.** Data from four patients who underwent dobutamine stress echocardiography at the John Radcliffe Hospital. *Left: Diastole and Right: Systole.*

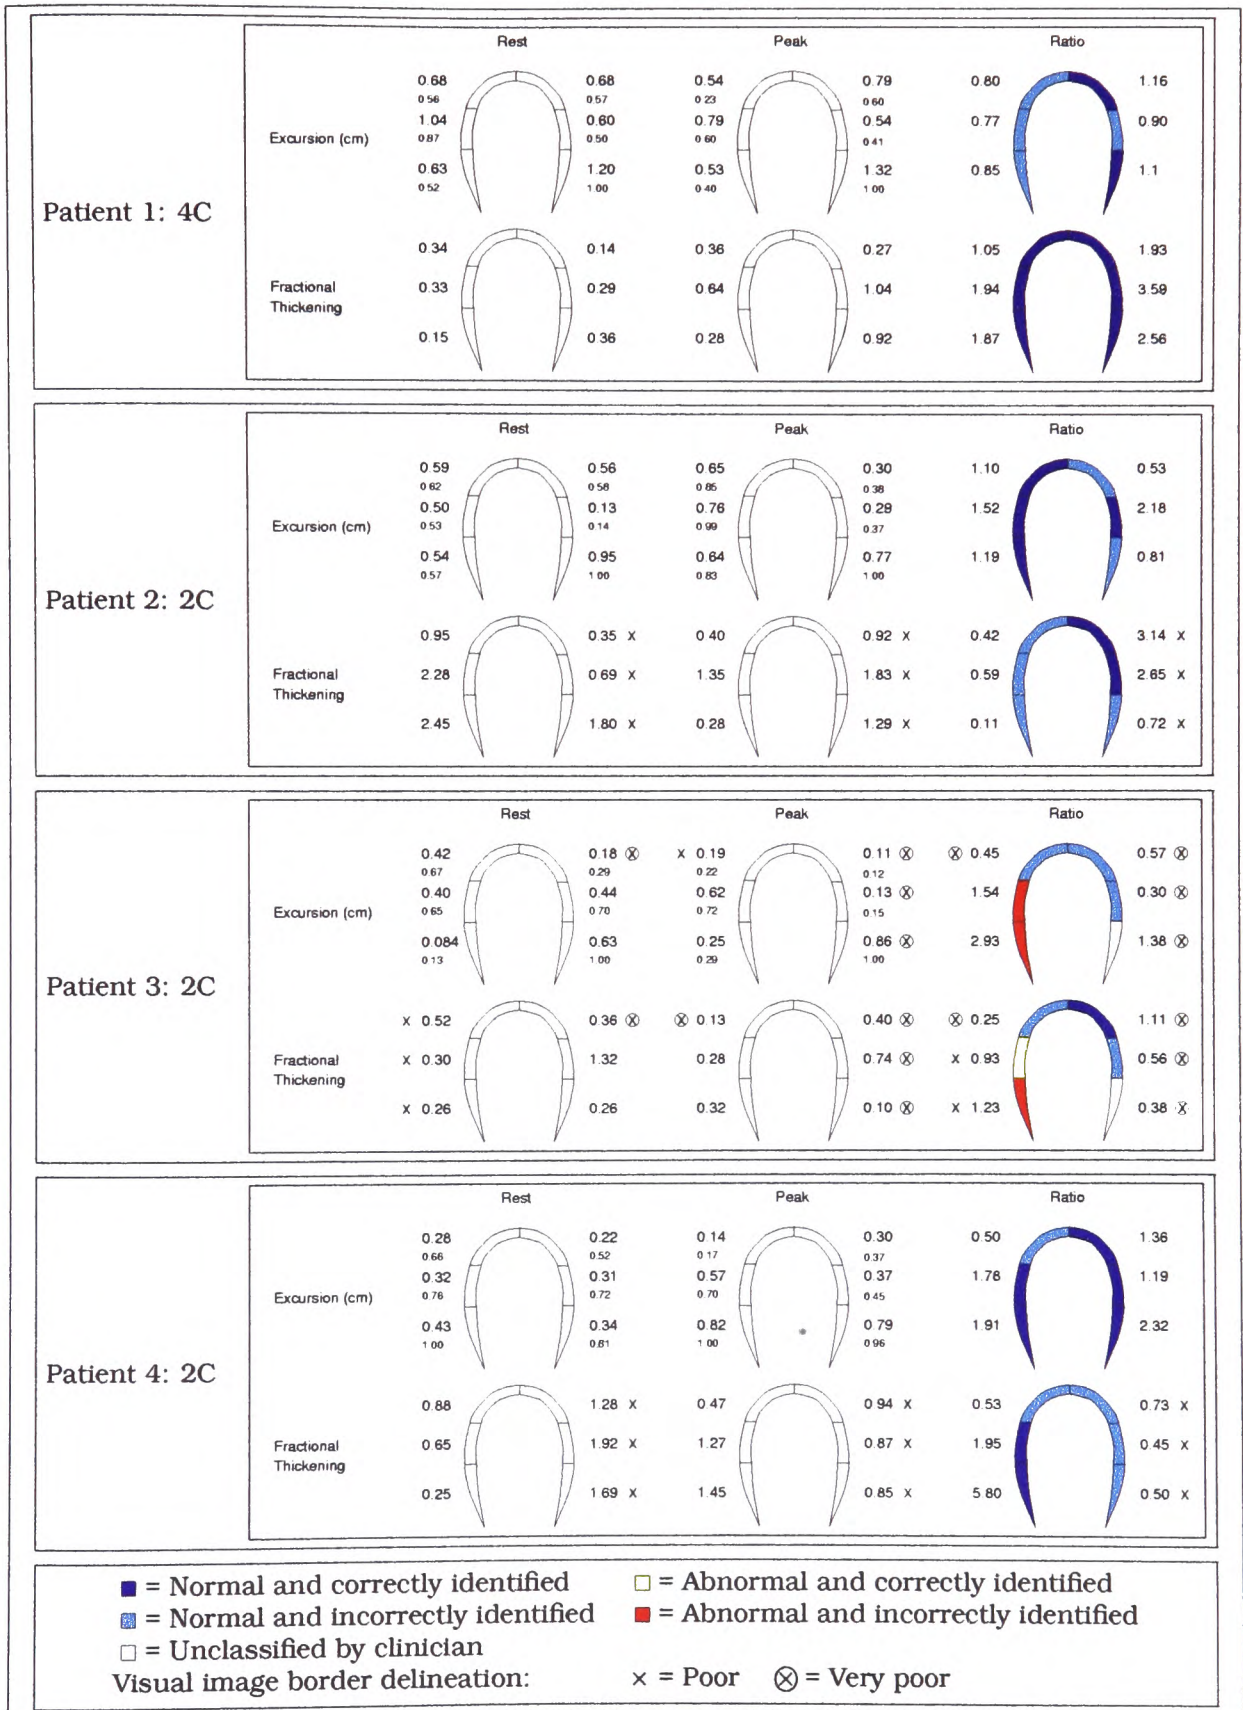


Figure 8.3: **Dobutamine stress echocardiography clinical test results.** The endocardial and epicardial boundaries of the heart were tracked. The observer and tracker were blind to their respective scores. For each patient, at rest and at peak, we calculated the endocardial wall excursion (in cm) and fractional myocardial wall thickening. In the figure, the normalised endocardial wall excursion is given in smaller font. The ratio of the peak to rest score should be greater than one for normal myocardial muscle, and less than one for ischemic myocardial muscle. Each segment is coloured coded according to the clinician's gradings




Figure 8.4: **Dobutamine stress echocardiography clinical test clinician scores.** The echocardiographic images were evaluated from the single-screen and quad-screen display by two observers and a consensus score was awarded. The clinician's score was based on the myocardial thickening of each segment. If however, the myocardium thickening was not clearly visible, but the endocardial border was, then the clinician's score was based on the observed endocardial excursion of each segment. Scores: 1=Normal, 2=Hypokinetic, 3=Akinetic, 4=Dyskinetic and ?=no score was awarded.

	Wall excursion			Wall thickening			
	Tracker			Tracker			
	Correct	Incorrect	Total	Correct	Incorrect	Total	
Clinician	Normal	10	11	21	11	10	21
	Abnormal	0	2	2	1	1	2
	Unclassified	(1)	0	1	0	(1)	1
	Total	10+(1)	13	24	12	11+(1)	24

Table 8.5: **Dobutamine stress echocardiography clinical test — summary of results.** A comparison of the classifications of the tracker against the clinician for the 24 segments in Figure 8.3. Here, we have assumed that the clinician's gradings apply for wall excursion and wall thickening. The number in the bracket indicates that the clinician did not grade a segment, but the tracker did.

### 8.3.3 Difficulties in the study

In this section, we discuss some of the difficulties we faced in analysing the stress echocardiography data and explain why the results were not as good as in the earlier example used in this thesis.

The stress echocardiography clinical service at the John Radcliffe is very new, and the routine protocol still being tuned. Importantly, this had a bearing on the quality of images recorded, which was poor in some cases. This is because, imaging patients during stress echocardiography requires additional skills than during standard echocardiography. For instance the echocardiographer must cope with imaging the patient who is breathing harder than at rest, whilst at the same time monitoring the patient's heart rate and blood pressure. As previously mentioned, logistically, the room was too small for the patient to be imaged while lying on their side, which can cause inferior image quality. When the images were very poor, tracking became very difficult because features were masked by significant clutter and noise.

It is also worth noting that the subjectivity of stress echocardiography is higher with observers with less experience. This is important because the stress echocardiography clinical service has only been running since October 1998. In an interesting paper, Picano *et al.* (1991) notes that: "... the diagnostic accuracy of a stress echocardiography procedure depends on the specific experience of the physician interpreting the test". They show that extensive expertise in routine echocardiography at rest is necessary, but that it is not sufficient to adequately interpret stress echocardiographic studies. Their study suggested that 100 stress echocardiographic studies are more than adequate to train a clinician. We were analysing four of the first 11 data sets acquired by the John Radcliffe's new unit.

We now turn to a more formal comparison of the performance of the tracker against the clinician.

## 8.4 Clinical Evaluation

One of the central problems in the area of medical image analysis is the validation of results from a segmentation algorithm. In this section, we clinically evaluate the performance of the tracking techniques described in this thesis. We assume that ground truth is taken to be a clinician's segmentation.

Until now, much of the work aimed at comparing a segmentation from an algorithm and clinician has used global measures to calculate a correlation coefficient between the algorithm and the ground truth. For example, Sonka *et al.* (1995) computed the lumen areas of coronary vessels, and used the correlation coefficient to assess the segmentation. A number of similar approaches

have been used in the context of evaluating segmentations of the left ventricle. Chu *et al.* (1988) calculated the correlation coefficient between the diameters of the automated endocardial and epicardial boundaries with known phantom dimensions. Similarly, Chalana *et al.* (1996) and Dias & Leitão (1996) compared algorithm and manual measurements of the area and volume of the left ventricle. Giachetti (1998) compared the ejection fraction from the algorithm against clinicians. Chalana & Kim (1997) calculate the average distance and Hausdorff distance<sup>2</sup> (which measures the distance between points that differ the most) between the clinician's and algorithm segmentation.

Verification of this kind is potentially very beneficial, but these methods have serious shortcomings. The problem is that the average distance between the algorithm and clinician's segmentation, as well as the area and volume of the ventricle may be totally normal, and therefore the correlation coefficient high, but the segmentation may be very poor. Moreover, knowing that the average distance between a segmentation from a clinician and two different algorithms are, for example, 8.79mm and 14.79mm, is not very enlightening, as no indication is given as to *where* and *why* there is a difference. It may be that the segmentation with the larger average distance may be correct most of the way along the boundary, but very poor in a small part of the boundary.

In contrast, we consider below a number of measures for a regional comparison of the segmentation of the algorithm and clinician. We turn to the clinician's segmentation first.

### 8.4.1 Clinician's segmentation

We asked a clinician to delineate the endocardial and epicardial borders of the left ventricle at end diastole and end-systole from the four dobutamine stress echocardiography sequences used earlier in this chapter, and the dataset used in the previous chapters.

Since the clinician was not familiar with the tracker software, we gave him a number of echocardiographic images on paper. We also felt that because the clinician was very busy, it would be quicker for him to draw the boundaries on paper.

Unfortunately, there were difficulties with this approach. The clinician was only confident about small sections of the endocardial and epicardial borders, because he could not incorporate proper temporal knowledge into his assessment (Figure 8.5 (a)).

The clinician felt that he needed to sit down at the ultrasound machine (or with a video/animation) and then browse a series of cycles to evaluate where the endocardial and epicardial boundaries were. However, the clinician had limited access to the ultrasound machine at the John Radcliffe

---

<sup>2</sup>The distance to the closest point (DCP) for a point  $a_i$  on curve  $A$  to the curve  $B$  is  $d(a_i, B) = \min_j \|b_j - a_i\|$ . The Hausdorff distance between two curves is defined as the maximum of the DCPs between the two curves.

Hospital, and felt that it would be difficult to use a video recording of the data and then draw the boundaries on paper.

Instead, the clinician used software developed by Christian Behrenbruch, which allowed the clinician to play a sequence of frames and delineate the boundaries by using a pen to draw on a graphical tablet or using a mouse to draw onto the screen.

Figures 8.5 and 8.6 show the clinician's segmentation with paper and with the software. There are three points of interest. Firstly, note that it was not possible to draw on paper around some of the datasets, for example, patient 3 at peak. Whereas, using the software, the clinician was at least able to draw part of the boundary. Secondly, that the clinician found it difficult to delineate the complete boundaries on the paper. This is shown by looking at how much of the boundaries the clinician outlined on the paper compared to using the software. Thirdly, the clinician was not able to precisely identify the end diastole and end-systole frames. This explains why the frame numbers in the figure are not exactly the same for the paper and software.

We now turn to a regional comparison of the segmentation of the algorithm and clinician.




Figure 8.5: **Clinicians segmentation with paper and with the software.** A clinician delineated the endocardial and epicardial borders of the left ventricle from the four dobutamine stress sequences used in this chapter. (a) By drawing on paper. (b) By using software, which allowed him to play a sequence of frames. Green = endocardium, blue = epicardium and / = no boundary was delineated by the clinician.




Figure 8.6: **Clinicians segmentation with paper and with software.** (a) Paper. (b) Software The same as in Figure 8.5, but for the dataset used earlier in this thesis. Green = endocardium and blue = epicardium.

### 8.4.2 Contour comparison

In this section, we provide a preliminary regional comparison of the clinician's and tracker algorithm's segmentation. A number of quantitative measures of comparison are proposed.

A visual comparison of the clinician and tracker segmentation for the four dobutamine stress echocardiography patients and the dataset used earlier in this thesis is provided in Figures 8.8 to 8.12. Here, the segmentations are coloured as follows: tracker endocardium (red), tracker epicardium (yellow), clinician endocardium (green) and clinician epicardium (blue).

Each dataset was analysed in the following way.

For each patient, at rest and peak and for the endocardial and epicardial boundaries, the following measures were calculated, and are illustrated in Figure 8.7.

1. Area — the area (in  $\text{cm}^2$ ) of the left ventricle. The tracker by default generated closed contours. The contours drawn by the clinician were closed by drawing a line between the end-points of the segmentation (Figure 8.7 (a)).
2. Global error — the global absolute error (in cm) between the clinician and tracker segmentation. The global error measures the overall closeness of the two segmentations (Figure 8.7 (b)).
3. Regional error — the regional error (in cm) between the clinician and tracker segmentation, based on dividing the left ventricle into six segments. This was obtained by calculating the principal axis of the clinician's segmentation, then dividing its maximal extent into three parts. Normals are then projected from the ends of the three parts, and the positions of where the normals cut the two segmentations are recorded. This defines six segments. The regional error of each segment measures the closeness of the two segmentations over this part of the left ventricular boundary (Figure 8.7 (c)).
4. Regional division — the percentage of the regional error that is made up from the tracker

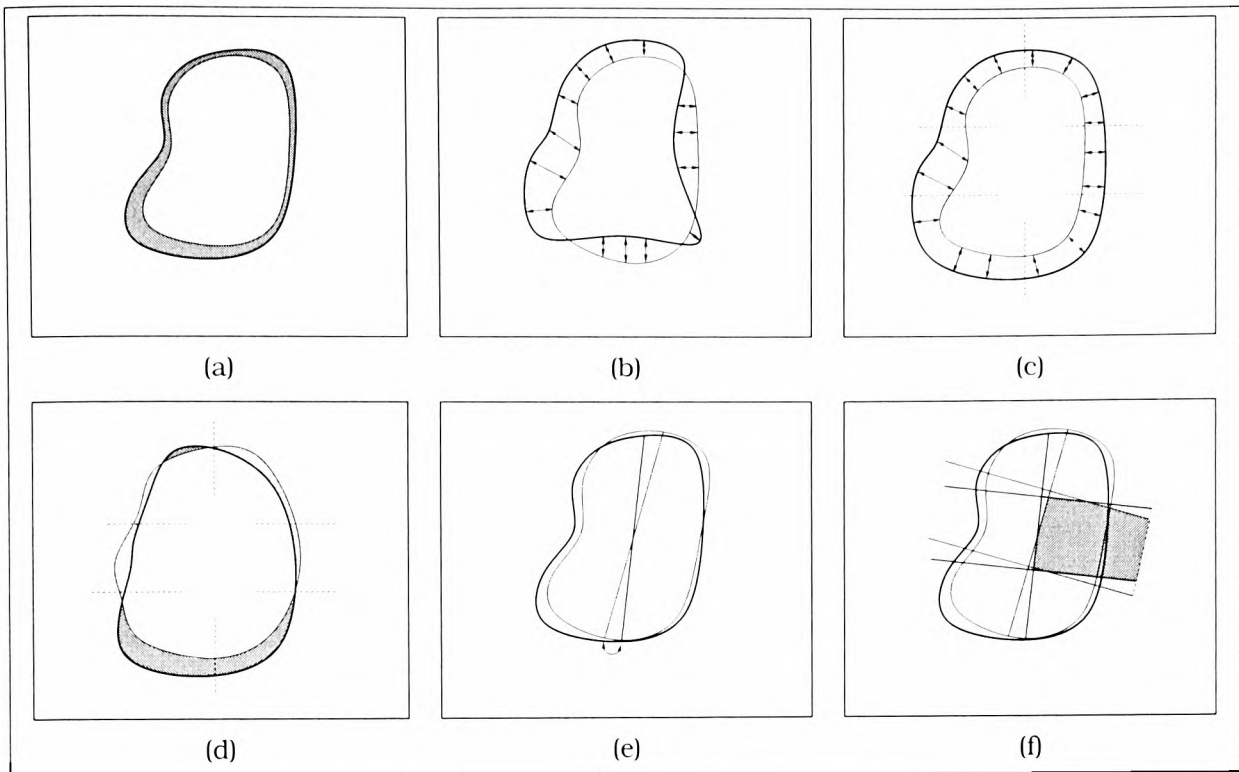


Figure 8.7: **Measures for a regional comparison of two segmentations.** (a) Area. The area (in  $\text{cm}^2$ ) of the left ventricle from the clinician (—) and tracker (—) segmentation. (b) Global error. The global absolute error ( $\leftrightarrow$ ) (in cm) between the clinician and tracker segmentation. (c) Regional error. The regional error (in cm) between the clinician and tracker segmentation, based on dividing the left ventricle into six segments. (d) Regional division. The percentage of the regional error that is from position of the tracker segmentation (—) (light shading) being outside the position of the clinician's segmentation (—) (dark shading). (e) Global rotation. The angular difference (in degrees) between the principal axis of the clinician (—) and tracker segmentation (—). (f) Regional overlap. For each segment, the percentage of the tracker and clinician's segmentation that overlap each other. i.e. the dark shaded region divided by the dark plus light shaded region.

boundary lying outside the clinician's boundary. High values indicate that the tracker boundary lies consistently outside the clinician's boundary. The regional division measures whether the tracker segmentation is biased in one direction (Figure 8.7 (d)).

A further two measures were also considered.

1. Global rotation — the angular difference (in degrees) between the principal axis of the clinician and tracker segmentation. The global rotation measures the difference in rotation between the two segmentations (Figure 8.7 (e)).
2. Regional overlap — the regional percentage overlap between the clinician and tracker segmentation. This measures the amount of overlap between the respective segments of the clinician's and tracker's segmentation. This is a comparative measure of horizontal and vertical extent (Figure 8.7 (f)).

The results for the four stress echocardiography data sets and running example are given in

Tables I.1 to Table I.3 of Appendix I.

For brevity, a selection of the results are provided here: Table 8.6 gives the area, global and regional errors between the clinician and tracker segmentation; the regional divisions are shown in Table 8.7; and global rotation and regional overlap in Table 8.8.

Following Sonka *et al.* (1995); Chalana *et al.* (1996) and Dias & Leitão (1996) a correlation coefficient between the area of the left ventricle from the algorithm and the clinician was obtained. Ignoring the inter-patient variability and variability due to stage of the cardiac cycle, boundary and drug, this was calculated to be 0.98 ( $n = 28$ ); similar values were obtained when coefficient were calculated for each patient separately. However, closer inspection of Figures 8.8 to 8.12 shows that some of the algorithm segmentations are not very close to the clinician's segmentations. In other words, simply basing the validation between the algorithm and clinician on the correlation coefficient of the two areas, leads to seemingly good results that are actually spurious.

A more detailed validation was performed using the 16-segment model. This showed that the absolute error between the clinician and tracker segmentation was badly affected by one or two segments of the algorithm segmentation being poor, for example, the endocardium estimation in diastole at rest sequence of patient 1 (Table 8.6). Further, larger errors between the clinician and tracker segmentation occurred where the tracker segmentation was consistently inside or outside the clinician segmentation; for example, the epicardium estimation in diastole at rest sequence of patient 5 (Table 8.7).

Lastly, larger errors between the clinician and tracker segmentation occurred when the horizontal and vertical extent of the two segmentations differed greatly; for example, the endocardium/epicardium estimation at rest sequence of patient 4 (Table 8.8).

The limited number of data sets and incomplete nature of the data, means that we cannot draw strong conclusion from this work. With more data, we could look at the differences between the tracker and clinician's segmentation, accounting for the variation due to the patients, stage of the cardiac cycle, boundary and drug. This would be a topic for future work.




Figure 8.8: **Patient 1. Clinician and tracker segmentations.** A clinician delineated the endocardial (green) and epicardial (blue) boundaries of the left ventricle from the dobutamine stress sequence of patient 1. For comparison, the endocardial (red) and epicardial (yellow) boundaries from the tracker are shown. RS = Rest systole, RD = Rest diastole, PS = Peak systole, PD = Peak diastole, EN = endocardium and EP = epicardium.




Figure 8.9: **Patient 2. Clinician and tracker segmentations.** A clinician delineated the endocardial (green) and epicardial (blue) boundaries of the left ventricle from the dobutamine stress sequence of patient 2. For comparison, the endocardial (red) and epicardial (yellow) boundaries from the tracker are shown. RS = Rest systole, RD = Rest diastole, PS = Peak systole, PD = Peak diastole, EN = endocardium and EP = epicardium.




Figure 8.10: **Patient 3. Clinician and tracker segmentations.** A clinician delineated the endocardial (green) and epicardial (blue) boundaries of the left ventricle from the dobutamine stress sequence of patient 3. For comparison, the endocardial (red) and epicardial (yellow) boundaries from the tracker are shown. RS = Rest systole, RD = Rest diastole, PS = Peak systole, PD = Peak diastole, EN = endocardium and EP = epicardium.




Figure 8.11: **Patient 4. Clinician and tracker segmentations.** A clinician delineated the endocardial (green) and epicardial (blue) boundaries of the left ventricle from the dobutamine stress sequence of patient 4. For comparison, the endocardial (red) and epicardial (yellow) boundaries from the tracker are shown. RS = Rest systole, RD = Rest diastole, PS = Peak systole, PD = Peak diastole, EN = endocardium and EP = epicardium.




Figure 8.12: **Example used in previous chapters. Clinician and tracker segmentations.** A clinician delineated the endocardial (green) and epicardial (blue) boundaries of the left ventricle from the dataset used earlier in this thesis. For comparison, the endocardial (red) and epicardial (yellow) boundaries from the tracker are shown. RS = Rest systole, RD = Rest diastole, EN = endocardium and EP = epicardium.

		Area (cm <sup>2</sup> )		Global Error (cm)	Regional Error (cm)						
		Track	Clin								
P1	En	RD	14.15	14.64	0.21(0.11)	0.16(0.072)	0.098(0.060)	0.33(0.17)	0.43(0.25)	0.073(0.038)	0.15(0.074)
	Ep	RD	18.17	21.37	0.33(0.13)	0.12(0.063)	0.12(0.040)	0.63(0.30)	0.77(0.27)	0.11(0.059)	0.18(0.054)
	En	PD	11.33	9.25	0.11(0.05)	0.12(0.12)	0.028(0.0032)	0.031(0.0053)	0.027(0.0019)	0.093(0.12)	0.33(0.076)
	Ep	PD	18.41	17.32	0.39(0.21)	0.24(0.15)	0.46(0.11)	0.58(0.31)	0.41(0.14)	0.27(0.16)	0.26(0.40)
P4	En	PS	10.00	12.10	0.29(0.19)	0.19(0.14)	0.19(0.065)	0.12(0.10)	0.52(0.23)	0.42(0.17)	0.29(0.44)
	Ep	PS	15.91	18.07	0.41(0.42)	0.30(0.44)	0.99(1.32)	0.26(0.35)	0.41(0.18)	0.28(0.10)	0.22(0.12)
P5	En	RD	40.03	36.99	0.31(0.13)	0.23(0.12)	0.23(0.057)	0.19(0.092)	0.32(0.18)	0.57(0.087)	0.30(0.24)
	Ep	RD	55.54	50.66	0.61(0.35)	0.95(0.32)	0.84(0.88)	0.43(0.29)	0.2(0.12)	0.64(0.21)	0.59(0.28)

Table 8.6: **Global and regional error.** The global and regional error between the clinician and tracker segmentation. The regional area is based on the 16-segment model. This measures the regional closeness of the two segmentations. P1 = patient 1, P2 = patient 2 etc. Patients one to four had a dobutamine stress echocardiography. Patient five is that used throughout this thesis. RS = Rest systole, RD = Rest diastole, PS = Peak systole, PD = Peak diastole, EN = endocardium, EP = epicardium and / = no results were calculated because only part of the boundary was delineated by the clinician. Table I.1 provides a complete set of results.

			Regional Division (%)						
P1	En	RD	100	100	11.76	4.41	38.52	95.21	
	Ep	RD	53.46	0	0	0	25.50	93.18	
	En	PD	99.73	38.71	13.67	100	100	100	
	Ep	PD	22.73	0	88.18	46.71	89.44	100	
P4	En	PS	53.43	100	25.51	0	0	1.35	
	Ep	PS	69.14	14.20	2.19	0	0	24.31	
P5	En	RD	90.35	100	23.37	19.84	100	98.49	
	Ep	RD	93.38	0	0.97	15.24	100	99.82	

Table 8.7: **Regional division.** The regional division between the the clinician and tracker segmentation. The regional division is the percentage of the regional error that is from the tracker segmentation being outside the clinician's segmentation. High values indicate that the tracker is consistently outside the clinician's segmentation. This measures whether the tracker segmentation is biased in one direction. P1 = patient 1, P2 = patient 2 etc. Patients one to four had a dobutamine stress echocardiography. Patient five is that used throughout this thesis. RS = Rest systole, RD = Rest diastole, PS = Peak systole, PD = Peak diastole, EN = endocardium, EP = epicardium and / = no results were calculated because only part of the boundary was delineated by the clinician. Table I.2 provides a complete set of results.

			Global Rotation (°)	Regional Overlap (%)						
P1	En	RD	1.92	70.03	54.50	45.89	54.34	63.01	76.08	
	Ep	RD	1.44	68.58	47.79	33.48	36.82	47.76	64.65	
	En	PD	1.52	72.72	83.16	87.86	67.87	69.06	63.51	
	Ep	PD	11.82	55.38	68.28	61.39	52.98	54.72	56.86	
P4	En	PS	0.67	65.19	60.38	55.98	54.39	60.90	67.54	
	Ep	PS	0.56	60.05	66.41	74.14	77.83	71.48	63.65	
P5	En	RD	2.70	69.29	68.19	73.84	73.71	71.04	65.99	
	Ep	RD	7.25	82.60	74.17	68.30	62.48	63.29	69.70	

Table 8.8: **Global rotation and regional overlap.** The global rotation and regional overlap between the clinician and tracker segmentation. The global rotation is the angular difference between the principal axis of the clinician and tracker segmentation. This measures the difference in rotation between the two segmentations. The regional percentage is the overlap between the clinician and tracker segmentation. This is a comparative measure of horizontal and vertical extent. P1 = patient 1, P2 = patient 2 etc. Patients one to four had a dobutamine stress echocardiography. Patient five is that used throughout this thesis. RS = Rest systole, RD = Rest diastole, PS = Peak systole, PD = Peak diastole, EN = endocardium, EP = epicardium and / = no results were calculated because only part of the boundary was delineated by the clinician. Table I.3 provides a complete set of results.

## 8.5 Conclusions

This chapter has assessed the performance of tracking algorithms in a retrospective clinical study using stress echocardiography. We analysed four patients at rest and at peak stress. We showed that we could correctly identify the increase in regional wall excursion and wall

thickening between rest and peak stress, provided that the image quality was not poor. If it was, then the tracking algorithm was not accurate and hence it incorrectly identified the regional wall function. We discussed some of the difficulties with the study.

We also evaluated the performance of the tracking techniques described in this thesis against ground truth, which was taken to be a clinician's segmentation.

We proposed a number of measures for a regional comparison of an algorithm and clinician segmentation and discussed these in relation to the stress echocardiography data available for our study.

# CHAPTER 9

---

## Discussion and Future Work

---

This chapter brings together the issues arising from the work in this thesis, and suggests some areas of future research which may prove fruitful.

### 9.1 The work in this thesis

This thesis has developed spatio-temporal tracking techniques for quantifying regional left ventricular function from echocardiographic sequences.

The approach consists of two parts: (1) tracking the endocardial and epicardial boundaries of the left ventricle (Chapters 5 and 7), and (2) using the estimated position of these boundaries to quantitatively measure regional endocardial wall excursion and myocardial wall thickening (Chapters 6 and 7). The entire approach has been used in a small retrospective clinical test to quantify changes in left ventricular function during stress echocardiography (Chapter 8).

The automatic quantitative approach detailed in this thesis, demonstrates a significant contribution to the field of quantitatively assessing regional left ventricular function. Until now, quantitative measures used were global, for example, the area of the left ventricular chamber. In this thesis, we have shown that we can provide clinically meaningful regional measures of wall motion. This work improves upon previous work, and the ideas have been applied to routine clinical datasets.

A visual tracker has been successfully adapted to track the motion of the endocardium. The tracker is based on a Kalman filtering framework. This approach has three key tracking ingredients: an appropriate model of shape; modelling and learning the motion of the object to guide a prediction; and using specific image feature detectors to enhance the measurement process, which in turn updates the position of the contour. We showed that a PCA shape basis was better than an affine basis at providing the correct deformations of the left ventricle. We went on to show that feature detection was improved using intensity-invariant based feature detection.

A quantitative measure of regional endocardial wall excursion, based on a 16-segment model of the heart, was obtained using the interpretational shape-space developed in Chapter 6. This provides a clinically meaningful quantitative way of assessing regional heart function.

The estimated position of the endocardium forms a starting point in searching for the epicardium. We showed that it was better to model the shape-deformation of the difference between the epicardium and endocardium, rather than just the epicardium. Features along the epicardium were found using a wavelet-based ridge detector, which assimilated wavelet information along the contour.

The time-varying difference between the endocardium and epicardium determines the myocardial wall thickening. We have shown that we can quantitatively measure regional thickening using the 16-segment model, and provide a clinically meaningful quantitative way of assessing regional heart function.

Our algorithms were applied to stress echocardiography. We analysed four patients at rest and at peak stress. We showed that we could correctly identify the increase in regional wall excursion and wall thickening between rest and peak stress, provided that the image quality was not poor.

We went on to clinically validate the results from the segmentation algorithm against the ground truth, which was taken to be a clinician's segmentation. A number of quantitative measures for a regional comparison of an algorithm and clinician segmentation were proposed.

The approach described in this thesis is aimed at situations in which quantitative rather than visual assessments of left ventricular function are needed. Furthermore, it is particularly suited to situations where inter- and intra-observer bias can occur, for example, in monitoring patients over time and in stress echocardiography.

To make these quantitative measures more clinically useful, we showed how they could be visualised as a number of colour-coded charts to depict wall motion and thickening. For instance, to visualise regional velocities of the endocardial wall. We also extended the idea of colour kinesis to myocardial wall thickening.

Further work aside, the techniques provide an original and partially evaluated methodology for analysing regional heart function, that has been shown to have enormous promise in the context of quantitative echocardiography, and an ideal application in stress echocardiography.

## 9.2 Future work

The work, as it has been presented in this thesis, is complete, but there is scope for further improvements to the techniques.

## 9.3 Building models for all views

The algorithms developed in this thesis have only been tested on the apical views of the left ventricle, and for five of the six datasets, the 2-chamber view was selected. In fact this is the hardest view for the clinician to obtain, and hence images tend to be poorest in quality. There, is however, nothing in the methodology discussed that prevents the approach from being applied to any of the four standard views of the left ventricle. Further work could build models for each view and also consider how to fuse information from a sparse set of views, for example, to measure the left ventricular mass.

## 9.4 Using Condensation

The tracker makes an implicit assumption that the measurement noise has a Gaussian distribution. This is in fact, false, because of the potential clutter and speckle noise along the search lines. Although in practice, this assumption does not seem to considerably affect tracking, there may be something to gain by dropping the Gaussian assumption in the filter.

The tracker is based on a Kalman filtering framework, which makes a single prediction of the boundary position on the next frame. Echocardiographic images are very noisy, and can contain significant imaging artifacts. It might be more sensible to use a filtering scheme that makes many predictions, which are then combined to form an estimate of the position of the object. The Condensation algorithm developed by Isard (1998) is one such algorithm that works on this philosophy. Isard (1998) and North (1998) have shown the benefit of using this filter, compared to the Kalman filter, for tracking human hands. It might be beneficial to apply this filter to echocardiographic sequences.

## 9.5 Using temporal information for feature detection

In practice, it is considerably easier for a clinician to delineate the endocardial and epicardial borders, after watching a sequence of echocardiographic images. This begs the question about using temporal information for feature detection, since structure should be preserved between frames. We have not currently considered this idea, but current work by Mulet-Parada (1999) is aimed at doing this. This could, in time, be incorporated into the measurement process of the tracker.

## 9.6 Epicardial feature detection

Image features along the epicardium were located using a wavelet-based ridge feature detector. Localisation was improved by reconstructing profiles with consideration given to the wavelet packet decomposition for previous profiles. This weighting reflected the similarity between successive profiles and for the evaluation performed in this thesis, was either set to zero (independent case) or unity (complete dependence on the previous localisation of the edge).

Ideally, localisation would be based on an adaptively varying evaluation of this statistical weighting function. This development is likely to be the focus of further work.

## 9.7 Shape space for epicardium estimation

In Section 7.4 it was shown that modelling the shape-deformation of the epicardium using a shape-space did not constrain the position of the epicardium. Instead, we modelled the difference between the epicardium and endocardium. The epicardium was fitted in the difference shape-space and a regulariser,  $\alpha$ , was used to control the trade off between high noise resistance and the bias towards a mean shape.

The problem with this method is that a high value of  $\alpha$ , will provides greater tolerance to noisy data, but then the difference between the endocardium and epicardium becomes close to a constant.

A curve-fitting problem solution would be to solve:

$$\min_{\mathbf{X}} T \quad \text{where} \quad T = (\mathbf{X} - \bar{\mathbf{X}})^T \bar{\mathbf{S}}_{\text{Ep}} (\mathbf{X} - \bar{\mathbf{X}}) + (\mathbf{X} - \bar{\mathbf{X}})^T \bar{\mathbf{S}}_{\text{Diff}} (\mathbf{X} - \bar{\mathbf{X}}) + \sum_{i=1}^N \frac{1}{\sigma_i^2} (\nu_i - \mathbf{h}(s_i)^T [\mathbf{X} - \bar{\mathbf{X}}])^2.$$

Here there are two regularisers that can be used to impose prior constraints on the likely shape of the contour:  $\bar{\mathbf{S}}_{\text{Ep}}$  imposes shape constraints on the epicardium position and  $\bar{\mathbf{S}}_{\text{Diff}}$  imposes

shape constraints on the difference between the endocardium and epicardium.

## 9.8 Quantitative scoring of wall function

At present, the quantitative scoring of wall function is based on two indices: endocardial wall excursion and (normalised) myocardial wall thickening. These indices are calculated using the maximal and minimal endocardial wall excursion and myocardial thickening of the entire echocardiographic sequence, not for each individual cardiac cycle. This is a limitation of the current approach which poses two difficulties. (1) There is no guarantee that the maximum and minimum wall excursion and thickening used to calculate these indices, come from the same cardiac cycle. (2) Spurious observations, or outliers have an effect on the scoring. Both of these effects could be minimised if the stage of the cardiac cycle were known, perhaps using the ECG information. The indices could then be calculated on a cycle-by-cycle basis, with the published score then the average of the scores on each cardiac cycle.

## 9.9 Towards a tracker-driven user-interface

At present, the data is collected at the John Radcliffe, and transferred to the laboratory for analysis. The results are subsequently then taken back to the clinicians. For the software to have clinical use, this middle stage needs to be eliminated. In this section, we shall address a number of practical issues relevant to turning the current tracker software into a clinically useful system for assessing regional heart function.

In this thesis, accurate tracking of the endocardial and epicardial borders was demonstrated on several different patients; however, the tracker was *patient dependent*, i.e. each heart tracker employed shape and motion models that were particular to the given patient. Although a basic tenet of modelling is to incorporate as much prior knowledge as possible, from a practical perspective, it is impractical to require that new models be developed from scratch every time a new patient dataset is to be tracked.

There are two potential ways of overcoming this problems. (1) Learn the shape and motion of each patient on-line, as at present, but consider ways of speeding up this process. (2) Construct a general generic model of the shape and motion of the left ventricle that is capable of tracking the majority of hearts. We will address each solution in turn.

### 9.9.1 Patient dependent tracker framework

We could attempt to make the tracker patient dependent. In this approach, our aim is to learn the shape and motion of each patient on-line. To do this, we need an initial template, shape-space and motion model.

#### Initialisation

Unlike Acoustic Quantification (AQ) (Mor-Avi *et al.*, 1995), the initial position of the tracker we use needs to be initialised, as do its initial parameters. In other words, the state of the system at  $t = 0$  was provided. It would be preferable to remove the necessity of this, and allow the system to learn the initial distribution of the state.

#### The shape-space model

We showed in Chapters 5 and 7 that the shape deformation of the endocardium and epicardium, respectively, should be modelled using a PCA shape-space for good tracking. This shape-space is built from a training set of manually segmented endocardium and epicardium boundaries. In practice, this would be too time-consuming for a clinician to do during an echocardiographic examination. The alternative is to use an affine shape-space, which can be directly constructed from the initial template, and hence requires no pre-training. The problem here is that the affine basis was shown both empirically and mathematically to be incapable of providing the correct deformations.

A possible solution is to use a bootstrapping technique as follows. Track a number of echocardiographic frames, for example, a complete cardiac cycle, using an affine basis which has been constructed from the initial template, or a PCA shape-space obtained from a previous study. Then using this tracked sequence, perform a PCA on the estimated contours, and generate a PCA shape-space. This PCA shape-space is then used to track a further cardiac cycle, and the results from this sequence are used to generate a new PCA shape-space. Bootstrap to convergence.

#### Learning the dynamics

In order to learn the dynamics of the heart, a training set needs to be collected, which in turn requires some form of un-trained tracker. Since the dynamics of any specific heart are unknown it is necessary to use a tracker with reasonable default dynamics. An untrained tracker is less powerful than a trained tracker, particularly in clutter and with rapid motions.

In visual tracking, a common technique is to bootstrap, that is: first to use a default tracker to learn a relatively slow motion, often in a noise free background, then incorporate that learnt motion into a tracker which should then be more agile than the default, and use it to learn a somewhat faster motion, and so on, repeating the cycle if necessary.

In the case of echocardiography, altering the left ventricle's speed in a noise free background seems implausible. A potential solution is to start with a standard motion model or for the clinician to choose one.

### 9.9.2 Generic tracker framework

To build a generic tracker framework, we could develop a library of endocardial and epicardial templates, shape-spaces, and motion models from patients with a wide range of left ventricular shapes and heart dynamics. During initialisation, a method would be needed for the clinician to determine the most appropriate model for the patient.

A more elegant solution is to build a generic tracker model based on this library. To do this, we need a generic shape-space model, and a generic motion model.

#### Generic shape-space estimation

Recall, that to obtain a PCA shape-space, we perform a PCA on a covariance matrix, obtained from a number of training examples.

In general, the dimension of the shape-space will be very different for different patients. Moreover, there is no guarantee that each individual deformation mode will be the same for different patients. Therefore, we need some way of combining together the information from the manually segmented endocardium and epicardium boundaries obtained from different patients.

One way of doing this is to stack together the training examples from each patient, construct a covariance matrix, and then perform a PCA on this matrix. The subtle problem with this is that the resulting co-variance matrix then has unwanted cross-correlation terms between subjects. A better solution is to perform a PCA on the average of the individual covariance matrices.

#### Generic motion model

Parameter estimation for a generic motion model could be based on a fixed-effects or random-effects model (Crowder & Hand, 1990), or simply using an 'average' model.

Let the parameters in the generic motion model be,  $A_1^g$ ,  $A_2^g$  and  $B_0^g$ .

### Fixed effects model

In the fixed effects model, we assume that there is no difference in the dynamics of the heart for different patients. The dynamical model then is,

$$\mathbf{X}_t - \bar{\mathbf{X}} = \mathbf{A}_1^g(\mathbf{X}_{t-1} - \bar{\mathbf{X}}) + \mathbf{A}_2^g(\mathbf{X}_{t-2} - \bar{\mathbf{X}}) + \mathbf{B}_0^g \mathbf{w}_t. \quad (9.1)$$

The parameters in the model are then estimated as before using maximum likelihood estimation, but the likelihood is then the product over all patients — see Appendix E.4.

The underlying assumption of the fixed effect model is that the motion of an object within a class is the same for all individuals within that class. This idea was used by Reynard *et al.* (1996) for tracking automatic crop spraying of cabbages .

We can avoid this assumption by introducing a random effects model to take account of variability within a class.

### Random effects model

A random effects model makes the assumption that each parameter,  $\theta$ , say, is specific to the individual within a class, and can be expressed in the form,  $\theta_i = \theta + e_i$ , where  $\theta$  is a constant, and the  $e_i$ 's model the variation within the class. We assume that model parameters are independent. The dynamical model then is,

$$\mathbf{X}_t - \bar{\mathbf{X}} = (\mathbf{A}_1^g + \alpha_1)(\mathbf{X}_{t-1} - \bar{\mathbf{X}}) + (\mathbf{A}_2^g + \alpha_2)(\mathbf{X}_{t-2} - \bar{\mathbf{X}}) + \mathbf{B}_0^g \mathbf{w}_t. \quad (9.2)$$

Here,  $\alpha_1$  and  $\alpha_2$  are patient specific parameters that model the class variation. The parameters can be estimated, non-trivially, using restricted maximum likelihood estimation (REML).

In practice, a third option is to estimate the generic model parameters by the average of the individual patient model parameters.

### Average model

Let the number of classes be  $\mathcal{C}$ . Let the number of objects within a class be  $\mathcal{K}_c$ . Suppose that each object within a class is observed in motion for  $\mathcal{M}_k$  timesteps. We seek to estimate the dynamics  $\mathbf{A}_0, \mathbf{A}_1$  and  $\mathbf{B}$ .

For each data set  $k$  we can estimate the system matrices  $\mathbf{A}_0^k, \mathbf{A}_1^k$  and  $\mathbf{B}^k$ . We can then construct a general model for tracking of the form,

$$\mathbf{X}_t - \bar{\mathbf{X}} = \mathbf{A}_1^g(\mathbf{X}_{t-1} - \bar{\mathbf{X}}) + \mathbf{A}_2^g(\mathbf{X}_{t-2} - \bar{\mathbf{X}}) + \mathbf{B}_0^g \mathbf{w}_t, \quad (9.3)$$

where,

$$\mathbf{A}_0^g = \sum_{k=1}^{\mathcal{K}_c} w_k \mathbf{A}_0^k, \quad \mathbf{A}_1^g = \sum_{k=1}^{\mathcal{K}_c} w_k \mathbf{A}_1^k, \quad \mathbf{B}_0^g = \sum_{k=1}^{\mathcal{K}_c} w_k \mathbf{B}_0^k, \quad w_k = 1/\mathcal{K}_c. \quad (9.4)$$

However, this does not take into account the number of timesteps each object was observed for. We could allow for this by setting the weight  $w_k$  to be,

$$w_k = \frac{\mathcal{M}_k}{\sum_{k=1}^{\mathcal{K}_c} \mathcal{M}_k}.$$

### 9.9.3 3-dimensional echocardiography

The work in this thesis considers 2-dimensional echocardiographic images. Recent interest has focused on using 3-dimensional echocardiography to view the valves and chambers of the heart (Bardinet *et al.*, 1996; Park *et al.*, 1996a; Rohling *et al.*, 1997; Sheehan *et al.*, 1998; Sanchez-Ortiz *et al.*, 1999a).

The tracking framework discussed in this thesis, readily extends to 3-dimensions. The challenge will be to find reliable feature detection for the endocardial and epicardial borders. The idea of using an interpretational shape-space generalises to 3-dimensions, and therefore we can calculate quantitative measure of regional endocardial wall excursion, based on a 16-segment model of the heart.

# APPENDIX A

---

## Glossary of clinical terms

---

<b>TERM</b>	<b>MEANING</b>
16-segment model	The ventricular boundary is divided into 16 segments, corresponding to significant blood supplies to the heart surface. This model is used to assess regional heart function (see Figure 6.4, page 85)
Accuracy	The ratio of the number of correct assessments (true positive and true negative) divided by the total number of assessments
Acoustic Quantification (AQ)	An automatic boundary detection scheme used in Hewlett Packard Ultrasound machines
Akinesia	Absence or lack of movement of a region of myocardium
Angiogram	An image of the blood vessels
Angiocardiology	An X-ray examination of the heart and coronary arteries; a liquid which shows up on X-rays is injected through a catheter into the heart. normally based on digital subtraction angiography
Aorta	The large artery leading out of the left side of the heart and feeding the whole body
Artery	A main blood vessel carrying blood under high pressure from the heart to the rest of the body

<b>TERM</b>	<b>MEANING</b>
Asyneresis	Localized decreased movement of a region of myocardium
Asynchrony	Disturbed temporal sequence of contraction
Atropine	A drug that produces tachycardia, which can then increase myocardial oxygen consumption. It produces tachycardia by blocking parasympathetic vagal output
Atrium	The two upper (left and right) chambers of the heart receiving blood on the right side from the body, and the left side from the lungs. They act as a collecting reservoir to fill the ventricles
Attenuation	Decrease in amplitude and intensity as a wave travels through a medium
Backscatter	Sound scattered back in the direction from which it originates
Capillaries	The smallest of the blood vessels joining the small arteries to the small veins where the majority of diffusion into tissues occurs — according to Starling's laws
Cardiomyopathy	Disease, usually of unknown cause, which weakens the heart muscle
Cardiovascular	Pertaining to the heart and blood vessels
Cardiovascular disease (CVD)	Cardiovascular disease includes all disease of the heart and blood vessels. The two main diseases in this category are coronary heart disease (CHD) and stroke, but CVD also includes congenital heart disease (i.e. the heart deformities present at birth), heart failure normally affecting older people and a range of other diseases of the heart and blood vessels. CHD and stroke have a similar cause — a blockage in an artery generally caused by atherosclerosis and thrombus formation
Chronotropic effects	Affects the time or rate of contraction of the heart
Colour kinesis (CK)	Colour kinesis analyzes backscatter in each acoustic frame in real time and classifies each pixel as either blood or myocardial tissue. Pixel transitions from tissue to blood during systole are detected and colour coded
Coronary angiogram	An image to show where the arteries are narrowed and how narrow they have become

<b>TERM</b>	<b>MEANING</b>
Coronary angioplasty	A treatment to improve the blood supply to the heart muscle. A catheter with a small inflatable balloon at its tip is inserted into an artery in the groin and threaded through the narrowed artery. The balloon is then inflated so that it squashes the fatty tissue (atherosclerotic plaques) responsible for the narrowing, and widens the artery
Coronary artery bypass surgery	An operation to bypass narrowed sections of the coronary arteries
Coronary arteries	Arteries arising from the beginning of the aorta and feeding the heart muscle with blood
Coronary artery disease (CAD)	Coronary arteries which are affected by a pathological process. Coronary artery disease usually occurs for many years before a disorder of the myocardial function develops.
Coronary heart disease (CHD)	Coronary heart disease comes in two main forms: Heart attack and angina. Angina is a pain in the chest brought on by exercise or emotion. It can be mild or severe and typically lasts less than 10 minutes. A heart attack causes similar pain but lasts longer and can be fatal. Angina is caused by a narrowing of the blood vessels to the heart muscle. A heart attack results when one of those vessels is entirely blocked by a blood clot
Diastole	The portion of the cardiac cycle that occurs when the heart is relaxed (thus dilated) and largely involves ventricular filling (passive)
Dobutamine	A drug that causes the heart to demand more oxygen
Doppler Ultrasound	A method of investigating the speed and direction of blood flow in the heart and blood vessels by high frequency sound waves
Dyskinesia	The impairment of the power of voluntary movement, resulting in fragmentary or incomplete movements of a region of myocardium
Echocardiogram	Recording of sound wave reflected from different parts of the heart to provide pictures of the heart structure and movement

<b>TERM</b>	<b>MEANING</b>
Echocardiography	A noninvasive, usually 2-dimensional, diagnostic method that utilizes high frequency sound to assess anatomical and functional abnormalities of the cardiovascular system
Electrocardiogram (ECG)	Recording of the electrical impulses produced by the heart
Ejection Fraction	The percentage or fraction of left ventricular diastolic volume that is ejected at systole. A normal heart has an ejection fraction of between 0.4–0.8 (Feigenbaum, 1994, page 669)
Endocardium	A membrane that forms a protective inner lining off the chambers and valves
Epicardium	A membrane that forms a protective outer covering of the chambers
Exercise (Stress) Test	A test of exercise capability
False Positive	Tests that are incorrectly announced positive
False Negative	Tests that are incorrectly announced negative
Hibernating myocardium	Dysfunctional myocardium due to chronic ischemia which can resolve following revascularisation
Hypertension	High blood pressure
Hypoxia	Lack of oxygen
Hypokinesia	Decreased muscular activity
Inotrope	A substance or drug which stimulates the force of the heart beat
Inotropic effects	The increase in contractility of the heart, i.e. the force of contraction and increase in cardiac output
Ischemia	Inadequate blood supply
Ischemic heart disease	A blood oxygen deficiency caused by a branch of a coronary artery becoming abnormally obstructed
Infarct	Myocardial infarction is necrosis (death) of a portion of heart muscle as a result of inadequate blood supply. This leads to scarring and loss of function
Left ventricle (LV)	Ventricular chamber that forces blood through the body

<b>TERM</b>	<b>MEANING</b>
Mitral Valve	Valve between left atrium and left ventricle
Myocardium	The muscle tissue (muscle) that is the middle layer of the heart wall. The function of the myocardium is to produce muscular contractions that force blood from the chambers
Myocardial Infarction	The process where a part of the heart muscle dies from lack of blood supply, i.e. a heart attack
Nuclear Magnetic resonance (NMR)	A technique where the body is put into a strong magnetic field which lines up the nuclei of certain atoms such as phosphorus or hydrogen in a regular manner. Another pulsed magnetic field is then introduced to disturb this alignment temporarily. Computer processed information can be obtained which will give images of the body (magnetic resonance image MRI) or information about the biochemical processes occurring in the body (magnetic resonance spectroscopy MRS)
Occlusion	An obstruction that may occur in a blood vessel
Papillary Muscle	A small mound of muscle within a chamber of the heart
Pericardium	Lubricating fibrous membrane that surrounds the heart
Pulmonary	Relating to the lungs
Radioisotope	Substance emitting very tiny amount of radioactivity that can be measured and localised from outside the body to provide diagnostic information
Right ventricle (RV)	Ventricular chamber that forces blood to the lungs
ROC	Receiver operating characteristic curve. A graph of the false positive fraction (1-specificity) versus the true positive fraction (sensitivity)
Sensitivity	The ratio of true positives to those that are positive (true positives + false negatives)
Septum	A muscular wall dividing two chambers of the heart

<b>TERM</b>	<b>MEANING</b>
Specificity	The ratio of true negatives to those that are negatives (true negatives + false positives)
Stress Echocardiography	A technique where the patient is exercised (physically or not), so that the heart demands more oxygen
Systole	The portion of the cardiac cycle that occurs when the heart is contracting
Tachycardia	A fast heart rate
Transoesophageal (TOE)	The transducer probe is inside the body
Transthoracic (TEE)	The transducer probe is outside the body
True positive	Tests that are correctly announced positive
True negative	Tests that are correctly announced negative
Vascular	Pertaining to the blood vessels
Vein	A vessel carrying blood back at low pressure from various parts of the body to the heart
Ventricle	Left and right; the two main pumping chambers of the heart

# APPENDIX B

---

## Anatomical Background

---

### B.1 The Heart

This chapter deals with the anatomy of the heart, and gives a brief account of the mechanism for heart contraction. The anatomical material is largely based on Hole JR. & Koos (1994); British Heart Foundation (1992); and Weatherall *et al.* (1996). The discussion on the cardiovascular system is largely taken from Navaratnam (1975); Berne & Levy (1986); Julian & Cowan (1992) and Levick (1995).

#### B.1.1 Functions of the cardiovascular system

The heart is part of the cardiovascular system which moves blood between the body cells and the organs of the integumentary, digestive, respiratory and urinary systems. The blood is set in motion within a continuous system of conduits or blood vessels by the regular pumping action of the muscular heart. The heart acts as a pump that forces blood through the blood vessels — Figure B.1 — and is located within the mediastinum of the thorax and resting upon the diaphragm, as shown in Figure B.2.

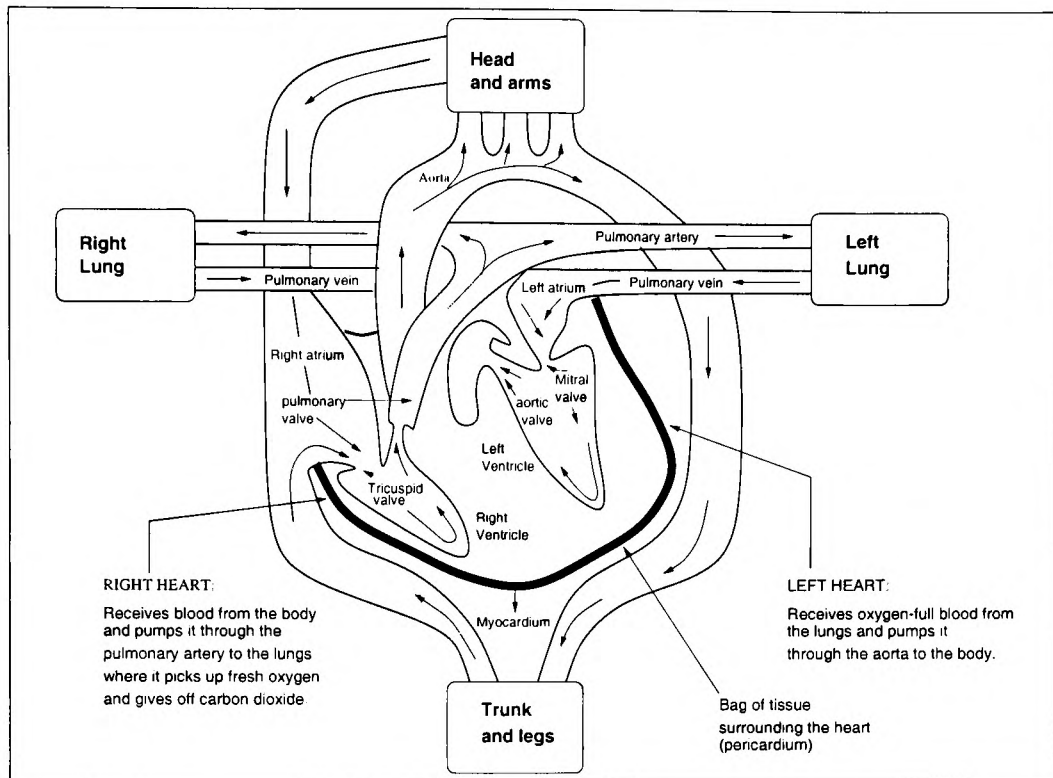


Figure B.1: **Blood flow to and from the heart.** The cardiovascular system functions to transport blood between the body cells and organs.




Figure B.2: **The location of the heart.** The heart is located behind the sternum, where it lies upon the diaphragm.

### B.1.2 Heart chambers and valves

The heart is divided into four chambers with two pumps in series; one (the right heart) propels blood through the lungs for exchange of oxygen and carbon dioxide (the *pulmonary circulation*) and the other (the left heart) propels blood to all other organs of the body (the *systemic circulation*). Unidirectional flow through the heart is achieved by the appropriate arrangement of effective flap valves (see Figure B.3 and Table B.1).

The upper *atria* chambers form a thin-walled collecting reservoir receiving blood from the veins. The lower thick-walled *ventricle* chambers force blood out of the heart into the arteries. (Note: veins are blood vessels that carry blood towards the heart; arteries carry blood away from the heart).

The atrium and ventricle on the right side are separated from those on the left by a septum. The atrium on each side communicates with its corresponding ventricle through an opening called the *atrioventricular orifice*, which is guarded by an *atrioventricular valve* (A-V valve).

The atrioventricular orifice between the right atrium and right ventricle is guarded by a large *tricuspid valve* — Figure B.3 — composed of three leaflets or cusps. This valve permits blood to move from the right atrium into the right ventricle and prevents it from passing in the opposite direction, i.e. during ventricular systole. The cusp folds passively out of the way when the blood pressure is greater on the atrial side; when the pressure is greater on the ventricular side, it closes passively.

The right ventricle has a much thinner muscular wall than the left ventricle. This is because the right chamber pumps blood a fairly short distance to the lungs. The left ventricle, on the other hand, must force blood to all other parts of the body. Pressures in the left ventricle reach 120mmHg, compared to 30mmHg in the right ventricle.

When the muscular wall of the right ventricle contracts, the blood inside its chambers is put under increasing pressure, and the tricuspid valve closes passively. The only exit is through the *pulmonary trunk*, which divides to form the left and right pulmonary arteries. At the base of this trunk, is a *pulmonary valve* that consists of three cusps, which are attached to the wall of the pulmonary trunk. The pulmonary valve prevents blood from moving from the pulmonary trunk into the right ventricle during ventricle relaxation. This valve opens as the right ventricle contracts, and the cusps are passively pushed aside. However, when the muscles relax, the blood begins to back up in the pulmonary trunk. This causes the cusps to fill and extend into the middle of the pulmonary artery, thus closing the valve. This action prevents a return blood flow into the ventricular chamber.

The left atrium receives blood from the lungs through four *pulmonary veins* — two from the right lung and two from the left. Blood passes from the left atrium into the left ventricle through the atrioventricular orifice, which is guarded by a valve. This valve consists of two leaflets, and is named the *bicuspid valve*, or, *mitral valve*. The bicuspid valve prevents blood from moving from the left ventricle into the left atrium during ventricular contraction.

When the left ventricle contracts, the bicuspid valve closes passively, and the only exit for blood is through a large artery called the *aorta*. Its branches distribute blood to all parts of the body.

At the base of the aorta is the *aortic valve* that consists of three cusps and is similar in structure and function to the pulmonary valve. The *aortic valve* prevents blood from moving from the aorta into the left ventricle during ventricular relaxation. Table B.1 summarises the location and function of the heart valves.

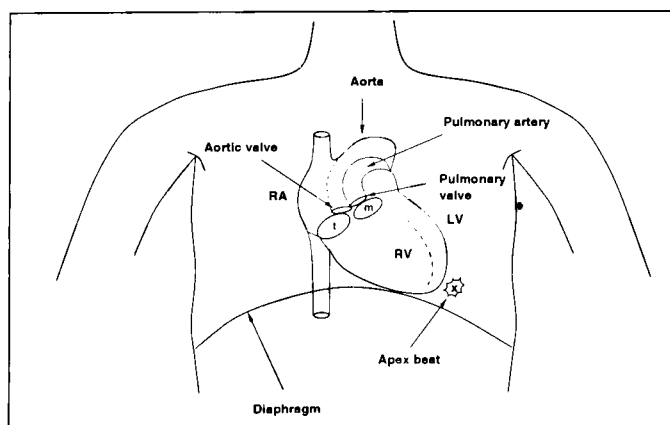


Figure B.3: **The heart valves.** The heart lies obliquely across the chest. The fibrotendinous ring acts as a base for the heart. It contains the tricuspid (t), mitral (m), aortic and pulmonary valves grouped in an oblique plane beneath the sternum. The apex of the heart is formed by the left ventricle (LV), and the anterior surface is formed by the right ventricle (RV) and right atrium (RA). The inferior surface of the heart and the pericardium rest on the central tendon of the diaphragm. Figure taken from Levick (1995).

Valve	Location	Function
Tricuspid valve	Right atrioventricular orifice	Prevents blood from moving from right ventricle to right atrium during ventricular contraction
Pulmonary valve	Entrance to pulmonary trunk	Prevents blood from moving from pulmonary trunk into right ventricle during ventricular relaxation
Bicuspid valve (mitral valve)	Left atrioventricular orifice	Prevents blood from moving from left ventricle to left atrium during ventricular contraction
Aortic valve	Entrance to aorta	Prevents blood from moving from aorta into left ventricle during ventricular relaxation

Table B.1: **The heart valves.** The heart wall consist of four valves: tricuspid, pulmonary, bicuspid (mitral valve) and aortic.

### B.1.3 Path of blood through the heart. The Circulation

Each heart contraction ejects blood forward into the arteries. The arteries divide into progressively smaller branches to supply a microscopic network of capillaries, so distributing the blood to every part of the body. Although the cardiac output is intermittent, continuous flow to the periphery occurs by distension of the aorta and its branches during ventricular contraction (*systole*) and elastic recoil of the walls of the large arteries with forward propulsion of the blood during ventricular relaxation (*diastole*).

Blood moves rapidly through the aorta and its arterial branches. The branches become narrower and their walls become thinner and change histologically toward the periphery. From a predominantly elastic structure, the aorta, the peripheral arteries become more muscular until at the arterioles the muscular layer predominates. Vascular smooth muscle allows individual tissues to control their own blood supply and also contributes to maintaining mean arterial pressure.

On its return to the heart from the capillaries, blood passes through venules and then through veins of increasing size with a progressive decrease in pressure until the blood reaches the right atrium. Blood entering the right ventricle via the right atrium is pumped through the pulmonary arterial system. The blood then passes through the lung capillaries where carbon dioxide is released and oxygen taken up. The oxygen-rich blood returns via the pulmonary veins to the left atrium and ventricle to complete the cycle. Thus in normal intact circulation the total volume of blood is constant. The velocity at which the blood circulates through the different regions of the body is determined by the output of the left ventricle and by the contractile state of the arterioles (resistance vessels) of these regions.

#### Blood vessels to the heart

Blood is supplied to the tissues of the heart by the first two branches of the aorta, called the *right* and *left coronary arteries*.

If a branch of a coronary artery becomes abnormally constricted or obstructed, the myocardium cell it supplies may experience a blood oxygen deficiency called *ischemia*. As a result of ischemia, a person may experience a painful condition called *angina pectoris*. This discomfort is more apparent during physical activity, when the myocardial cells need for oxygen exceeds the blood oxygen supply. Sometimes a portion of the heart dies because a blood clot forms and completely obstructs a coronary artery or one of its branches. This condition is called *myocardial infarction* (more commonly called a heart attack).

### B.1.4 Surface anatomy

The heart is enclosed by a double-layered *pericardial sac*, as shown in Figure B.4 (b). The outer layer of this loose-fitting sac is the tough *fibrous pericardium* and the inner layer is the *serous pericardium*. The *epicardium* (pericardium) is formed by the serous pericardium turning back on itself at the base of the heart. This is the thin, outer layer of the heart wall.

### B.1.5 Structure of the heart wall

Although the ventricular walls are much thicker than those of the atria, the layers of the cardiac wall are essentially similar in all four chambers. A ventricular wall is principally made up of muscle called *myocardium*. The myocardium is protected by an outer covering, called the *epicardium*, and by an protective inner lining, called the *endocardium*. Figures B.4 (a) and B.4 (b) illustrate this structure. Table B.2 (a) summarises the compositions and functions of these three distinct layers.

#### Epicardium

This is the visceral layer of serous pericardium, and is firmly tethered to the myocardium except where the vessels of the heart wall insinuate themselves and at sites of sub-epicardial fat decomposition.

#### Endocardium

The endocardium is thicker than the epicardium, and on its luminal surface is lined by a smooth endothelial layer.

#### Myocardium

The myocardium is relatively thick and consists largely of the cardiac muscle tissue responsible for forcing blood out of the heart chambers. The musculature contributes a near complete coat to the heart. During systole, as the heart propels blood from the left ventricle, the thickness of the myocardium increases — so called myocardial wall thickening. This is mainly perceived in ultrasound by the movement of the endocardium.

Layer	Composition	Function
Epicardium (visceral pericardium)	Serous membrane of connective tissue covered with epithelium and including blood capillaries, lymph capillaries and nerve fibers	Forms a protective outer covering
Myocardium	Cardiac muscle tissue separated by connective tissues and including blood capillaries, lymph capillaries and nerve fibers	Produces muscular contractions that force blood from the heart chambers
Endocardium	Membrane of epithelium and connective tissues, including elastic and collagenous fibers, blood vessels and specialized muscle fibers	Forms a protective inner lining of the chambers and valves

Table B.2: **The heart wall.** The heart wall consist of three layers: endocardium, myocardium and epicardium.




Figure B.4: **Layers of the heart.** (a) *Top:* The wall of the heart consists of three layers: endocardium, myocardium and epicardium. *Bottom:* The heart is enclosed by a layered pericardium. (b) Echocardiographic images showing myocardial thickening during systole. Note how the myocardium is clear in the top image, but not clear at all in the bottom image.

## B.2 Contraction and excitation of a cardiac myocyte

Each heart beat is initiated by an electrical system, called the pacemaker conduction system, which is situated within the wall of the heart. This system is composed of muscle cells, which form the bulk of the heart. When this electrical signal reaches these muscle cells, the cells are excited and fire off an action potential. This leads to a rise in intracellular calcium concentration, which in turn activates the contractile machinery of the cell.

The vast majority of cardiac muscle cells are designed to perform mechanical work, i.e. to contract and subsequently relax. These cells are called *myocytes*. These cells do not contract, however, unless stimulated electronically. A minority of the muscle cells are designed to initiate and conduct an electrical pulse in order to excite the myocytes. These cells form a specialized *cardiac electrical system*, the task of which is to initiate and co-ordinate the heart beat.

### B.2.1 The myocyte

The contractile cell, or (human) myocyte is typically  $10\text{--}20\mu\text{m}$  in diameter and  $50\text{--}100\mu\text{m}$  long, with a single central nucleus. The cell is branched, and is attached in an end-to-end fashion, as shown in Figure B.5.

The end-to-end junction, or interlaced disc, contains two smaller junctions. One of these called the *gap junction*, is thought to be an electronically-conductive region, through which ionic currents can pass from one cell to another. The gap ( $2\text{--}4\text{nm}$ ) is spanned by protein particles, called *connexons*, which have a central channel running through them. Ions can pass through these tiny channels, allowing transmission of impulses from cell to cell. As a result of these electrical transmissions, the myocardium acts as a electrically continuous sheet, and this enables excitation to reach every cell.

In patients with ischemia, these gap junctions are progressively uncoupled, and this is thought to contribute to poor conduction in ischaemic myocardium.




Figure B.5: **The Myocyte.** A section of myocardium parallel to the fibre axis. The width of the sarcomere ( $2\mu\text{m}$ ) and the red cell ( $7\mu\text{m}$ ) indicate the scale (Figure taken from (Levick, 1995, page 26)).

#### The myofibril and the sarcomere

The myocyte is packed with long contractile bundles called *myofibril-like units*. They are about  $1\mu\text{m}$  in thickness and are arranged linearly. Each myofibril consist of longitudinally repeating units, called *sarcomeres*, joined end-to-end. The sarcomere measure  $1.8\text{--}2\mu\text{m}$  in resting myocyte and are aligned across the cell. The sarcomere is defined as the material between two Z-lines (see Figure B.5), and is the basic contractile unit. The sarcomere contains two types of filament: a thick filament composed of *myosin* and thin filaments called *actin*. The myosin and actin filaments overlap each other. The myosin, of diameter,  $11\text{nm}$ , and length  $1.6\mu\text{m}$ , lie in parallel in a central region of the sarcomere called the A-band. The actin, of diameter,  $6\text{nm}$ , and length  $1.05\mu\text{m}$ , are rooted in the Z-line. A cell also contains non-contractile filaments, which contribute to its stiffness.

#### Sarcoplasmic reticulum

Within the cell, there is a system of tubular structures called the *sarcoplasmic reticulum*. The sarcoplasmic reticulum contains the store of  $\text{Ca}^{2+}$  ions. This store is released into the sarcoplasm following electrical excitation, and activates the contractile machinery there. We now look at how this contraction is produced.

## B.2.2 Contraction

Contraction of the myocyte is caused by shortening of sarcomeres. In fact the I-band shortens, but the A-band does not (see Figure B.5). This observation indicates that contraction is caused by the thin filaments of the I-band sliding into the spaces between the thick filaments of the A-band. This is called the *sliding mechanism*. The filaments are propelled past each other by the repeated creation and breaking of crossbridges between the thin and thick filaments. These crossbridges are actually the heads of myosin molecules, which protrude from the side of the thick filament, as illustrated in Figure B.6. Contraction is initiated by a sudden rise in the concentration of free intracellular calcium ions. The force and movement is produced by a subsequent change in the angle of the crossbridge (i.e. the attached myosin head), after which the head disengages and the process repeats itself at a new actin site. This process occurs at numerous similar sites along the filament, and in this way the thick filament 'rows' itself into the space between the thin filaments. The force of contraction depends directly on the concentration of free calcium ions within the myocyte and therefore, the number of X-bridges formed. In systole, there is a 50-fold increase in intracellular calcium concentration.




Figure B.6: **The crossbridge cycle.** (a) The resting state. The myosin head is disengaged. (b) The myosin head cross-links to the exposed actin binding site. (c) The angle of the myosin head changes, 'rowing' the thick filament towards the Z line. The head then disengages, and the cycle repeats at a new actin site. Only one myosin is shown here, but the thick filament contains approximately 400. Similarly, there are many binding sites along the actin filament (Figure taken from (Levick, 1995, page 29)).

## B.2.3 The mechanics of the myocardium

Two important characteristics of heart muscle are the force with which it contracts, and the velocity with which it shortens. These determine the blood expelled during systole. Three factors affect the force and velocity of cardiac contraction:

- Preload — if all other factors are held constant, then the extent to which the heart contracts depends on the extent to which it has been stretched prior to contraction, as illustrated in Figure B.7. i.e. venous return and ventricular filling.
- Afterload – the load that the ventricle faces during contraction.
- The contractile state of the myocardium — if the end-diastolic volume and aortic impedance are held constant, the force and velocity of contraction of the myocardium depend upon its contractile state, as shown in Figure B.7.




Figure B.7: **Fibre length.** The relationship between end-diastolic fibre length and the left ventricular stroke work. For any given fibre length, increased contractility produces more stroke work. Likewise, for any given amount of stroke work, increased contractility enables this to be performed at a smaller end-diastolic length.

# APPENDIX C

---

## Basic principles of ultrasound

---

Although the focus of our work is not image formation, the construction of a robust segmentation and tracking scheme requires a good understanding of ultrasonic data. For completeness, in this appendix we review some of the fundamentals of ultrasound image formation and acquisition. A more thorough account is given in Webb (1988) and Kremkau (1993). In practice, conventional echocardiography typically exhibits sound in the range 2–7MHz. By comparison, the audible range for the human ear is 20–20,000Hz.

### C.1 Properties of Ultrasound

As Feigenbaum (1994) points out, ultrasound has a number of advantages as a diagnostic tool; ultrasound can be directed in a beam; it obeys the laws of reflection and refraction; and it is reflected by small objects. The principal disadvantage of ultrasound is that it propagates poorly through a gaseous medium. The amount of ultrasound reflected depends on the acoustic mismatch. Therefore, when the beam travels through tissue containing gases and solids, almost all of the sound is reflected and penetration is poor. For example, a dense substance, such as bone, calcium or metal, will reflect almost all the energy. Thus, the ribs and lungs (which are situated next to the heart) offer distinct imaging challenges for the echocardiographer, and can affect the quality of the images. This is, of course, patient dependent. In practice, because the transducer should have airless contact with the body during the examination, an aqueous gel or mineral oil is used as a coupling medium. Moreover, in general, the patient is imaged while lying on their (left) side. This has the result of moving the lungs slightly, so that the heart becomes more visible.

The velocity,  $v$ , at which sound waves travel through a particular medium is, to a first approximation constant, and is equal to the product of frequency,  $f$ , and wavelength,  $\lambda$ ,

$$v = f\lambda.$$

Clearly, as  $v$  is constant, the higher the frequency, the shorter the wavelength. Higher frequency ultrasound with short cycle lengths is most desirable for echocardiographic images, allowing reflection from small targets, which enhances resolution. However, high frequency ultrasound waves are rapidly attenuated resulting in lower penetration.

The velocity at which sound travels through a medium depends on the density and stiffness properties of that medium. In general, sound travels slower through a less dense substance, such as a gas, than it does through a more dense substance, such as liquid, than it does through an even more dense substance, such as solid. This increasing sequence is not caused by the

increasing density (this reduces the velocity), but by the increasing stiffness; the differences in stiffness outweighs those of the differences in density.

The average velocity of sound is fairly constant for human tissue, around 1540m/s. The difference in the velocity is significant if the sound passes through a solid structure such as bone. Although velocity also depends on temperature, the small change in body temperature are not important in echocardiography.

### C.1.1 Acoustic impedance

The acoustic impedance,  $Z$ , of a medium characterises how sound travels through that medium. It is defined to be the density of the medium,  $\rho$ , times the velocity that sound travels through that medium.

$$z = \rho v$$

Impedance is determined by the density and stiffness of the medium, but not frequency; impedance increases if the density or stiffness increases.

Diagnostic ultrasound is based on the principle that ultrasound is reflected (and refracted) at an interface between two media of two different impedances. The greater this acoustic mismatch, the greater the amount of sound reflected. For example, more sound is reflected from an interface between gaseous and solid media than from a liquid-solid interface. For instance, when a pulse crosses the boundary between muscle tissue ( $Z = 1.7 \times 10^6$  rayls) and blood ( $Z = 1.6 \times 10^6$  rayls) about 0.1% of the acoustic energy is reflected (Karrer & Dickey, 1983).

Reflected echoes, can be classed in to two types, specular echoes and scattered echoes. Specular echoes are produced by objects that are fairly large with respect to the wavelength and that present a relatively smooth surface to the ultrasound beam. The reflection can be predicted by the spatial orientation and shape of reflecting object or interface. In contrast, scattered echoes originate from relatively small objects with irregular surfaces. The resulting echoes are reflected in multiple directions, and only a small percentage of the ultrasound energy returns to the transducer. This means that even if the ultrasound beam is incorrectly targeted, and no specular echoes are returned, energy will return to the transducer from the scattered echoes. Scattered echoes, however, play an important role in helping to visualize objects which are parallel to the ultrasound beam, for example, the lateral or anterior walls (see Figure 6.4, page 85) of the left ventricle (Kremkau, 1993, pages 49-51).

### C.1.2 Attenuation

Attenuation describes the loss of ultrasound energy as it traverses a medium, a combination of absorption (sound to heat), reflection and scattering. Absorption is normally the dominant contribution to attention in soft tissue. The longer the path over which sound travels, the greater the attenuation. All media attenuate ultrasound, so that the intensity of a plane wave,  $I(x)$ , propagating in the  $x$  direction decreases exponentially with the distance,

$$I(x) = I_0 e^{-\mu x},$$

where,  $\mu$  is the intensity attenuating coefficient. If,  $\mu_a$  is the intensity absorption coefficient, and  $\mu_s$  the intensity scattering coefficient, then,  $\mu = \mu_a + \mu_s$ . The relative contributions are, however,, not known for many tissues. Many soft tissues of the body attenuate ultrasound to a similar degree, and display a nearly linear frequency dependence. The attention of ultrasound increases with frequency. This is why it is mostly the base that can be heard when your neighbour is playing loud music!

Attenuation is higher in lung than in other soft tissues, and it is higher in bone than in soft tissues. Lung and bone attenuations are not proportionally dependent on frequency. In practice this provides a practical constraint on imaging depth; imaging depth decreases as frequency increases because attenuation increases.

### C.1.3 Scattering

The usefulness of ultrasound as an imaging tool is primarily the result of reflection and scattering at organ and tissue boundaries and scattering within heterogeneous tissues.

Structures within tissues, which may potentially scatter ultrasound range over at least four orders of magnitude in size, from cells (at about  $10\mu\text{m}$ , or  $0.03\lambda$  at 5MHz) to organ boundaries (up to 10cm, or  $300\lambda$  at 5MHz). Different kinds of scattering phenomena occur at different levels of structure. Structures within tissues, which may potentially scatter ultrasound range over at least four orders of magnitude in size, from cells (at about  $10\mu\text{m}$ , or  $0.03\lambda$  at 5MHz) to organ boundaries (up to 10cm, or  $300\lambda$  at 5MHz). Different kinds of scattering phenomena occur at different levels of structure.

At a large-scale boundary, representing the interface between two homogeneous media, the visual law of reflection, and Snell's law for refraction can be applied to predict the direction of the reflected and refracted sound.

$$\frac{\sin(\theta_i)}{\sin(\theta_t)} = \frac{c_1}{c_2} \quad \theta_i = \theta_r,$$

where,  $c_1$  is the speed in of sound in medium 1 and  $c_2$  is the speed in of sound in medium 2.  $\theta_i$  is the angle of the incident ray,  $\theta_t$  the angle of the refracted ray and  $\theta_r$  the angle of the refracted ray.

The intensity of the reflected sound beam, relative to the incident intensity, is given by the intensity reflection coefficient,  $R_i$ ,

$$R_i = \left( \frac{Z_2 \cos \theta_i - Z_1 \cos \theta_t}{Z_2 \cos \theta_i + Z_1 \cos \theta_t} \right)^2,$$

where,  $Z_1$  is the acoustic impedance of medium 1 and  $Z_2$  is the acoustic impedance of medium 2.

For perpendicular incidence, if the media are the same, there is no reflected sound, and the refracted intensity is equal to the incident intensity:

$$R_i = \frac{\sqrt{K_2} - \sqrt{K_1}}{(\sqrt{K_2} + \sqrt{K_1})^2},$$

where,  $K_1$  is the adiabatic bulk elastic modulus.

Interfaces between media that are separated by the greatest difference in speed of sound provide the largest reflection coefficient, and therefore echo imaging is a technique with provides excellent contrast for depicting boundaries between media with different speeds. This is why lung attenuates the ultrasound so much, and why it is difficult for ultrasound to visualize beyond bone or gas. It also explains why a coupling gel is used to provide a good sound path from the transducer to the skin (eliminating the thin layer of air).

## C.2 Image formation

The transducer is the key component in any medical ultrasound system. Ultrasound transducers produce and receive ultrasound energy using a piezoelectric crystal (Wells, 1993). Piezoelectric crystals produce sound waves by rapidly changing their shape when subjected to an alternating current.

The ultrasound beam consists primarily of a series of longitudinal waves moving parallel to the direction of propagation of the sound waves. The width of the ultrasound beam is uniform closest to the transducer. Resolution is less in the far field. The actual ultrasound beam shape strictly speaking is three-dimensional. The shape is affected by the number of crystal elements in the transducer (Feigenbaum, 1994, page 7). In practice there is a range of ultrasound transducers of different size, shape and frequency to deal with the body habitus of the patient. With a young child or a thin adult, preference is for a higher frequency (3.5MHz or even 7MHz) transducer.

Attenuation is more pronounced with high frequency ultrasound signal, indicating that lower frequency transducers provide better imaging in large patients. Hence, a thick chested individual would be imaged using a 2 to 2.5MHz transducer is necessary for adequate penetration. In the case of a new born child, the choice would be a 7 or 7.5MHz transducer.

In general, larger transducers have good resolution because the larger the size of the transducer, the longer the near field and lesser the degree of beam divergence in the far field. Detrimentally, large transducers are more difficult to manipulate, and images may be poorer for patients with relatively narrow intercostal spaces and whose ribs are calcified (i.e limited access windows).

When the ultrasound energy returns as an echo signal and strikes the piezoelectric element of the transducer, an electrical impulse is created and transmitted back in the form of radio frequency (RF) to the ultrasound unit. In two-dimensions, digital manipulation consists of digital scan converters to reconstruct and to display cardiac images. The scan conversion may degrade the ultrasonic image if there is inadequate density of pixels or number of grey levels, and some of the processed information is repeated to fill in the missing gaps (Leavitt *et al.*, 1983).

# APPENDIX D

---

## Miscellaneous Matrix Results

---

### D.1 Matrix Operators

For completeness, some matrix operations used in this thesis are stated without proof below (Rogers, 1980; Magnus & Neudecker, 1988).

#### D.1.1 The Kronecker product

In Section 4.1, page 33, the Kronecker product is used to map a matrix,  $\mathbf{U}(s)$ , to a control point vector  $\mathbf{Q}$  to image-curve  $\mathbf{r}(s)$ .

Let  $\mathbf{A}$  be a  $m \times n$  matrix and  $\mathbf{B}$  be an  $p \times q$  matrix. The Kronecker product of  $\mathbf{A}$  and  $\mathbf{B}$ , denoted,  $\mathbf{A} \otimes \mathbf{B}$ , is defined as the  $mp \times nq$  matrix,

$$\begin{pmatrix} a_{11}\mathbf{B} & \cdots & a_{1n}\mathbf{B} \\ \vdots & & \vdots \\ a_{m1}\mathbf{B} & \cdots & a_{mn}\mathbf{B} \end{pmatrix}$$

#### D.1.2 The Moore-Penrose (MP) Inverse

An  $n \times m$  matrix  $\mathbf{A}^\dagger$  is the MP inverse of a real  $m \times n$  matrix  $\mathbf{A}$  if,

$$\begin{aligned} \mathbf{A}\mathbf{A}^\dagger\mathbf{A} &= \mathbf{A}, \\ \mathbf{A}^\dagger\mathbf{A}\mathbf{A}^\dagger &= \mathbf{A}^\dagger, \\ (\mathbf{A}\mathbf{A}^\dagger)^\top &= \mathbf{A}\mathbf{A}^\dagger, \\ (\mathbf{A}^\dagger\mathbf{A})^\top &= \mathbf{A}^\dagger\mathbf{A}. \end{aligned}$$

### D.2 Matrix Differentiation

Some proofs concerning mainly vector and matrix differentiation are stated below without proof. The results are used in Appendix E.2.2 to differentiate the likelihood function.

### D.2.1 Definition of Matrix differentiation

If  $f(x)$  is a real function of an  $m \times n$  matrix  $\mathbf{X} = (x_{ij})$ , then the partial differential of  $f$  with respect to  $\mathbf{X}$  is defined as the  $m \times n$  matrix of partial differentials  $\partial f / \partial x_{ij}$ :

$$\frac{\partial f(\mathbf{X})}{\partial \mathbf{X}} = \begin{pmatrix} \frac{\partial f}{\partial x_{11}} & \cdots & \frac{\partial f}{\partial x_{1n}} \\ \vdots & & \vdots \\ \frac{\partial f}{\partial x_{m1}} & \cdots & \frac{\partial f}{\partial x_{mn}} \end{pmatrix}$$

If  $\mathbf{x}$  is an  $n$ -vector and  $\mathbf{A}$  be a symmetric  $n \times n$  matrix. Then,

$$\frac{\partial}{\partial \mathbf{x}} \mathbf{x}^\top \mathbf{A} \mathbf{x} = 2\mathbf{A} \mathbf{x}. \quad (\text{D.1})$$

If  $\mathbf{x}$  is a  $p$ -vector,  $\mathbf{A}$  a symmetric  $n \times n$  matrix, and  $\mathbf{C}$  a  $n \times p$  matrix. Then,

$$\frac{\partial}{\partial \mathbf{C}} \mathbf{x}^\top \mathbf{C}^\top \mathbf{A} \mathbf{C} \mathbf{x} = 2\mathbf{A} \mathbf{C} \mathbf{x} \mathbf{x}^\top. \quad (\text{D.2})$$

If  $\mathbf{X}^\top$ ,  $\mathbf{A}$  and  $\mathbf{Y}$  are 3 matrices. Then,

$$\frac{\partial}{\partial \mathbf{X}} \mathbf{X}^\top \mathbf{A} \mathbf{Y} = \mathbf{A} \mathbf{Y}^\top. \quad (\text{D.3})$$

If  $\mathbf{X}$  is a  $k \times k$  symmetric nonsingular matrix of independent real variables (except  $x_{ij} = x_{ji}$ ). Then,

$$\frac{\partial}{\partial \mathbf{X}} \log |\mathbf{X}| = 2\mathbf{X}^{-1} - \text{Diag} \mathbf{X}^{-1}, \quad (\text{D.4})$$

where,  $\text{Diag} \mathbf{X}^{-1}$  is a diagonal matrix with  $i$ -th diagonal element equal to the  $i$ -th diagonal element of  $\mathbf{X}^{-1}$ .

If  $\mathbf{a}$  is a  $k \times 1$  vector of constants, and  $\mathbf{X}$  a  $k \times k$  symmetric matrix of independent real variables (except  $x_{ij} = x_{ji}$ ). Then,

$$\frac{\partial}{\partial \mathbf{X}} \mathbf{a} \mathbf{X} \mathbf{a}^\top = 2\mathbf{a} \mathbf{a}^\top - \text{Diag}(\mathbf{a} \mathbf{a}^\top), \quad (\text{D.5})$$

where,  $\text{Diag}(\mathbf{a} \mathbf{a}^\top)$  is a  $k \times k$  matrix whose  $i$ -th diagonal element is equal to the diagonal of the matrix  $\mathbf{a} \mathbf{a}^\top$ .

# APPENDIX E

---

## Derivations and proofs

---

### E.1 Proof that $\mathbf{W}^\dagger = (\mathbf{W}^\top \mathcal{H} \mathbf{W})^{-1} \mathbf{W}^\top \mathcal{H}$

This proof is relevant to the theory in Section 4.3.1, page 36. A shape-space is defined by,

$$\mathbf{Q} = \mathbf{W}\mathbf{X} + \mathbf{Q}_0. \quad (\text{E.1})$$

We wish to find a matrix,  $\mathbf{W}^\dagger$ , that transforms a control vector,  $\mathbf{Q}$ , in control-point space,  $\mathcal{S}_Q$  into a shape-space vector,  $\mathbf{X}$ , in shape-space  $\mathcal{S}$ . More formally, we seek  $\mathbf{W}^\dagger$ , such that,

$$\mathbf{X} = \mathbf{W}^\dagger(\mathbf{Q} - \mathbf{Q}_0). \quad (\text{E.2})$$

The projection of  $\mathbf{Q}$  onto each of the basis vectors of  $\mathbf{W}$  (the columns of which are  $\mathbf{w}_i$ ) is given by the inner product  $\langle \mathbf{w}_i, \mathbf{Q} \rangle$ .

$$\begin{aligned} \langle \mathbf{w}_i, \mathbf{Q} \rangle &= \int_{s=0}^L \mathbf{w}_i^\top \begin{pmatrix} \mathbf{B}(s)^\top & \mathbf{0} \\ \mathbf{0} & \mathbf{B}(s)^\top \end{pmatrix} \mathbf{Q} ds \\ &= \mathbf{w}_i^\top \mathcal{H} \mathbf{Q}, \end{aligned} \quad (\text{E.3})$$

where,

$$\mathcal{H} = \begin{pmatrix} 0 & 0 \\ 0 & 0 \end{pmatrix} \otimes \frac{1}{L} \int_{s=0}^L \mathbf{B}(s) \mathbf{B}(s)^\top ds. \quad (\text{E.4})$$

Here  $L$  is the number of spans on B-spline curve, and  $\mathbf{B}(s)$  is a vector of B-spline blending functions parametrized by  $s$

The projection of  $\mathbf{Q}$  onto the basis vectors in  $\mathbf{W}$  can be written as  $\mathbf{W}^\top \mathcal{H} \mathbf{Q}$ .

$\mathbf{X}$  is not directly given by this projection — the columns of  $\mathbf{W}$  are not necessarily orthogonal to each other. However, for some transformation,  $\mathbf{T}$ ,

$$\mathbf{X} = \mathbf{T} \mathbf{W}^\top \mathcal{H} (\mathbf{Q} - \mathbf{Q}_0). \quad (\text{E.5})$$

Finding this  $\mathbf{T}$  will allow us to find  $\mathbf{W}^\dagger$ .

Now for any  $\mathbf{X}$  in  $\mathcal{S}$ , transforming it up into the control-point space,  $\mathcal{S}_Q$ , and back again, should

leave it unchanged. Hence,

$$\begin{aligned} \mathbf{Q} &= (\mathbf{T}\mathbf{W}^\top \mathcal{H})\mathbf{W}\mathbf{Q}, \forall \mathbf{Q} \in \mathcal{S}_Q, \\ \therefore \mathbf{T}\mathbf{W}^\top \mathcal{H}\mathbf{W} &= \mathbf{I}, \text{ (identity)}, \\ \therefore \mathbf{T} &= (\mathbf{W}^\top \mathcal{H}\mathbf{W})^{-1}. \end{aligned} \quad (\text{E.6})$$

Therefore,

$$\mathbf{W}^\dagger = (\mathbf{W}^\top \mathcal{H}\mathbf{W})^{-1} \mathbf{W}^\top \mathcal{H}. \quad (\text{E.7})$$

□

## E.2 Maximum Likelihood Estimation for parameters $\mathbf{A}_0$ , $\mathbf{A}_1$ and $\mathbf{B}$

In Section 4.5.1, page 44, we assumed a second order autoregressive process for the dynamical tracking model. This takes the form,

$$\mathbf{X}_t - \bar{\mathbf{X}} = \mathbf{A}_1(\mathbf{X}_{t-1} - \bar{\mathbf{X}}) + \mathbf{A}_2(\mathbf{X}_{t-2} - \bar{\mathbf{X}}) + \mathbf{B}_0 \mathbf{w}_t \quad \text{and} \quad \mathbf{w}_t \sim \text{i.i.d } \mathcal{N}(\mathbf{0}, \mathbf{I}). \quad (\text{E.8})$$

Here  $\mathbf{A}_0$ ,  $\mathbf{A}_1$  and  $\mathbf{B}$  are matrices, and  $\mathbf{w}(n)$  is Gaussian noise.  $\mathbf{B}$  is the size of the variation that we observe (tracking error). The univariate equivalent scalars are  $a_0$ ,  $a_1$  and  $b$ .

The discrete-time system parameters are estimated via Maximum Likelihood estimation (MLE). Assuming that the noise is Gaussian, it is straightforward to set up and minimise the likelihood function. We assume that the mean configuration is  $\mathbf{0}$ .

### E.2.1 Univariate case

Since the noise is assumed to be distributed  $\mathcal{N}(0, b^2)$ , the probability density function takes the form,

$$f(x|a_0, a_1, b) = \frac{1}{\sqrt{2\pi} b} \exp \left\{ -\frac{1}{2b^2} (x_t - a_1 x_{t-1} - a_2 x_{t-2})^2 \right\}. \quad (\text{E.9})$$

The corresponding likelihood function is then the product of the individual probability densities,

$$L(\mathbf{x}|a_0, a_1, b) = \prod_{t=3}^M \frac{1}{\sqrt{2\pi} b} \exp \left\{ -\frac{1}{2b^2} (x_t - a_1 x_{t-1} - a_2 x_{t-2})^2 \right\},$$

where,  $M$ , is the number of frames in the training sequence.

Maximising the likelihood function is entirely equivalent to maximising the log-likelihood function, which takes the form,

$$\ln L(\mathbf{x}|a_0, a_1, b) = l(\mathbf{x}|a_0, a_1, b) = \sum_{t=3}^M -\ln \sqrt{2\pi} - \ln b - \frac{1}{2b^2} (x_t - a_2 x_{t-2} - a_1 x_{t-1})^2, \quad (\text{E.10})$$

$$l(\mathbf{x}|a_0, a_1, b) \propto -\sum_{t=3}^M \frac{1}{2b^2} (x_t - a_1 x_{t-1} - a_2 x_{t-2})^2 - (m-2) \ln b. \quad (\text{E.11})$$

The MLE for the coefficients  $a_0, a_1, b$  is obtained by maximising equation E.11 with respect to

each parameter in turn.

$$\begin{aligned}
 \frac{\partial l}{\partial a_0} &= -\frac{1}{2b^2} \sum_{t=3}^M (x_t - a_1 x_{t-1} - a_2 x_{t-2}) x_{t-2} \\
 &= 0 \text{ when,} \\
 \sum_{t=3}^M (x_t - \hat{a}_1 x_{t-1} - \hat{a}_0 x_{t-2}) x_{t-2} &= 0, \\
 \sum_{t=3}^M x_t x_{t-2} &= \hat{a}_0 \sum_{t=3}^M x_{t-2} x_{t-2} + \hat{a}_1 \sum_{t=3}^M x_{t-1} x_{t-2}.
 \end{aligned} \tag{E.12}$$

Now if we let,

$$s_{ij} = \sum_{t=3}^M x_{t+i} x_{t+j} \quad i, j = 0, 1, 2, \quad s_{ij} = s_{ji}, \tag{E.13}$$

then,

$$s_{20} = \hat{a}_0 s_{00} + \hat{a}_1 s_{10}. \tag{E.14}$$

Differentiating with respect to  $a_1$ ,

$$\begin{aligned}
 \frac{\partial l}{\partial a_1} &= -\frac{1}{2b^2} \sum_{t=3}^M (x_t - a_1 x_{t-1} - a_2 x_{t-2}) x_{t-1} \\
 &= 0 \text{ when,} \\
 \sum_{t=3}^M (x_t - \hat{a}_1 x_{t-1} - \hat{a}_0 x_{t-2}) x_{t-1} &= 0, \\
 \sum_{t=3}^M x_t x_{t-2} &= \hat{a}_0 \sum_{t=3}^M x_{t-2} x_{t-1} + \hat{a}_1 \sum_{t=3}^M x_{t-2} x_{t-1}.
 \end{aligned} \tag{E.15}$$

Using Equation E.13 this becomes,

$$s_{21} = \hat{a}_0 s_{01} + \hat{a}_1 s_{11}. \tag{E.16}$$

We now simultaneously solve Equations E.14 and E.16 to obtain the MLE of the parameters  $a_0$  and  $a_1$ .

The MLE estimate for  $b$  is obtained by maximising the log-likelihood function (equation E.11) with respect to  $b$ , holding  $a_0$  and  $a_1$  constant:

$$\begin{aligned}
 \frac{\partial l}{\partial b} &= -\frac{M}{b} + \frac{1}{b^3} \sum_{t=3}^M (x_t - a_1 x_{t-1} - a_2 x_{t-2})^2 \\
 &= 0 \text{ when,} \\
 \hat{b}^2 &= \frac{1}{M} \sum_{t=3}^M (x_t - \hat{a}_1 x_{t-1} - \hat{a}_0 x_{t-2})^2
 \end{aligned} \tag{E.17}$$

## E.2.2 Multivariate case

Multivariate estimation follows analogously. The noise is assumed to be distributed  $N(0, \mathbf{B})$ . Letting  $\mathbf{B}^{-1} = \mathbf{V}$ , the probability density function takes the form,

$$f(\mathbf{x}|\mathbf{A}_1, \mathbf{A}_2, \mathbf{B}) = |2\pi|^{-\frac{1}{2}} |\mathbf{V}|^{\frac{1}{2}} \exp\left\{-\frac{1}{2}(\mathbf{X}_t - \mathbf{A}_1 \mathbf{X}_{t-1} - \mathbf{A}_2 \mathbf{X}_{t-2})^\top \mathbf{V}(\mathbf{X}_t - \mathbf{A}_1 \mathbf{X}_{t-1} - \mathbf{A}_2 \mathbf{X}_{t-2})\right\}. \quad (\text{E.18})$$

The likelihood function has the form,

$$L(\mathbf{x}|\mathbf{A}_1, \mathbf{A}_2, \mathbf{B}) = |2\pi|^{-\frac{M}{2}} |\mathbf{V}|^{\frac{M}{2}} \exp\left\{-\frac{1}{2} \sum_{t=3}^M (\mathbf{X}_t - \mathbf{A}_1 \mathbf{X}_{t-1} - \mathbf{A}_2 \mathbf{X}_{t-2})^\top \mathbf{V}(\mathbf{X}_t - \mathbf{A}_1 \mathbf{X}_{t-1} - \mathbf{A}_2 \mathbf{X}_{t-2})\right\}. \quad (\text{E.19})$$

The MLE estimates for the parameters  $\mathbf{A}_1$ ,  $\mathbf{A}_2$  and  $\mathbf{B}$  are obtained by maximising this likelihood or the corresponding log-likelihood function, given by,

$$\begin{aligned} \ln L(\mathbf{x}|\mathbf{A}_1, \mathbf{A}_2, \mathbf{B}) &= l(\mathbf{x}|\mathbf{A}_1, \mathbf{A}_2, \mathbf{B}) \\ &= -\frac{M}{2} \ln|2\pi| + \frac{M}{2} \ln|\mathbf{V}| \\ &\quad - \frac{1}{2} \sum_{t=3}^M (\mathbf{X}_t - \mathbf{A}_2 \mathbf{X}_{t-2} - \mathbf{A}_1 \mathbf{X}_{t-1})^\top \mathbf{V}(\mathbf{X}_t - \mathbf{A}_2 \mathbf{X}_{t-2} - \mathbf{A}_1 \mathbf{X}_{t-1}). \end{aligned} \quad (\text{E.20})$$

Then using Equation D.2,

$$\begin{aligned} \frac{\partial l}{\partial \mathbf{A}_2} &= -\frac{1}{2} \sum_{t=3}^M \mathbf{V}(\mathbf{X}_t - \mathbf{A}_1 \mathbf{X}_{t-1} - \mathbf{A}_2 \mathbf{X}_{t-2}) \mathbf{X}_{t-2}^\top \\ &= 0 \text{ when,} \\ &\quad \sum_{t=3}^M (\mathbf{X}_t - \mathbf{A}_1 \mathbf{X}_{t-1} - \mathbf{A}_2 \mathbf{X}_{t-2}) \mathbf{X}_{t-2}^\top = 0 \quad (\text{pre-multiplying by } \mathbf{V}^{-1}). \end{aligned} \quad (\text{E.21})$$

Now analogously with Equation E.13, we can define  $\mathbf{S}$  as,

$$\mathbf{S}_{ij} = \sum_{t=3}^M \mathbf{X}_{t+i} \mathbf{X}_{t+j}^\top \quad i, j = 0, 1, 2, \quad \mathbf{S}_{ij}^\top = \mathbf{S}_{ji}. \quad (\text{E.22})$$

Therefore,

$$\mathbf{S}_{20} - \hat{\mathbf{A}}_0 \mathbf{S}_{00} - \hat{\mathbf{A}}_1 \mathbf{S}_{10} = 0. \quad (\text{E.23})$$

$$\begin{aligned} \frac{\partial l}{\partial \mathbf{A}_1} &= -\frac{1}{2} \sum_{t=3}^M \mathbf{V}(\mathbf{X}_t - \mathbf{A}_1 \mathbf{X}_{t-1} - \mathbf{A}_2 \mathbf{X}_{t-2}) \mathbf{X}_{t-1}^\top \\ &= 0 \text{ when,} \\ &\quad \sum_{t=3}^M (\mathbf{X}_t - \mathbf{A}_1 \mathbf{X}_{t-1} - \mathbf{A}_2 \mathbf{X}_{t-2}) \mathbf{X}_{t-1}^\top = 0 \quad (\text{pre-multiplying by } \mathbf{V}^{-1}). \end{aligned} \quad (\text{E.24})$$

Therefore,

$$\mathbf{S}_{21} - \hat{\mathbf{A}}_0 \mathbf{S}_{01} - \hat{\mathbf{A}}_1 \mathbf{S}_{11} = 0. \quad (\text{E.25})$$

We can now simultaneously solve Equations E.23 and E.25 to obtain the MLE of the parame-

ters  $\mathbf{A}_2$  and  $\mathbf{A}_1$ .

$$\hat{\mathbf{A}}_0 = (\mathbf{S}_{20} - \mathbf{S}_{21}\mathbf{S}_{11}^{-1}\mathbf{S}_{10})(\mathbf{S}_{00} - \mathbf{S}_{01}\mathbf{S}_{11}^{-1}\mathbf{S}_{10})^{-1}, \quad (\text{E.26})$$

$$\hat{\mathbf{A}}_1 = (\mathbf{S}_{20}\mathbf{S}_{01}^{-1}\mathbf{S}_{10} - \mathbf{S}_{20})(\mathbf{S}_{10} - \mathbf{S}_{11}\mathbf{S}_{01}^{-1}\mathbf{S}_{10})^{-1}. \quad (\text{E.27})$$

The MLE estimate for  $\mathbf{V}$  is obtained by maximising the log-likelihood function (Equation E.20) with respect to  $\mathbf{V}$ , holding  $\mathbf{A}_2$  and  $\mathbf{A}_1$  constant.

Using Equations D.4 and D.5,

$$\begin{aligned} \frac{\partial l}{\partial \mathbf{V}} &= \frac{M}{2} \left( 2\mathbf{V}^{-1} - \text{Diag}\mathbf{V}^{-1} \right) - \frac{1}{2} \left\{ \sum_{t=3}^M \left[ 2(\mathbf{X}_t - \mathbf{A}_1\mathbf{X}_{t-1} - \mathbf{A}_2\mathbf{X}_{t-2})(\mathbf{X}_t - \mathbf{A}_1\mathbf{X}_{t-1} - \mathbf{A}_2\mathbf{X}_{t-2})^\top \right. \right. \\ &\quad \left. \left. - \text{Diag} \left( (\mathbf{X}_t - \mathbf{A}_1\mathbf{X}_{t-1} - \mathbf{A}_2\mathbf{X}_{t-2})(\mathbf{X}_t - \mathbf{A}_1\mathbf{X}_{t-1} - \mathbf{A}_2\mathbf{X}_{t-2})^\top \right) \right] \right\} \\ &= \frac{M}{2} \left( 2\mathbf{V}^{-1} - \frac{1}{M} \sum_{t=3}^M (\mathbf{X}_t - \mathbf{A}_1\mathbf{X}_{t-1} - \mathbf{A}_2\mathbf{X}_{t-2})(\mathbf{X}_t - \mathbf{A}_1\mathbf{X}_{t-1} - \mathbf{A}_2\mathbf{X}_{t-2})^\top \right) \\ &\quad - \frac{M}{2} \left\{ \text{Diag}\mathbf{V}^{-1} - \frac{1}{M} \text{Diag} \left( \sum_{t=3}^M (\mathbf{X}_t - \mathbf{A}_1\mathbf{X}_{t-1} - \mathbf{A}_2\mathbf{X}_{t-2})(\mathbf{X}_t - \mathbf{A}_1\mathbf{X}_{t-1} - \mathbf{A}_2\mathbf{X}_{t-2})^\top \right) \right\} \\ &= \frac{M}{2} \left( 2\mathbf{M} - \text{Diag}\mathbf{M} \right), \end{aligned} \quad (\text{E.28})$$

where,  $\mathbf{M} = \mathbf{V}^{-1} - \frac{1}{M} \sum_{t=3}^M (\mathbf{X}_t - \mathbf{A}_2\mathbf{X}_{t-2} - \mathbf{A}_1\mathbf{X}_{t-1})(\mathbf{X}_t - \mathbf{A}_2\mathbf{X}_{t-2} - \mathbf{A}_1\mathbf{X}_{t-1})^\top$ .

The MLE of  $\mathbf{V}$  is obtained when  $\frac{\partial l}{\partial \mathbf{V}} = 0$ , which gives that  $2\mathbf{M} = \text{Diag}\mathbf{M}$ . This implies that  $\mathbf{M} = 0$ , and that the MLE of  $\mathbf{V}$  is given by,

$$\mathbf{V}^{-1} = \frac{1}{M} \sum_{t=3}^M (\mathbf{X}_t - \mathbf{A}_1\mathbf{X}_{t-1} - \mathbf{A}_2\mathbf{X}_{t-2})(\mathbf{X}_t - \mathbf{A}_1\mathbf{X}_{t-1} - \mathbf{A}_2\mathbf{X}_{t-2})^\top. \quad (\text{E.29})$$

Recall that we let  $\mathbf{B} = \mathbf{V}^{-1}$ , so the MLE of  $\mathbf{B}$  can be obtained.

### E.3 Myocardial thickening

This proof is relevant to the theory in Section 7.7 page 127. It enables the segmental distance between the epicardium and endocardium contours to be calculated using the difference shape-space (Section 7.4),  $\mathbf{W}^{\text{Diff}}$ , and the associated shape-space vector,  $\mathbf{X}^{\text{Diff}}$ .

Let  $\mathbf{Q}_{\text{En}}$  be the endocardium contour and  $\mathbf{Q}_{\text{Ep}}$  be the epicardium contour. The associated curves are given by  $r_{\text{En}}(s)$  and  $r_{\text{Ep}}(s)$ , respectively. Let  $w(s)$  be a weight function that selects segments of along boundary. Then the regional myocardial thickening, is given by,

$$\int_{s_1}^{s_2} (r_{\text{Ep}}(s) - r_{\text{En}}(s))^2 w(s) ds, \quad (\text{E.30})$$

where,  $s_1$  and  $s_2$  are the start and end points of the region of interest. For example, in the context of a 16-segment model of the heart (Section 6.2),  $s_1$  and  $s_2$  could represent the start and end of the basal inferior segment, for instance. Here,  $w(s)$  is a weighting function which ensures that only selected regions are included.

From Equation 4.4, page 33,  $r(s) = \mathbf{U}(s)\mathbf{Q}$ . Then,

$$r_{\text{Diff}}(s) = r_{\text{Ep}}(s) - r_{\text{En}}(s) \quad (\text{E.31})$$

$$= \mathbf{U}\mathbf{Q}_{\text{Diff}} \quad (\text{E.32})$$

$$= \mathbf{U}(s)\mathbf{W}_{\text{Diff}}\mathbf{X}_{\text{Diff}}. \quad (\text{E.33})$$

Therefore,

$$\frac{1}{s_2 - s_1} \int_{s_1}^{s_2} r_{\text{Diff}}^2(s) ds = \int_{s_1}^{s_2} \mathbf{X}_{\text{Diff}}^\top \mathbf{W}_{\text{Diff}}^\top \mathbf{U}(s)^\top \mathbf{U}(s) \mathbf{W}_{\text{Diff}} \mathbf{X}_{\text{Diff}} ds \quad (\text{E.34})$$

$$= \mathbf{X}_{\text{Diff}}^\top \mathbf{W}_{\text{Diff}}^\top \left( \int_{s_1}^{s_2} \mathbf{U}(s)^\top \mathbf{U}(s) ds \right) \mathbf{W}_{\text{Diff}} \mathbf{X}_{\text{Diff}} \quad (\text{E.35})$$

□

## E.4 Generic tracker model — parameter estimation

In Section 4.5.1, page 44, we assumed a second order autoregressive process for the dynamical tracking model. This takes the form,

$$\mathbf{X}_t - \bar{\mathbf{X}} = \mathbf{A}_1(\mathbf{X}_{t-1} - \bar{\mathbf{X}}) + \mathbf{A}_2(\mathbf{X}_{t-2} - \bar{\mathbf{X}}) + \mathbf{B}_0 \mathbf{w}_t \quad \text{and} \quad \mathbf{w}_t \sim \text{i.i.d } \mathcal{N}(\mathbf{0}, \mathbf{I}). \quad (\text{E.36})$$

Here  $\mathbf{A}_0$ ,  $\mathbf{A}_1$  and  $\mathbf{B}$  are matrices, and  $\mathbf{w}(n)$  is Gaussian noise.  $\mathbf{B}$  is the size of the variation that we observe (tracking error). The univariate equivalent scalars are  $a_0$ ,  $a_1$  and  $b$ .

The discrete-time system parameters are estimated via Maximum Likelihood estimation (MLE). Assuming that the noise is Gaussian, it is straightforward to set up and minimise the likelihood function. We assume that the mean configuration is  $\mathbf{0}$ .

### E.4.1 Constant dynamics within a class — Fixed effects model

We now want to obtain a generic tracker model. Let the number of classes be  $\mathcal{C}$ . Let the number of objects within a class be  $\mathcal{K}_c$ . Suppose that each object within a class is observed in motion for  $\mathcal{M}_k$  timesteps. We seek to estimate the dynamics  $\mathbf{A}_0$ ,  $\mathbf{A}_1$  and  $\mathbf{B}$ . We assume that  $\mathbf{B}$  does not vary much between individuals.

#### Univariate case

Since the noise is assumed to be distributed  $\mathcal{N}(0, b^2)$ , the probability density function takes the form,

$$f(x|a_0, a_1, b) = \frac{1}{\sqrt{2\pi} b} \exp \left\{ -\frac{1}{2b^2} (x_{t+2} - a_1 x_{t+1} - a_0 x_t)^2 \right\}. \quad (\text{E.37})$$

The corresponding likelihood function is then the product of the individual probability densities,

$$L(\mathbf{x}|a_0, a_1, b) = \prod_{k=1}^{\mathcal{K}_c} \prod_{n=1}^{\mathcal{M}_k-2} \frac{1}{\sqrt{2\pi} b} \exp \left\{ -\frac{1}{2b^2} ({}^k x_{t+2} - a_1 {}^k x_{t+1} - a_0 {}^k x_t)^2 \right\}. \quad (\text{E.38})$$

Maximising the likelihood function is entirely equivalent to maximising the log-likelihood function, which takes the form,

$$\log L(\mathbf{x}|a_0, a_1, b) = l(\mathbf{x}|a_0, a_1, b) = \sum_{k=1}^{\mathcal{K}_c} \sum_{n=1}^{\mathcal{M}_k-2} -\log \sqrt{2\pi} - \log b - \frac{1}{2b^2} ({}^k x_{t+2} - a_1 {}^k x_{t+1} - a_0 {}^k x_t)^2, \quad (\text{E.39})$$

$$l(\mathbf{x}|a_0, a_1, b) \propto \sum_{k=1}^{\mathcal{K}_c} \sum_{n=1}^{\mathcal{M}_k-2} \frac{1}{2b^2} ({}^k x_{t+2} - a_1 {}^k x_{t+1} - a_0 {}^k x_t)^2 - \mathcal{K}_c(\mathcal{M}_k - 2) \log b. \quad (\text{E.40})$$

The MLE for the coefficients  $a_0, a_1, b$  is obtained by maximising equation E.39 with respect to each parameter in turn.

$$\begin{aligned} \frac{\partial l}{\partial a_0} &= -\frac{1}{2b^2} \sum_{k=1}^{\mathcal{K}_c} \sum_{n=1}^{\mathcal{M}_k-2} ({}^k x_{t+2} - a_1 {}^k x_{t+1} - a_0 {}^k x_t) {}^k x_t \\ &= 0 \text{ when,} \\ &\sum_{k=1}^{\mathcal{K}_c} \sum_{n=1}^{\mathcal{M}_k-2} ({}^k x_{t+2} - \hat{a}_1 {}^k x_{t+1} - \hat{a}_0 {}^k x_t) {}^k x_t = 0 \\ &\sum_{k=1}^{\mathcal{K}_c} \sum_{n=1}^{\mathcal{M}_k-2} {}^k x_{t+2} {}^k x_t = \hat{a}_0 \sum_{k=1}^{\mathcal{K}_c} \sum_{n=1}^{\mathcal{M}_k-2} {}^k x_t {}^k x_t + \hat{a}_1 \sum_{k=1}^{\mathcal{K}_c} \sum_{n=1}^{\mathcal{M}_k-2} {}^k x_{t+1} {}^k x_t. \end{aligned} \quad (\text{E.41})$$

Now if we let,

$$s_{ij} = \sum_{k=1}^{\mathcal{K}_c} \sum_{n=1}^{\mathcal{M}_k-2} {}^k x_{n+i} {}^k x_{n+j} \quad i, j = 0, 1, 2, \quad s_{ij} = s_{ji}, \quad (\text{E.42})$$

then,

$$s_{20} = \hat{a}_0 s_{00} + \hat{a}_1 s_{10}. \quad (\text{E.43})$$

Differentiating with respect to  $a_1$ , similarly gives,

$$s_{21} = \hat{a}_0 s_{01} + \hat{a}_1 s_{11}. \quad (\text{E.44})$$

We now simultaneously solve Equations E.43 and E.44 to obtain the MLE of the parameters  $a_0$  and  $a_1$ .

The MLE estimate for  $b$  is obtained by maximising the log-likelihood function (equation E.39) with respect to  $b$ , holding  $a_0$  and  $a_1$  constant, giving,

$$\hat{b}^2 = \frac{1}{\mathcal{K}_c(\mathcal{M}_k - 2)} \sum_{k=1}^{\mathcal{K}_c} \sum_{n=1}^{\mathcal{M}_k-2} (x_{t+2} - \hat{a}_1 x_{t+1} - \hat{a}_0 x_t)^2. \quad (\text{E.45})$$

### Multivariate case

Multivariate estimation follows analogously. The noise is assumed to be distributed  $N(\mathbf{0}, \mathbf{B})$ . Letting  $\mathbf{B}^{-1} = \mathbf{V}$ , the probability density function takes the form,

$$f(\mathbf{X}|\mathbf{A}_0, \mathbf{A}_1, \mathbf{B}) = |2\pi|^{-\frac{1}{2}} |\mathbf{V}|^{\frac{1}{2}} \exp \left\{ -\frac{1}{2} ({}^k \mathbf{X}_{t+2} - \mathbf{A}_1 {}^k \mathbf{X}_{t+1} - \mathbf{A}_0 {}^k \mathbf{X}_t \mathbf{V} ({}^k \mathbf{X}_{t+2} - \mathbf{A}_1 {}^k \mathbf{X}_{t+1} - \mathbf{A}_0 {}^k \mathbf{X}_t)^T \right\}. \quad (\text{E.46})$$

The likelihood function has the form,

$$L(\mathbf{X}|\mathbf{A}_0, \mathbf{A}_1, \mathbf{B}) = \prod_{k=1}^{\mathcal{K}_c} \prod_{n=1}^{\mathcal{M}_c-2} |2\pi|^{-\frac{1}{2}} |\mathbf{V}|^{\frac{1}{2}} \exp\left\{-\frac{1}{2}({}^k\mathbf{X}_{t+2} - \mathbf{A}_1 {}^k\mathbf{X}_{t+1} - \mathbf{A}_0 {}^k\mathbf{X}_t) \mathbf{V} ({}^k\mathbf{X}_{t+2} - \mathbf{A}_1 {}^k\mathbf{X}_{t+1} - \mathbf{A}_0 {}^k\mathbf{X}_t)^T\right\}. \quad (\text{E.47})$$

The MLE estimates for the parameters  $\mathbf{A}_2, \mathbf{A}_1$  and  $\mathbf{B}$  are obtained by maximising this likelihood or the corresponding log-likelihood function, given by,

$$\begin{aligned} \log L(\mathbf{X}|\mathbf{A}_0, \mathbf{A}_1, \mathbf{B}) &= l(\mathbf{X}|\mathbf{A}_0, \mathbf{A}_1, \mathbf{B}) \\ &= -\mathcal{K}_c \frac{\mathcal{M}_c - 2}{2} \log|2\pi| + \mathcal{K}_c \frac{\mathcal{M}_c - 2}{2} \log|\mathbf{V}| \\ &\quad - \frac{1}{2} \sum_{k=1}^{\mathcal{K}_c} \sum_{n=1}^{\mathcal{M}_c-2} ({}^k\mathbf{X}_{t+2} - \mathbf{A}_1 {}^k\mathbf{X}_{t+1} - \mathbf{A}_0 {}^k\mathbf{X}_t)^T \mathbf{V} ({}^k\mathbf{X}_{t+2} - \mathbf{A}_1 {}^k\mathbf{X}_{t+1} - \mathbf{A}_0 {}^k\mathbf{X}_t). \end{aligned} \quad (\text{E.48})$$

Now analogously with Equation E.42, we can define  $\mathbf{S}$  as,

$$\mathbf{S}_{ij} = \sum_{k=1}^{\mathcal{K}_c} \sum_{n=1}^{\mathcal{M}_c-2} {}^k\mathbf{X}_{n+i} {}^k\mathbf{X}_{n+j} \quad i, j = 0, 1, 2, \quad \mathbf{S}_{ij}^T = \mathbf{S}_{ji}. \quad (\text{E.49})$$

Therefore,

$$\mathbf{S}_{20} - \hat{\mathbf{A}}_0 \mathbf{S}_{00} - \hat{\mathbf{A}}_1 \mathbf{S}_{10} = \mathbf{0}. \quad (\text{E.50})$$

Differentiating with respect to  $\mathbf{A}_1$ , similarly gives,

$$\mathbf{S}_{21} - \hat{\mathbf{A}}_0 \mathbf{S}_{01} - \hat{\mathbf{A}}_1 \mathbf{S}_{11} = \mathbf{0}. \quad (\text{E.51})$$

The MLE estimate for  $\mathbf{V}$  is obtained by maximising the log-likelihood function (Equation E.48) with respect to  $\mathbf{V}$ , holding  $\mathbf{A}_1$  and  $\mathbf{A}_2$  constant.

$$\mathbf{V}^{-1} = \frac{1}{\mathcal{K}_c (\mathcal{M}_c - 2)} \sum_{k=1}^{\mathcal{K}_c} \sum_{n=1}^{\mathcal{M}_c-2} ({}^k\mathbf{X}_{t+2} - \mathbf{A}_1 {}^k\mathbf{X}_{t+1} - \mathbf{A}_0 {}^k\mathbf{X}_t) ({}^k\mathbf{X}_{t+2} - \mathbf{A}_1 {}^k\mathbf{X}_{t+1} - \mathbf{A}_0 {}^k\mathbf{X}_t)^T. \quad (\text{E.52})$$

Recall that we let  $\mathbf{B} = \mathbf{V}^{-1}$ , so the MLE of  $\mathbf{B}$  can be obtained.

# APPENDIX F

---

## Tracking results


---


This chapter shows a number of sets of tracking results to justify statements made in the main body of the thesis.


### F.1 Comparing affine and PCA tracking shape-spaces


This section is relevant to the conclusions drawn in Section 5.1.4, page 63. Here we show results to provide a visual comparison of tracking performance using an affine and PCA shape-space. For each frame illustrated below, the top image represents tracking using the affine shape-space and the bottom images the equivalent frame for tracking with a PCA shape-space. Both trackers use the same default dynamics.


One cardiac cycle (80 frames) is shown below. The data was digitally captured and stored on optical disk. The frequency of the data was 61 Hz.
















## F.2 Comparing feature detection schemes



This Section is relevant to the conclusions drawn in Section 5.1.4, page 63. Here we show results to provide a visually comparison of tracking performance using intensity and intensity-invariant feature detection schemes.



For each frame illustrated below, the images represent; (top) gradient-based feature detection on the original data (the original visual image tracker); (middle) gradient-based feature detection on TDLMS/IBS enhanced images; and (bottom) local-phase based feature detection on TDLMS/IBS enhanced images.



Two cardiac cycles are shown below. The data was recorded on VHS video and then digitised. Because of this, the acquisition frequency of the data was 25 Hz. The images shown start at frame 2, because the TDLMS filter we used considered 2 frames before, and 2 frames after, the frame at time  $t$ . The first frame was frame 0, and hence the first TDLMS frame is frame 2.











# APPENDIX G

---

## Stress Echocardiography

---

### G.1 Motivation for stress echocardiography

Under resting conditions, patients with coronary artery disease may have normal left ventricular function. If, at the time of the examination, no permanent myocardial damage has occurred, and if the ventricle is not ischemic, the study will not reflect any underlying coronary artery disease.

There has been considerable interest in combining standard echocardiography with stress intervention to induce ischemia (Limacher *et al.*, 1983; Marwick *et al.*, 1992; Aldrich & Reichek, 1993). This is because exercise stress increases cardiac oxygen consumption, and may provoke ischemia in patients with coronary stenosis (Pozen *et al.*, 1981).

Early studies were based on exercise electrocardiograms (ECG), but the usefulness of this method in terms of sensitivity and specificity is questionable (Epstein, 1976). Ryan *et al.* (1988) argues that exercise echocardiography is superior to treadmill ECG for the diagnosis of coronary artery disease in those patients with abnormal baseline ECG, and provides valuable prognostic information in patients recovering from acute myocardial infarction. Currently, there are a number of methods available to the clinician to induce ischemia — Table G.1 summarises these.

The most common methods are exercise and pharmacological stress. These are discussed next.

#### G.1.1 Traditional stress echocardiography

Traditionally, studies used supine or upright bicycle exercise (Hecht *et al.*, 1993) with patients being imaged throughout exercise. There are, however, a number of limitations associated with these techniques, such as leg fatigue and potential difficulties in obtaining an adequate acoustic window. Moreover, some patients are not fit to cycle.

A treadmill test is now the most popular form of stress testing. An excellent review is given by Armstrong (1997). With treadmill exercise (Armstrong, 1997) the patient is imaged using the four standard views in the supine position, creating a baseline for reference. Then the patient exercises on the treadmill. The speed and inclination of the treadmill is increased according to a defined protocol (Bruce & Hornstein, 1969; McHenry *et al.*, 1972), so that in practice the patient remains on the treadmill for around 8–10 minutes. Throughout this time, the patient's ECG, blood pressure and heart rate is monitored, although no imaging takes place. When the patient indicates that they can no longer exercise or a target heart rate is reached, they are immediately returned to the examination couch and re-imaged post exercise using the same four standard views. Ryan *et al.* (1988); Limacher *et al.* (1983) and Quiñones *et al.* (1992), in their studies, marked optimal transducer positions on the patients chest before exercise. This

---

Exercise  
   Post Treadmill  
   Supine bicycle  
   Upright bicycle  
 Pharmacologic stress  
   Adrenergic stimulation  
     Dobutamine  
     Isuprel  
     Arbutamine  
     Epinephrine  
   Vasodilation  
     Dypridamole  
     Adenosine  
 Pacing  
   Esophageal  
   Atrial  
   Ventricular  
 Handgrip  
 Coldpress

---

Table G.1: **Stress Echocardiography.** Methods available to induce ischemia. Based on Feigenbaum (1994).

is because human bodies differ and therefore it can take as much as 20 seconds to obtain an optimal position for a given view. In the context of stress echocardiography, the idea is to image the patient while they are under full stress. After leaving the treadmill, the heart begins to recover. Any time then lost in re-finding optimal positions may mean that abnormalities do not show up. Armstrong (1997) suggests that the greatest advantage of post treadmill imaging is that *"in many countries (United States being the most obvious), walking or treadmill exercise is more familiar to patients than is bicycle exercise"*.

Although treadmill exercise may be considered the most vigorous form of stress, there are some key limitations.

1. Many patients are not suitable for, or unable to, exercise.
2. It is not possible to image the patient during exercise, which contradicts with the idea of observing left ventricular wall motion under stress. Only baseline and post exercise images are acquired for analysis. The clinician is not afforded the stage of onset of the ischemic wall motion abnormality and hence information about the severity of a coronary stenosis.
3. The degree of cardiac stress can not be controlled; dobutamine may provide a more accurate assessment of physiological significance of a coronary artery lesion.
4. The patients is only at "full exercise" as soon as they leave the treadmill, and this rapidly decreases as the heart rate falls. This ventricular function recovery depends on the duration and severity of the ischemia. Mild degrees of ischemia may resolve rapidly and be missed, even if the recommendation that the patient is re-imaged within 60–120 seconds of leaving the treadmill is achieved (Aldrich & Reichek, 1993). Many groups claim that they can re-image patients in 60–180 seconds (Quiñones *et al.*, 1992; Marwick *et al.*, 1992; Hoffmann *et al.*, 1993). Since imaging takes place in 4 standard views, the recommendation implies that the clinician has, at best, 15-30 seconds to re-image each view. Taking into account the fact that the patient needs to be moved from the treadmill back to the examination couch and the optimal transducer position needs to be re-found, 30 seconds is an overestimate. Moreover, since the patient will be breathing very heavily, increased lung movement often means that the images are very unclear (because ultrasound does not pass through air). So recording an 'appropriate' cardiac cycle will add more time to the imaging process.
5. Exaggerated chest and lung movement following exercise can obscure the patient's acoustic window (Picano *et al.*, 1987). Due to this, Armstrong (1997) comments that during this re-

imaging it is “*highly advantageous to have the patient briefly hold their breath*” — perhaps not an easy task for a patient with ischemia who has just exercised for 10 minutes!

6. Clearly, the notion of “full exercise” is patient dependent, but it is not quantifiable.

Many of the limitations of exercise stress echocardiography can be overcome by using drugs to mimic the exercise process. This approach, used by the cardiologists at the John Radcliffe Hospital, Oxford, is discussed next.

### G.1.2 Pharmacologic stress testing

In Pharmacologic stress testing a pharmacologic drug agent is intravenously injected into a patient, which increases oxygen demand, and hence the heart needs to work harder.

Pharmacologic stress testing is probably the most popular non exercise form of stress echocardiography (Previtali *et al.*, 1993), and there are two general approaches; nuclear cardiology perfusion techniques and techniques using a pharmacological agent.

Nuclear cardiology perfusion techniques (Huang *et al.*, 1997), such as multi-gated acquisition (MUGA) scanning (Harbert & Da Rocha, 1984; McKiddie *et al.*, 1996), generally use vasodilators, such as dipyridamole (Picano & Lattanzi, 1991) or adenosine (Marwick, 1997). Vasodilators dilate the normal arteries at the expense of the obstructed arteries, resulting in blood being shunted towards the normal myocardium. The patient is then imaged with a MUGA scanner, which is ECG gated. By measuring blood counts (using a gamma camera) following injection of a radiopharmaceutical, information is supplied about cardiac function. This is then averaged over many cycles. Wall motion abnormalities can then be detected as the blood is shunted away from the obstructed arteries. Likewise, adenosine has been shown to induce regional wall motion abnormalities.

Nuclear cardiology perfusion, although widely used, is not without its problems. The most obvious is the safety aspect of injecting a radioisotope into the patient. Further, because cycles are averaged, any patient movement causes misregistration errors (Falseti *et al.*, 1981). Importantly, MUGA is not capable of imaging myocardial thickening. Lastly, from the point of view of a diagnosis, is that the majority of information MUGA scanning provides is global, whereas ischemia tends to be regionalised.

An alternative to nuclear cardiology perfusion is to use a pharmacologic agent such as dobutamine (Ryan, 1991; Secknus & Marwick, 1997). This is the drug used at the John Radcliffe Hospital, Oxford.

A thorough review of all dobutamine stress echocardiography studies is given by Geleijnse *et al.* (1997). The technique was first described by Tuttle & Mills (1975). It was first reported as a cardiac stress agent by Mason *et al.* (1984), and it has gained more widespread use recently.

During a dobutamine infusion, both the patient's heart rate and blood pressure is increased, mimicking exercise (Carlson *et al.*, 1988). This increase in heart function, is observed echocardiographically as an increase in both endocardial wall excursion and wall thickening. That is unless ischemia develops (Pozen *et al.*, 1981).

The normal response to dobutamine consists of hyperdynamic motion at peak dose, with all segments thickening and moving inward to a greater extent than in the corresponding rest image. However, in the presence of significant coronary artery disease, a wall motion abnormality develops in the region of the myocardium supplied by the critically stenosed vessel.

As well as a 'normal' and 'abnormal' response to dobutamine, there are two other types of response: Hibernating myocardium and Stunned myocardium.

“Hibernating myocardium” causes a “biphasic” response of wall thickening during low and high dose infusion (Senior & Lahiri, 1995). In patients with myocardial infarction with regional asynchrony, wall thickening improves with administration of low dose dobutamine then deteriorates with administration of high dose dobutamine (15 to 40Mg/Kg/min) — the so called “biphasic response” indicating hibernating myocardium (Krauss *et al.*, 1994; Haque *et al.*, 1995).

In contrast, "Stunned myocardium" (Chen *et al.*, 1995) refers to viable cardiac muscle that has been threatened by ischemia (e.g. in the zone immediately surrounding an infarct), but which fails to contract for a variable period (sometimes 24–48 hours or more) after re-establishment of perfusion. It represents the combined results of ischemic and reperfusion injury. It has been postulated that the contractile failure is related to intracellular accumulation of calcium, activating calcium-dependent proteases, thereby injuring myofilaments. Dobutamine stress echocardiography has been reported to identify stunned myocardium by demonstrating improvement in myocardial thickening and regional wall motion compared to a baseline exam (Chen *et al.*, 1995).

In this section we have considered the use of pharmacologic drugs as a way of mimicking the effects of exercise on the heart. We now compare which pharmacologic agent should be used to diagnose coronary artery disease.

### Comparison of Exercise, Dobutamine, Dipyridamole and Adenosine

As was noted, exercise causes myocardial ischemia through a marked increase in myocardial oxygen demand, resulting from an increase in heart rate, systolic blood pressure and contractility. Dobutamine acts in exactly the same way. On the other hand, dipyridamole-induced ischemia arises through a slight increase in myocardial oxygen demand, but more importantly a decrease in oxygen supply, thus predominantly affecting the supply part of the supply-demand ratio.

Therefore, dobutamine stress, more closely approximates exercise than does dipyridamole (Previtali *et al.*, 1991), and is the choice drug for stress echocardiography for the clinician.

In the next Section, we consider dobutamine stress echocardiography in more detail, first looking at the procedure in Section G.1.3 and then going on to look at how results of the procedures are analysed.

### G.1.3 Dobutamine stress echocardiography procedure

In this section we shall look at the standard infusion and imaging protocols involved in dobutamine stress echocardiography.

#### Drug protocol

Dobutamine infusion is used with increasing doses of 5, 10, 20, 30 and 40Mg/Kg/min. Each stage may last as much as 5 minutes, allowing time for the heart rate to stabilise at a 'within dose' level. At peak dose, if the patient has not reached the target heart rate — the chronotropic response is inadequate — up to 20Mg of atropine (i.e. glaucoma) may be given (Rallidis *et al.*, 1997; Owen *et al.*, 1998). If patients are unable to receive atropine, a dobutamine dose of 50Mg/Kg/min may be given.

There are a number of reasons why dobutamine should be gradually infused. Firstly, in comparison to exercise stress, it allows the site and extent of the ischemia to be fully recognised. Secondly, it is safer bearing in mind potential side effects. Thirdly, it allows hibernating myocardium to be diagnosed.

The protocol must include the end point of the study, such as, a new wall motion abnormality, diagnostic ST segment shift (ECG), achieving a target heart rate, significant symptoms, significant change in blood pressure or completion of the protocol (Mazeika *et al.*, 1992; Previtali *et al.*, 1993).

#### Imaging Protocol

The patient is studied while lying on the examination couch in a left decubitus position. ECG (12 lead) recordings are made at each drug dose and in recovery. The patient's heart rate and blood pressure is also recorded at each dose. During the examination, ventricular wall motion is visualised in the standard parasternal long and short axis, and the apical two and four chamber views.

At baseline, each infusion dose level and recovery, the clinician will record echocardiographic images of each view onto video tape (Attenhofer *et al.*, 1994). Usually 8 images in systole are then digitised, and arranged and displayed as cine loop. Nowadays, newer ultrasound machines are able to directly digitally store the data from the ultrasound screen. For efficiency, in the case of treadmill exercise, where images need to be acquired rapidly, video is still used, and appropriate cycles are subsequently digitized to disk.

Since only early systole is recorded, potential important endocardial motion during the rest of the cardiac cycle is ignored (Mazeika *et al.*, 1992). Each echocardiographic view can then be shown as part of a quad-screen format with each quadrant showing a cine loop of a single cycle played synchronously. Generally, though, the clinician will view a cine loop of a single view recorded at four different infusion times — the baseline and three levels of the drug, or the baseline, two levels of the drug and recovery.

## G.2 Stress echocardiography analysis

Stress echocardiography, whether this be bicycle, exercise, or pharmacology based, provides the technique to induce ischemia, but this change in regional wall function needs to be analysed in some way. There are generally three ways of doing this — using M-mode echocardiography, electrocardiogram (ECG) analysis or two-dimensional echocardiography.

Traditionally, M-mode echocardiography was used in early work (Corya *et al.*, 1975) but has been superseded by two-dimensional echocardiography, because it provides a more complete spatial picture of the wall motion. A number of groups (Previtali *et al.*, 1991; Marcovitz & Armstrong, 1992) have used electrocardiograms (ECG) to provide supplementary information, although Rallidis *et al.* (1997) comment that the dobutamine stress test leads to uninterpretable ECG results. Marwick *et al.* (1992) compares exercise ECG and exercise echocardiography and concludes that exercise echocardiography had a significantly higher sensitivity than exercise ECG. The reason given is that exercise ECG is significantly influenced by the performance of maximal exercise.

### The 16-segment model of the heart

One way to qualitatively assess wall motion in real-time is to indicate which segments are moving abnormally. The obvious problem with a total qualitative assessment is that it is subjective, and that the clinician is more likely to notice the most striking abnormalities and may overlook subtle effects taking place.

The systematic way to assess regional wall motion is to divide the ventricular boundary into segments corresponding to significant blood supplies to the heart surface. The number of segments varies from 9 to 20 among different groups (Stamm *et al.*, 1983; Previtali *et al.*, 1991, 1993; Mazeika *et al.*, 1992; Hecht *et al.*, 1993; Beleslin *et al.*, 1994; Senior & Lahiri, 1995).

The generally accepted model adopted by The American Society of Echocardiography is the 16-segment approach (Schiller *et al.*, 1989) proposed by Bourdillon *et al.* (1989) and detailed in Section 6.2, page 84. The model is shown in Figure 6.4, page 85. This model is presently used by a number of groups (Krahwinkel *et al.*, 1997; Marwick, 1997; Rallidis *et al.*, 1997; Cokkinos *et al.*, 1998). The model is also used by the John Radcliffe Hospital.

### G.2.1 Qualitative analysis of wall motion

The standard scoring scheme proposed by the American Society of Echocardiography has been adopted by a number of groups (Krahwinkel *et al.*, 1997; Cokkinos *et al.*, 1998) and that shown in Table 6.2, page 6.2. A normal segment is given the value 1. Abnormal segments have higher numbers. Occasionally, 5 is awarded to an aneurysmal region. The clinician's score is based on the myocardial thickening of each segment. If however, the myocardium thickening is not clearly visible, but the endocardial border is, then the clinician's score is based on the observed endocardial excursion of each segment.

Test outcome	Wall motion		Wall thickening	
	Low Dose	High Dose	Low Dose	High Dose
Normal	increase	increase	increase	increase
Positive	increase/decrease	increase/decrease	increase/decrease	increase/decrease
Biphasic	increase	decrease	increase	decrease

Table G.2: **Dobutamine stress echocardiography outcomes.** The main outcomes of a dobutamine stress echocardiography test.

### Definition of test result

There are a number of outcomes to a dobutamine stress echocardiography procedure outlined in Table G.2.

**Normal test result:** A normal stress echocardiogram is defined by a uniform increase in wall motion and systolic wall thickening, with a reduction in end-systolic cavity area.

**Positive test result:** A positive test result is generally defined as the appearance of a new or worsening wall motion abnormality (Sawada *et al.*, 1991; Afridi *et al.*, 1994; Cokkinos *et al.*, 1998).

**A biphasic response:** In patients with rest wall abnormalities, initial improvement of wall motion and thickening at low dose followed by worsening at higher dose.

**Other results:** Tardokinesia (delayed excursions) and relative failure to augment wall thickening.

Elhendy *et al.* (1998) looked at the optimal criteria for the diagnosis of coronary artery disease (CAD). They compared using one, and two or more ischemic segments as a marker for CAD. They found that using one ischemic segment is diagnostically sufficient for patients without previous myocardial infarction. With patients with previous myocardial infarction, the use of two or more segments reduces the sensitivity without significant increase in specificity.

Several groups have tried to assign a score that characterise the heart at rest (pre-infusion) and each of the levels of the drug (Sawada *et al.*, 1991; Cokkinos *et al.*, 1998). An overall wall motion score was defined as the summation of individual segment scores divided by the number of segments. Sawada *et al.* (1991) uses these dose specific scores to look for significant differences between patients with one and multivessel disease. A potential problem of this is that the average score for a patient with a regional wall abnormality can be masked by the rest of the heart being normal.

Until now, all the methods aimed at looking at regional heart function have been qualitative. In the next Section, we consider quantitative approaches.

## G.2.2 Quantitative analysis of wall motion

The previous Sections have dealt with qualitative approaches. To date, there has been little work on quantitative analysis.

Research into a quantitative (partially automated) scoring system for dobutamine stress echocardiography has been slow. Little mention is made in the literature about quantitative analysis. Even if it is, the authors have tended to dismiss the possibility (Marcovitz & Armstrong, 1992). Mazeika *et al.* (1992) comments, without reference, that "*Quantitative analysis of regional function was not performed in this study because of the known confounding effects of cardiac translation and rotation*". Comments like these are often supplemented with "*The qualitative approach adopted by us and other workers resulted in good interobserver agreements*".

A possible reason for the lack of interest in a quantitative approach has been the poor quality of images obtainable from an ultrasound machine. Marwick *et al.* (1992), in their study, graded images into four categories according to endocardial definition. Their "A quality" was one that showed good endocardial definition with good image quality in all views. Whereas, their "D

quality” reflected severely compromised images with failure to visualize one or more coronary territories.

With the advent of improved ultrasound machine in the last five years, particular the capabilities to save data digitally, endocardial (if maybe at times not epicardial) border definition has been significantly improved. Of course, whatever the capabilities of the machine, no machine can ever combat a patient with a poor acoustic window.

Even with the newer technology, recent published work on dobutamine stress testing (Cokkinos *et al.*, 1998; Krahwinkel *et al.*, 1997) still uses subjective scoring, and no attempt, or even mention, was paid to quantifying the procedure, or at the very least noting possible failings of a qualitative approach.

The most obvious reason for this is that no one has been able to robustly and reliably track the endocardium and epicardium boundaries echocardiographic images. We have seen in Chapters 5 and 7 that we can track these borders.

### Qualitative or quantitative analysis

One reason for clinicians still favouring a qualitative analysis, is that they are approximately able to eyeball abnormalities, and the extra accuracy achievable by a quantitative analysis may not be important is their diagnosis.

While this may be true, there are a number of specific advantages of a quantitative over a qualitative dobutamine stress analysis.

1. If patients are to be monitored over time, qualitative scoring will suffer from subjective interobserver and intraobserver errors.
2. The clinician will want to be certain that when they classify a segment of the heart as abnormal, it really is. A simple examination of these scores will show a constant level of score, then an immediate step change in score. Clearly, the progress of the disease is more likely to be continuous — only a continuous quantifiable technique can detect the change in perfusion.
3. Precise quantification of the magnitude of the ischemic may be necessary for prognostic purposes or for evaluating the results of an intervention.
4. Much of the literature has been focused on asking: how does stress echocardiography compare to nuclear perfusion; how dobutamine compares to dipyridamole; and, consequently how they compare to exercise stress and angiography. We are not going to look at these issues.

To answer these questions, the sensitivity, specificity and accuracy of these tests are often compared and quoted. Different groups use different basis to calculate these numbers — a quantitative analysis will remove some of the bias from the comparisons.

What little work there is on quantitating stress echocardiography, has dealt with trying to partially automate area and volume measurements. The clinician then uses these (global) measurements to assign scores to heart function. This still leaves the subjective wall scoring to the clinician. The novel difference with our work is that we will use our endocardial/myocardial border detection to assign the score *directly*. More importantly, area and volume measurements are global measures that can only look at global systolic function and hence *only* global heart disease. Whereas, coronary artery disease, tends to be regionalised. As such, we need regional measures to detect these abnormalities.

In some early work Rifkin & Koito (1990) tried to estimate the ejection fraction using a segmental geometrical scoring system. One notable aspect to their method was that wall motion in each segment was scored on a continuous scale. Normal contraction = 100, akinesis = 0, and various degrees of hypokinesis were represented by wall motion scores between 0 and 100 — mild (75), moderate (50) and severe (25) hypokinesis. Dyskinesis and hyperkinesis were represented by scores below 0 and above 100, respectively.

Peréz *et al.* (1992c) following work in their group, describes a system “allowing automatic detection and tracking of these [endocardial/blood interface] in real time using a refinement of a previous

*approach and novel software*". They apply the border detection technique to obtain area (and fractional change in area) of the cavity to examine global cardiac response. No attempt was made to consider regional wall motion or provide any (quantitative) score of wall motion. This system was also used by Vitarelli *et al.* (1997) to find various volumetric measurements during dobutamine stress echocardiography.

### G.2.3 Safety

As with other stress tests in coronary artery disease, dobutamine stress echocardiography has certain risks. These risks can be segregated into risks due to the coronary artery disease itself, and those associated with the side effects of the drug.

Krahwinkel *et al.* (1997) reports on five other studies in which complications occurred. They also report on their own experience of 312 patients, and complications encountered. They note that the test can easily be interrupted if any complications occur.

The most serious side effects due to the coronary artery disease are transient arrhythmias or hypertension, none of which has been life threatening.

More common, although still around 10/25% of patients (Mertes *et al.*, 1993), are non-cardiac side effects. These are listed in Table G.3. They do not, though, generally result in the test being terminated.

Palpitations	Nausea
Vertigo	Headache
Tremor	Anxiety
Flush	Urgency
Allergic reactions	Intravenous access malfunction

Table G.3: **Non-cardiac dobutamine side effects.** Non-cardiac side effects of dobutamine stress echocardiography.

### G.2.4 Sensitivity, Specificity and Accuracy

To gauge their diagnostic value, the various stress echocardiography techniques need to be compared in some way. The adopted method amongst cardiology groups is to report sensitivity, specificity and accuracy values, with coronary angiography taken as ground truth (Hoffmann *et al.*, 1993).

The reported sensitivity values for a number of recent papers are shown in Table G.4.

Author	year	Patient No.	Sensitivity (%)	Specificity (%)	Accuracy (%)
Cohen	1991	70	86	95	89
Mazeika	1992	50	78	93	84
Marcovitz	1992	141	96	66	N/A
Segar	1992	85	95	82	92
Marwick	1993	217	72	83	76
Hoffman	1993	64	79	81	80
Previtali	1993	80	79	83	80
Beleslin	1994	136	82	77	82

Table G.4: **Dobutamine sensitivity, specificity and accuracy values.** Dobutamine stress echocardiography sensitivity, specificity and accuracy values obtained by various groups.

## Bias

The problem with merely quoting these values is that the sensitivity, in particular, depends on the inclusion in the study of patients with wall abnormalities at rest, which then has implications for comparing these values between groups; direct comparison of the results of studies is difficult as they involve heterogeneous groups of patients.

Inclusion of such patients with a high likelihood of coronary artery disease biases the test and prevents distinction between stable preexisting wall motion abnormalities and exercise-induced ischemia. Ideally patients with rest wall motion abnormalities should be excluded from the study (Ryan *et al.*, 1988). Another way of thinking about this is that patients with a low pretest probability of CAD are rarely referred for invasive evaluation — therefore those patients with a normal coronary angiogram are a selected population and do not represent a truly normal population (Aldrich & Reichek, 1993).

Sensitivity-specificity is a binary classification, which Demer *et al.* (1989) states causes three main limitations:

1. Coronary artery disease is not an all-or-none condition: binary classification requires arbitrary threshold criteria and creates artificial distinctions in CAD that in actuality have a continuous spectrum of severity.
2. Sensitivity and specificity values are determined by the disease distribution of the study population: a sample distribution with a high frequency of mild disease, where scatter is more likely to lower sensitivity and specificity.
3. Percent diameter narrowing is not an adequate standard for quantifying stenosis in clinical studies.

Picano & Lattanzi (1991) further notes that a positive CAD response can be stratified along temporal and spatial coordinates that can greatly strengthen the diagnostic and prognostic information derived from a positive test.

## Standard errors and odd ratios

Beleslin *et al.* (1994) goes one stage further and as well as quoting the sensitivity, specificity and accuracy of the stress techniques they use, they also give 95% confidence intervals based on assumption that these follow a binomial distribution.

If  $p$  is the sensitivity rate, where,

$$p = \frac{\text{True Positives}}{\text{Total number positive}} = \frac{TP}{M} \quad (\text{G.1})$$

then,  $p \sim \text{Bin}(M, \hat{p})$ , and,  $E(\hat{p}) = p$  and  $\text{Var}(\hat{p}) = \frac{p(1-p)}{M}$  •

Whether the difference between the techniques is statistically significant can then be found from the confidence intervals.

## G.3 Comparison of Dobutamine echocardiography, Dipyridamole echocardiography and Exercise echocardiography

### Dobutamine

Segar *et al.* (1992) shows that dobutamine stress echo has a high sensitivity and specificity for the detection and localisation of coronary artery disease.

### Dobutamine vs Exercise

In patients unable to exercise adequately, the diagnostic accuracy of dobutamine stress echocardiography has led to enthusiasm for its general use.

Marcovitz & Armstrong (1992) in their study, found that the specificity, sensitivity and accuracy values of dobutamine were equivalent to exercise but that it was hindered by patients with resting wall abnormalities.

Previtali *et al.* (1993) directly compared dobutamine stress echocardiography with exercise stress testing. They found that dobutamine had a higher sensitivity and significantly higher accuracy than exercise. Hoffmann *et al.* (1993) there to be no significant difference between sensitivity, specificity and accuracy values for exercise and dobutamine. Rallidis *et al.* (1997) in their randomised crossover study, directly compared the effects of echo immediately after exercise with dobutamine echo in the same patients and on the same day in patients with known coronary artery disease. Their results show that the incidence and magnitude of myocardial ischemia were greater with exercise echo than dobutamine. They conclude that "*Because of its superior hemodynamic effects, exercise echocardiography should be chosen over dobutamine for diagnosing ischemia, when possible*"

### Dobutamine vs Dipyridamole

In patients unable to exercise adequately, echocardiographic imaging can also be performed with dipyridamole.

Previtali *et al.* (1991) demonstrates that there is a similar overall sensitivity and specificity between dobutamine and dipyridamole for the diagnosis of coronary artery disease. Their findings in Previtali *et al.* (1993) indicate that dobutamine is superior to dipyridamole and exercise; they comment that the different mechanisms by which each induces ischemia may explain their differences in sensitivity.

### Dobutamine vs Exercise ECG

Hennessy *et al.* (1997) finds that dobutamine echocardiography sensitivity is better than treadmill exercise ECG, but it has lower specificity, which may mean that more patients being referred for diagnostic coronary angiography.

## G.3.1 Inter-observer and Intra-observer agreement

Using a qualitative scoring scheme in dobutamine stress echo is often justified by groups (Previtali *et al.*, 1991) on the basis that studies often use two blinded reviewers to assess the recordings, and that disagreement is accomplished using a third observer. Often the correlation between the observers is noted.

Reported (Previtali *et al.*, 1991, 1993; Beleslin *et al.*, 1994) intraobserver and interobserver agreement for ischemia ranges from 95%–98% and from 92%–96%, respectively.

Rallidis *et al.* (1997) states the average wall motion score index for two observers and states that there are no significant differences between the scores, although the standard errors of the scores are very large.

Picano *et al.* (1991) address the question: "*How does the diagnostic accuracy of a stress echocardiography procedure depend on the specific experience of the physician interpreting the test*". They show that extensive expertise in routine echocardiography at rest is necessary, but that not sufficient to adequately interpret stress echocardiographic studies.

Intraobserver variability is a more difficult, but equally important, question to address; In the study by Previtali *et al.* (1993), one reviewer reviewed a random sample of studies twice, with 96% agreement. The problem with this method occurs on the second review, which will bias the intraobserver agreement. Namely, that even if the studies were shown to the observer in a

random order (this is not clear from the paper), the observer is very likely to remember some of the data from their first viewing.

# APPENDIX H

## John Radcliffe Hospital dobutamine stress scoring sheet

The John Radcliffe Hospital dobutamine stress echocardiography scoring sheet.

John Radcliffe Hospital Cardiology Department					
Dobutamine stress echo					
Date of test					
Hospital number		OP:			
Patient name					
Patient address					
Date of birth					
Referring physician					
Clinical details					
History:					
Result of previous ETT:					
Current medication:					
Dobutamine stress results					
Stage	Time	HR	BP	Symptoms	ECG
Rest					
5mcg/kg/min					
10mcg/kg/min					
20mcg/kg/min					
30mcg/kg/min					
40mcg/kg/min					
atropine 0.6mg					
Reason for stopping:					

Wall motion analysis (Normal score: 1, Hypokinetic: 2, Aknetic: 3, Dyskinetic: 4)	
REST	
PEAK STRESS	
Report	
Date: _____ Reported by: _____	

Figure H.1: John Radcliffe Hospital dobutamine stress scoring sheet.

## APPENDIX I

---

### Contour Comparison Results

---

The results in this section are relevant to the preliminary regional comparison of the segmentation of an algorithm and clinician in Section 8.4.2.

For each patient, at rest and peak and for the endocardial and epicardial boundaries, the following area based measures were calculated:

1. Area — the area (in  $\text{cm}^2$ ) of the left ventricle from the clinician and tracker segmentation. The contours drawn by the clinician were closed by drawing a line between the end-points of the segmentation.
2. Global error — the global absolute error (in cm) between the clinician and tracker segmentation. The global error measures the overall closeness of the two segmentations.
3. Regional error — the regional error (in cm) between the clinician and tracker segmentation, based on dividing the left ventricle into six segments.

This was obtained by calculating the principal axis of the clinician's segmentation, then dividing its maximal extent into three parts. Normals are then projected from the ends of the three parts, and the positions of where the normals cut the two segmentations are recorded. This defines six segments. The regional error measures the regional closeness of the two segmentations.

4. Regional division — the percentage of the regional error that is from position of the tracker segmentation being outside the position of the clinician's segmentation. High values indicate that the position of the tracker segmentation is consistently outside the position of the clinician's segmentation. The regional division measures whether the tracker segmentation is biased in one direction.

A further two measures were tested.

1. Global rotation — the angular difference (in degrees) between the principal axis of the clinician and tracker segmentation. The global rotation measures the difference in rotation between the two segmentations.
2. Regional overlap — the regional percentage is the overlap between the clinician and tracker segmentation. This is a comparative measure of horizontal and vertical extent.

		Area (cm <sup>2</sup> )		Global Error (cm)	Regional Error (cm)						
		Track	Clín								
P1	En	RS	9.13	9.73	0.17(0.14)	0.15(0.11)	0.07(0.026)	0.26(0.22)	0.30(0.19)	0.11(0.11)	0.12(0.17)
		RD	14.15	14.64	0.21(0.11)	0.16(0.072)	0.098(0.060)	0.33(0.17)	0.43(0.25)	0.073(0.038)	0.15(0.074)
	Ep	RS	13.08	17.04	0.34(0.14)	0.24(0.17)	0.34(0.09)	0.58(0.19)	0.57(0.22)	0.22(0.096)	0.089(0.048)
		RD	18.17	21.37	0.33(0.13)	0.12(0.063)	0.12(0.040)	0.63(0.30)	0.77(0.27)	0.11(0.059)	0.18(0.054)
	En	PS	7.62	6.72	0.21(0.11)	0.16(0.10)	0.049(0.047)	0.28(0.22)	0.39(0.21)	0.27(0.08)	0.099(0.029)
		PD	11.33	9.25	0.11(0.05)	0.12(0.12)	0.028(0.0032)	0.031(0.0053)	0.027(0.0019)	0.093(0.12)	0.33(0.076)
	Ep	PS	13.70	15.70	0.28(0.12)	0.22(0.096)	0.35(0.066)	0.22(0.15)	0.13(0.11)	0.65(0.25)	0.13(0.058)
		PD	18.41	17.32	0.39(0.21)	0.24(0.15)	0.46(0.11)	0.58(0.31)	0.41(0.14)	0.27(0.16)	0.26(0.40)
P2	En	RS	10.91	10.34	0.26(0.16)	0.19(0.15)	0.35(0.13)	0.29(0.16)	0.29(0.23)	0.21(0.12)	0.26(0.15)
		RD	18.51	16.03	0.47(0.14)	0.65(0.16)	0.55(0.14)	0.38(0.16)	0.44(0.18)	0.34(0.035)	0.47(0.15)
	Ep	RS	18.01	18.78	0.45(0.17)	0.51(0.22)	0.18(0.056)	0.22(0.21)	0.73(0.22)	0.53(0.12)	0.52(0.22)
		RD	24.37	24.79	0.53(0.34)	0.41(0.31)	0.22(0.37)	0.23(0.12)	0.69(0.19)	0.71(0.78)	0.92(0.26)
	En	PS	17.04	13.87	0.53(0.37)	0.59(0.20)	1.29(1.14)	0.59(0.26)	0.32(0.22)	0.12(0.044)	0.27(0.37)
		PD	8.81	6.68	0.34(0.20)	0.48(0.30)	0.55(0.075)	0.41(0.13)	0.21(0.12)	0.079(0.037)	0.33(0.52)
	Ep	PS	/	/	/	/	/	/	/	/	/
		PD	/	/	/	/	/	/	/	/	/
P3	En	RS	/	/	/	/	/	/	/	/	/
		RD	10.90	7.14	0.40(0.21)	0.33(0.19)	0.38(0.13)	0.17(0.085)	0.52(0.26)	0.45(0.079)	0.58(0.51)
	Ep	RS	/	/	/	/	/	/	/	/	/
		RD	16.16	13.45	0.40(0.16)	0.60(0.15)	0.18(0.23)	0.27(0.18)	0.52(0.13)	0.32(0.075)	0.48(0.21)
	En	PS	/	/	/	/	/	/	/	/	/
		PD	/	/	/	/	/	/	/	/	/
	Ep	PS	/	/	/	/	/	/	/	/	/
		PD	/	/	/	/	/	/	/	/	/
P4	En	RS	15.70	14.23	0.45(0.23)	0.80(0.45)	0.23(0.06)	0.44(0.19)	0.51(0.14)	0.22(0.13)	0.52(0.40)
		RD	11.78	10.00	0.40(0.22)	0.85(0.38)	0.14(0.064)	0.37(0.27)	0.32(0.22)	0.25(0.07)	0.48(0.30)
	Ep	RS	23.17	20.88	0.45(0.23)	0.54(0.43)	0.34(0.084)	0.34(0.19)	0.39(0.16)	0.33(0.16)	0.77(0.38)
		RD	/	/	/	/	/	/	/	/	/
	En	PS	10.00	12.10	0.29(0.19)	0.19(0.14)	0.19(0.065)	0.12(0.10)	0.52(0.23)	0.42(0.17)	0.29(0.44)
		PD	9.86	8.16	0.35(0.20)	0.64(0.43)	0.12(0.066)	0.22(0.16)	0.28(0.17)	0.32(0.13)	0.54(0.25)
	Ep	PS	15.91	18.07	0.41(0.42)	0.30(0.44)	0.99(1.32)	0.26(0.35)	0.41(0.18)	0.28(0.10)	0.22(0.12)
		PD	16.64	14.34	0.36(0.20)	0.47(0.33)	0.077(0.19)	0.26(0.11)	0.24(0.11)	0.53(0.15)	0.60(0.30)
P5	En	RS	25.60	23.29	0.051(0.030)	0.11(0.15)	0.04(0.004)	0.041(0.0057)	0.043(0.007)	0.041(0.006)	0.040(0.0046)
		RD	40.03	36.99	0.31(0.13)	0.23(0.12)	0.23(0.057)	0.19(0.092)	0.32(0.18)	0.57(0.087)	0.30(0.24)
	Ep	RS	41.82	42.11	0.72(0.18)	1.00(0.27)	0.91(0.12)	0.88(0.17)	0.43(0.19)	0.4(0.20)	0.69(0.16)
		RD	55.54	50.66	0.61(0.35)	0.95(0.32)	0.84(0.88)	0.43(0.29)	0.2(0.12)	0.64(0.21)	0.59(0.28)

Table 1.1: **Global and regional error.** The global and regional error between the clinician and tracker segmentation. The regional area is based on the 16-segment model. This measures the regional closeness of the two segmentations. P1 = patient 1, P2 = patient 2 etc. Patients one to four had a dobutamine stress echocardiography. Patient five is that used throughout this thesis. RS = Rest systole, RD = Rest diastole, PS = Peak systole, PD = Peak diastole, EN = endocardium, EP = epicardium and / = no results were calculated because only part of the boundary was delineated by the clinician.

		Regional Division (%)						
P1	En	RS	95.51	100	7.15	22.99	52.44	39.95
		RD	100	100	11.76	4.41	38.52	95.21
	Ep	RS	0.30	0	0	0	0	60.38
		RD	53.46	0	0	0	25.50	93.18
	En	PS	89.44	74.51	93.19	93.82	0	81.70
		PD	99.73	38.71	13.67	100	100	100
	Ep	PS	0.30	0	0	0	0	60.38
		PD	22.73	0	88.18	46.71	89.44	100
P2	En	RS	100	100	83.78	3.90	5.95	28.74
		RD	100	100	96.06	0	0	57.95
	Ep	RS	99.91	100	22.53	0	0	75.29
		RD	99.91	100	48	0	0	79.52
	En	PS	100	100	70.06	0	97.39	26.56
		PD	98.04	100	76.24	0	68.29	88.36
	Ep	PS	/	/	/	/	/	/
		PD	/	/	/	/	/	/
P3	En	RS	/	/	/	/	/	/
		RD	100	100	31.74	99.94	100	100
	Ep	RS	/	/	/	/	/	/
		RD	65.31	0	31	100	100	100
	En	PS	/	/	/	/	/	/
		PD	/	/	/	/	/	/
	Ep	PS	/	/	/	/	/	/
		PD	/	/	/	/	/	/
P4	En	RS	99.43	100	30.60	0	0	98.63
		RD	98.63	93.33	4.42	2.37	100	100
	Ep	RS	96.56	100	22.63	0	0	96.07
		RD	/	/	/	/	/	/
	En	PS	53.43	100	25.51	0	0	1.35
		PD	86.01	85.12	5.71	0.36	100	100
	Ep	PS	69.14	14.20	2.19	0	0	24.31
		PD	95.08	0	0	5.65	100	100
P5	En	RS	53.03	48.39	2.06	57.73	100	89.84
		RD	90.35	100	23.37	19.84	100	98.49
	Ep	RS	81.11	0	0	0	10.43	18.20
		RD	93.38	0	0.97	15.24	100	99.82

Table 1.2: **Regional division.** The regional division between the the clinician and tracker segmentation. The regional division is the percentage of the regional error that is from the tracker segmentation being outside the clinician's segmentation. High values indicate that the tracker is consistently outside the clinician's segmentation. This measures whether the tracker segmentation is biased is one direction. P1 = patient 1, P2 = patient 2 etc. Patients one to four had a dobutamine stress echocardiography. Patient five is that used throughout this thesis. RS = Rest systole, RD = Rest diastole, PS = Peak systole, PD = Peak diastole, EN = endocardium, EP = epicardium and / = no results were calculated because only part of the boundary was delineated by the clinician.

		Global Rotation (°)	Regional Overlap (%)						
P1	En	RS	1.57	78.24	62.99	44.15	38.09	52.38	70.70
		RD	1.92	70.03	54.50	45.89	54.34	63.01	76.08
	Ep	RS	2.14	59.43	46.29	34.42	44.76	63.94	76.94
		RD	1.44	68.58	47.79	33.48	36.82	47.76	64.65
	En	PS	6.15	66.64	49.98	37.73	46.59	66.56	82.44
		PD	1.52	72.72	83.16	87.86	67.87	69.06	63.51
	Ep	PS	2.49	69.66	74.37	76.74	86.62	80.31	78.93
		PD	11.82	55.38	68.28	61.39	52.98	54.72	56.86
P2	En	RS	5.72	44.69	37.79	37.47	40.92	47.94	52.50
		RD	3.94	38.57	39.45	47.90	36.50	40.99	42.67
	Ep	RS	6.40	37.40	33.53	38.95	35.92	37.23	33.98
		RD	6.76	34.86	33.88	44.31	36.29	36.11	29.70
	En	PS	0.62	57.45	50.16	42.68	49.53	59.29	69.03
		PD	2.51	31.74	29.04	26.90	49.89	60.15	70.45
	Ep	PS	/	/	/	/	/	/	/
		PD	/	/	/	/	/	/	/
P3	En	RS	/	/	/	/	/	/	/
		RD	0.07	47.14	61.01	76.24	54.99	47.40	37.52
	Ep	RS	/	/	/	/	/	/	/
		RD	0.81	42.06	49.75	38.37	47.83	60.61	48.48
	En	PS	/	/	/	/	/	/	/
		PD	/	/	/	/	/	/	/
	Ep	PS	/	/	/	/	/	/	/
		PD	/	/	/	/	/	/	/
P4	En	RS	0.30	14.08	24.15	33.56	31.15	27.47	24.67
		RD	4.05	17.34	26.80	33.69	26.63	21.27	20.70
	Ep	RS	3.54	34.82	37.24	42.19	46.10	48.55	46.09
		RD	/	/	/	/	/	/	/
	En	PS	0.67	65.19	60.38	55.98	54.39	60.90	67.54
		PD	3.09	37.03	38.14	42.60	41.39	39.03	37.21
	Ep	PS	0.56	60.05	66.41	74.14	77.83	71.48	63.65
		PD	0.69	48.42	52.65	60.29	47.65	43.49	38.80
P5	En	RS	2.68	65.46	67.32	62.64	49.60	58.86	62.97
		RD	2.70	69.29	68.19	73.84	73.71	71.04	65.99
	Ep	RS	2.11	55.19	52.14	48.15	47.18	48.53	52.92
		PS	/	/	/	/	/	/	/
		RD	7.25	82.60	74.17	68.30	62.48	63.29	69.70

Table I.3: **Global rotation and regional overlap.** The global rotation and regional overlap between the clinician and tracker segmentation. The global rotation is the angular difference between the principal axis of the clinician and tracker segmentation. This measures the difference in rotation between the two segmentations. The regional percentage is the overlap between the clinician and tracker segmentation. This is a comparative measure of horizontal and vertical extent. P1 = patient 1, P2 = patient 2 etc. Patients one to four had a dobutamine stress echocardiography. Patient five is that used throughout this thesis. RS = Rest systole, RD = Rest diastole, PS = Peak systole, PD = Peak diastole, EN = endocardium, EP = epicardium and / = no results were calculated because only part of the boundary was delineated by the clinician.

---

## Bibliography

---

- Afridi, I., Quiñones, M. A., Zoghbi, W. A. and Cheirif, J. (1994) Dobutamine stress echocardiography: sensitivity, specificity, and predictive value for future cardiac events. *American Heart Journal* **127**(6), 1510–1515.
- Ahmad, T., Taylor, C. J., Lanitis, A. and Cootes, T. F. (1995) Tracking and recognising hand gestures using statistical shape models. In *Proc. British Machine Vision Conf.*, ed. D. Pycock, pp. 403–412. BMVA Press.
- Aisen, A. M. and Buda, A. J. (1985) Magnetic resonance imaging of the heart. In *Digital cardiac imaging*, eds A. J. Buda and E. J. Delp, Chapter 15, pp. 271–292. Boston, Lancaster: Nijhoff.
- Aldrich, H. R. and Reichek, N. (1993) Stress echocardiography. *Current Opinion in Cardiology* **8**, 978–987.
- Amico, A. F., Lichtenberg, G. S., Reisner, S. A., Stone, C. K., Schwartz, R. G. and Meltzer, R. S. (1989) Superiority of visual versus computerized echocardiographic estimation of radionuclide left ventricular ejection fraction. *American Heart Journal* **118**(6), 1259–1265.
- Amini, A., Chen, Y., Sun, J. and Mani, V. (1998a) SPAMM-MRI with constrained elastic splines: implementation and validation. In *1st International Conference on Medical Image Computing and Computer Assisted Intervention*, pp. 167–175. Cambridge, MA: Springer.
- Amini, A., Curwen, R., Constable, R. T. and Gore, J. (1994) MR physics-based snake tracking and dense deformations from tagged cardiac images. In *AAAI Symposium on Applications of Computer Vision to Medical Image Processing*.
- Amini, A., Radeva, P., Elayyadi, M. and Li, D. (1998b) Measurement of 3D motion of myocardial material points from explicit B-surface reconstruction of tagged MRI data. In *1st International Conference on Medical Image Computing and Computer Assisted Intervention*, pp. 110–129. Cambridge, MA: Springer.
- Amit, Y., Grenander, U. and Piccioni, M. (1991) Structural image restoration through deformable templates. *Journal of the American Statistical Association* **84**(414), 376–387.
- Armstrong, W. F. (1997) Treadmill exercise echocardiography: methodology and clinical role. *European Heart Journal* **Supplement D**, 2–8.
- Ashton, E. A. and Parker, K. J. (1995) Multiple resolution Bayesian segmentation of ultrasound images. *Ultrasonic imaging* **17**, 291–304.
- Astrom, K. and Wittenmark, B. (1984) *Computer Controlled Systems*. Addison Wesley.
- Attenhofer, C. H., Pellikka, P. A., Oh, J. K. and Roger, V. L. (1994) Comparison of frame grabbed cine loop images and videotape record in stress echocardiography — a prospective study. *Circulation* **90**(4 Pt2), 391.

- Ayache, N., Cohen, I. and Herlin, I. (1992) Medical image tracking. In *Active Vision*, eds A. Blake and A. Yuille, pp. 285–302. MIT.
- Azhari, H., Sideman, S., Beyar, R., Grenadier, E. and Dinnar, U. (1987) An analytical descriptor of three-dimensional geometry: application to the analysis of the left ventricle shape and contraction. *IEEE Trans. on Biomedical Engineering* **34**(5), 345–355.
- Bamber, J. C. (1993) Speckle reduction. In *Advances in Ultrasound Techniques and Instrumentation*, ed. P. N. T. Wells, pp. 55–67. Churchill Livingstone, NY.
- Bamber, J. C. and Daft, C. (1986) Adaptive filtering for reduction of speckle in ultrasound pulse-echo images. *Ultrasonics* pp. 41–44.
- Bar-Shalom, Y. and Fortmann, T. (1988) *Tracking and Data Association*. Academic Press.
- Bardinet, E., Cohen, L. and Ayache, N. (1996) Tracking and motion analysis of the left ventricle with deformable superquadrics. *Medical Image Analysis* **1**(2), 129–149.
- Barnett, S. (1990) *Matrices: Methods and Applications*. Oxford University Press.
- Bartels, R., Beatty, J. and Barsky, B. (1987) *An Introduction to Splines for use in Computer Graphics and Geometric Modeling*. Morgan Kaufmann.
- Bascle, B. and Blake, A. (1998) Separability of pose and expression in facial tracking and animation. In *Proc. 6th Int. Conf. on Computer Vision*, pp. 323–328. Bombay, India.
- Bascle, B., Bouthemy, P., Deriche, R. and Meyer, F. (1994) Tracking complex primitives in an image sequence. Technical Report 2428, INRIA.
- Bascle, B. and Deriche, R. (1994) Region tracking through image sequences. Technical Report 2439, INRIA.
- Basu, S., Oliver, N. and Pentland, A. (1998) 3D modelling and tracking of human lip motions. In *Proc. 6th Int. Conf. on Computer Vision*, pp. 337–343. Bombay, India.
- Baumberg, A. and Hogg, D. (1994a) An efficient method for contour tracking using active shape models. In *IEEE Workshop on motion of non-rigid and articulate objects*. Austin, Texas.
- Baumberg, A. and Hogg, D. (1994b) Learning flexible models from image sequences. In *Proc. 3rd European Conf. Computer Vision*, ed. J.-O. Eklundh, pp. 299–308. Springer-Verlag.
- Baumberg, A. and Hogg, D. (1995) Generating spatio-temporal models from examples. In *Proc. British Machine Vision Conf.*, volume 2, pp. 413–422.
- Baxley, W. A. and Reeves, T. J. (1971) Abnormal regional myocardial performance in coronary artery disease. *Progress in cardiovascular diseases* **13**, 405–421.
- Behrenbruch, C., Brady, M., Jacob, G. and Noble, J. A. (1999) Adaptive multiresolution detection of the epicardial wall in 2D cardiac ultrasound. In *Third British Conference on Medical Image Understanding and Analysis*, pp. 89–92.
- Beleslin, B. D., Ostojic, M., Stepanovic, J., Djordjevic-Dikic, A., Stojkovic, S., Nedeljkovic, M., Stankovic, G., Petrasinovic, Z., Gojkovic, L., Vasiljevic-Pokrajcic, Z. and Nedeljkovic, S. (1994) Stress echocardiography in the detection of myocardial ischemia. Head-to-head comparison of exercise, dobutamine and dipyridamole tests. *Circulation* **90**(3), 1169–1176.
- Berne, R. M. and Levy, M. N. (1986) *Cardiovascular system*. 6th Edition. Mosby Year Book.
- Besag, J. E. (1986) On the statistical analysis of dirty pictures (with discussion). *Journal Royal Stat. Soc. Ser. B* **48**, 259–302.
- Besag, J. E. (1989) Towards Bayesian image analysis. *Journal of Applied Statistics* **16**, 395–407.
- Besag, J. E., Green, P. J., Higdon, D. and Mengersen, K. (1995) Bayesian computation and stochastic systems. *Statistical Science* **10**, 3–66.

- Black, M. and Anandan, P. (1993) A framework for the robust estimation of optical flow. In *Proc. 4th Int. Conf. on Computer Vision*, pp. 231–236. Berlin, Germany.
- Blake, A., Curwen, R. and Zisserman, A. (1993) A framework for spatio-temporal control in the tracking of visual contours. *Int. J. Computer Vision* **11**, 127–145.
- Blake, A. and Isard, M. (1994) 3D position, attitude and shape input using video tracking of hands and lips. In *Proc. Siggraph*, pp. 185–192.
- Blake, A. and Isard, M. (1998) *Active contours*. Springer.
- Blake, A., Isard, M. and Reynard, D. (1995) Learning to track the visual motion of contours. *Artificial Intelligence* **78**, 179–212.
- Bookstein, F. L. (1989) Principal warps: thin plate splines and the decomposition of deformations. *IEEE Trans. on Pattern Analysis and Machine Intelligence* **11**(6), 567–585.
- Boukerroui, D., Basset, O., Guirin, N. and Baskurt, A. (1998) Multiresolution texture based adaptive clustering algorithm for breast lesion segmentation. *European Journal of Ultrasound* **8**(2), 135–144.
- Bourdillon, P. D. V., Broderick, T. W., Sawada, S. G., Armstrong, W. F., Ryan, T., Dillon, J. C., Fineberg, N. S. and Feigenbaum, H. (1989) Regional wall motion index for infarct and non infarct regions after perfusion in acute myocardial infarction: comparison with global wall motion index. *Journal of the American Society of Echocardiography* **2**, 398–407.
- Bregler, C. and Omohundro, S. M. (1994) Surface learning with applications to lipreading. In *Advances in Neural Information Processing Systems*, eds J. D. Cowan, G. Tesauro and J. Alspector, volume 6, pp. 1–2. Morgan Kaufmann Publishers.
- British Heart Foundation (1992) The heart. Technical terms explained.
- Bruce, R. A. and Hornstein, T. R. (1969) Exercise stress testing in evaluation of patients with ischemic heart disease. *Progress in cardiovascular diseases* **11**, 371–390.
- Brundage, B. H. (1985) Clinical applications of cardiac CT. In *Digital cardiac imaging*, eds A. J. Buda and E. J. Delp, Chapter 10, pp. 140–156. Boston, Lancaster: Nijhoff.
- Bucy, R. S. (1969) Bayes theorem and digital realizations for non-linear filters. *J. Astronautical Sciences* **17**(2), 80–94.
- Buda, A. J. and Delp, E. J. (1985) *Digital cardiac imaging*. Boston, Lancaster: Nijhoff.
- Burckhardt, C. B. (1978) Speckle in ultrasound B-scans. *IEEE Trans. on Sonics and Ultrasonics* **1**, 1–6.
- Byrne, N. J., Baumberg, A. M. and Hogg, D. C. (1994) Using shape and intensity to track non-rigid objects. Technical Report 94.14, School of Computer Studies, University of Leeds, Leeds, UK.
- Canny, J. F. (1986) A computational approach to edge detection. *IEEE Transactions on Pattern Analysis and Machine Intelligence* **8**, 679–698.
- Carlson, R. E., Kavanaugh, K. M. and Buda, A. J. (1988) The effect of different mechanisms of myocardial ischemia on left ventricular motion. *American Heart Journal* **116**, 536–545.
- Cauce, A. and Taylor, C. J. (1998) 3D point distribution models of the cortical sulci. In *Proc. 6th Int. Conf. on Computer Vision*, pp. 402–407. Bombay, India.
- Chalana, V. and Kim, Y. (1997) A methodology for evaluation of boundary detection algorithms on medical images. *IEEE Trans. Medical Imaging* **16**(5), 643–652.
- Chalana, V., Linker, D. T., Haynor, D. R. and Kim, Y. (1996) A multiple active contour model for cardiac boundary detection on echocardiographic sequences. *IEEE Trans. Medical Imaging* **15**(3), 290–298.

- Chambolle, A., DeVore, R. A., Lee, N. Y. and Lucier, B. J. (1998) Nonlinear wavelet image processing: variational problems, compression and noise removal through wavelet shrinkage. *IEEE Trans. Medical Imaging* **7**(3), 319–335.
- Chen, C., Huang, T. S. and Arrott, M. (1994) Modelling, analysis and visualization of left ventricle shape and motion by hierarchical decomposition. *IEEE Trans. on Pattern Analysis and Machine Intelligence* **16**(4), 342–356.
- Chen, C. G., Chen, L. L., Hall, E., Dyckman, W., Waters, D. and Gillam, L. (1995) The responses of stunned and of hibernating myocardium to dobutamine differ — implications for dobutamine stress echocardiography. *Circulation* **92**(8 SS), 1269. Meeting Abstract.
- Chen, F. and Suter, D. (1997) Left ventricular motion reconstruction based on elastic vector splines. *Submitted to Special Issue on Deformable Models and Techniques for Image and Signal Processing*.
- Chu, C. H., Delp, E. J. and Buda, A. J. (1988) Detecting left ventricular endocardial and epicardial boundaries by digital two-dimensional echocardiography. *IEEE Trans. Medical Imaging* **7**(2), 81–90.
- Cipolla, R. and Blake, A. (1990) The dynamic analysis of apparent contours. In *Proc. 3rd Int. Conf. on Computer Vision*, pp. 616–625.
- Clarysse, P., Friboulet, D. and Magnin, I. E. (1997) Tracking geometrical descriptors on 3D deformable surfaces: application to the left-ventricular surface of the heart. *IEEE Trans. Medical Imaging* **16**(4), 393–404.
- Clifford, P. (1989) Markov random fields in statistics. *Journal of Applied Statistics* **16**, 19–32.
- Cohen, L. and Cohen, I. (1990) A finite element method applied to new active contour models and 3D reconstruction from cross sections. In *Proc. 3rd Int. Conf. on Computer Vision*, pp. 587–591.
- Coifman, R. R. and Wickerhauser, M. V. (1992) Entropy based algorithms for best basis selection. *IEEE Trans. Inf. Theor.* **32**, 712–771.
- Cokkinos, P., Kurbaan, A. S. and Nihoyannopoulos, P. (1998) Dobutamine stress echocardiography. *The British Journal of Cardiology* **5**(5), 276–285.
- Cootes, T. F., Hill, A., Taylor, C. J. and Haslam, J. (1994) The use of active shape models for locating structures in medical images. *Image and Vision Computing* **12**(6), 355–366.
- Cootes, T. F. and Taylor, C. J. (1992) Active shape models — ‘smart snakes’. In *Proc. British Machine Vision Conf.*, pp. 266–275. Springer-Verlag.
- Cootes, T. F. and Taylor, C. J. (1993a) Active shape model search using local grey level models: a quantitative evaluation. In *Proc. British Machine Vision Conf.*, ed. J. Illingworth, pp. 639–648. BMVA Press.
- Cootes, T. F. and Taylor, C. J. (1993b) Active shape model search using local grey-level models. In *BMVC*, ed. J. Illingworth, pp. 639–648. BMVA Press.
- Cootes, T. F. and Taylor, C. J. (1995) Combining point distribution models with shape models based on finite element analysis. *Journal of Image and Vision Computing* **13**(5), 403–410.
- Cootes, T. F. and Taylor, C. J. (1997) A mixture model for representing shape variation. In *Proc. British Machine Vision Conf.*, ed. A. F. Clark, volume 1, pp. 110–119. Colchester, Essex: BMVA Press.
- Cootes, T. F., Taylor, C. J., Cooper, D. H. and Graham, J. (1995) Active shape models — their training and application. *Computer Vision and Image Understanding* **61**(1), 38–59.
- Cootes, T. F., Taylor, C. J., Cooper, D. H. and Graham, P. (1992) Training models of shape from sets of examples. In *Proc. British Machine Vision Conf.*, pp. 9–18. Springer-Verlag.
- Cootes, T. F., Taylor, C. J., Lanitis, A., Cooper, D. H. and Graham, P. (1993) Building and using flexible models incorporating grey-level information. In *Proc. 4th Int. Conf. on Computer Vision*, pp. 242–246. Berlin, Germany.

- Coppini, G., Poli, R. and Valli, G. (1995) Recovery of the 3D shape of the left ventricle from echocardiographic images. *IEEE Trans. Medical Imaging* **14**(2), 301–317.
- Corya, B. C., Rasmussen, S., Fiegenbaum, H., Knoebel, S. B. and Black, M. J. (1977) Systolic thickening and thinning of the septum and posterior wall in patients with coronary artery disease, congestive cardiomyopathy. *Circulation* **55**, 109–114.
- Corya, B. C., Rasmussen, S., Knoebel, S. B. and Feigenbaum, H. (1975) Echocardiography in acute myocardial infarction. *The American Journal of Cardiology* **36**, 1–10.
- Cox, I. J. (1993) A review of statistical data association techniques for motion correspondence. *Int. J. Computer Vision* **10**(1), 53–66.
- Crowder, M. J. and Hand, D. J. (1990) *Analysis of repeated measures*. Chapman & Hall.
- Curwen, R. and Blake, A. (1992) Dynamic contours: real-time active splines. In *Active Vision*, eds A. Blake and A. L. Yuille, pp. 39–58. MIT.
- Curwen, R., Blake, A. and Cipolla, R. (1991) Parallel implementation of Lagrangian dynamics for real-time snakes. In *Proc. British Machine Vision Conf.*, pp. 29–36.
- Curwen, R. W. M. (1993) *Dynamic and adaptive contours*. Ph.D. thesis, University of Oxford.
- Dalton, B., Kaucic, R. and Blake, A. (1995) Automatic speechreading using dynamic contours. In *Proceedings NATO ASI Conference on Speechreading by Man and Machine: Models, Systems and Applications*.
- Daubechies, I. (1988) Orthonormal bases of compactly supported wavelets. In *Communications on Pure and Applied Mathematics XLI*, pp. 909–996.
- Daubechies, I. (1994) Ten lectures on wavelets. In *CMBS, SIAM*, pp. 258–261.
- DeCarlo, D. and Metaxas, D. (1998) Deformable model-based shape motion analysis from images using motion residual error. In *Proc. 6th Int. Conf. on Computer Vision*, pp. 113–119. Bombay, India.
- Declerck, J., Feldmar, J. and Ayache, N. (1997) Definition of a 4D continuous polar transformation for the tracking and the analysis of LV motion. *Lecture notes in computer science* **1205**, 33–42.
- Declerck, J., Feldmar, J. and Ayache, N. (1998) Definition of a four-dimensional continuous planispheric transformation for the tracking and the analysis of left-ventricle motion. *Medical Image Analysis* **2**(2), 197–213.
- Demer, L. L., Gould, L., Goldstein, R., Kirkeeide, R. L., Mullani, N. A., Smalling, R. W., Nishikawa, A. and Merhige, M. E. (1989) Assessment of coronary artery disease severity by positron emission tomography. comparison with quantitative arteriography in 193 patients. *Circulation* **79**(4), 825–835.
- Deriche, R. (1987) Using Canny's criteria to derive a recursively implemented optimal edge detector. *Int. J. Computer Vision* **1**(2), 167–187.
- Dias, J. M. B. and Leitão, J. M. N. (1996) Wall position and thickness estimation from sequences of echocardiographic images. *IEEE Trans. Medical Imaging* **15**(1), 25–38.
- Dickinson, R. J. (1982) Reduction of speckle in ultrasound B-scans by digital processing. In *Acoustic Imaging*, eds E. A. Ash and C. R. Hill, pp. 213–224. Plenum Press, New York.
- Edwards, G. J. and Taylor, C. J. (1998) Learning to identify and track faces in image sequences. In *Proc. 6th Int. Conf. on Computer Vision*, pp. 317–322. Bombay, India.
- Elhendy, A., Geleijnse, M. L., vanDomburg, R. T., Nierop, P. R., Bax, J. J., ElSaid, G. M., Ibrahim, M. M. and Nosir, Y. F. M. (1998) Optimal criteria for the diagnosis of coronary artery disease by dobutamine stress echocardiography. One or more ischemic segments? *Journal of the American College of Cardiology* **31**(2 SA), A40. Meeting Abstract.

- Epstein, S. E. (1976) Editorial: Whither the ST segment during exercise? *Circulation* **54**(703–706).
- Evans, A. N. and Nixon, M. S. (1996) Biased motion-adaptive temporal filtering for speckle reduction in echocardiography. *IEEE Trans. Medical Imaging* **15**(1), 39–50.
- Falseti, H. L., Marcus, M. L., Kerber, R. E. and Skorton, D. J. (1981) Editorial: Quantification of myocardial ischemia and infarction by left ventricular imaging. *Circulation* **63**(4), 747–751.
- Feigenbaum, H. (1994) *Echocardiography, 5th edition*. Philadelphia, PA: Lea & Febiger.
- Ferrier, N. J., Rowe, S. M. and Blake, A. (1994) Real-time traffic monitoring. In *Proc. 2nd IEEE Workshop on Applications of Computer Vision*, pp. 81–88.
- Folland, E. D., Parisi, A. F., Moynihan, P. F., Jones, P. F., Feldman, D. R. and Tow, D. E. (1979) Assessment of left ventricular ejection fraction and volumes by real-time, two-dimensional echocardiography. *Circulation* **60**, 760–767.
- Friedland, N. and Adam, D. (1989) Automatic ventricular cavity boundary detection from sequential ultrasound images using simulated annealing. *IEEE Trans. Medical Imaging* **8**(4), 344–353.
- Gelb, A. (1974) *Applied Optimal Estimation*. MIT Press.
- Gelberg, H. J., Brundage, B. H., Glantz, S. and Parmley, W. W. (1979) Quantitative left ventricular wall motion analysis: a comparison of area, chord and radial methods. *Circulation* **59**(5), 991–1000.
- Geleijnse, M. L., Fioretti, P. M. and Roelandt, J. R. T. C. (1997) Methodology, feasibility, safety and diagnostic accuracy of dobutamine stress echocardiography. *Journal of the American College of Cardiology* **30**(3), 595–606.
- Gelfand, A. E. and Smith, A. F. M. (1990) Sampling-based approaches to calculating marginal densities. *J. Amer. Statist. Assoc* **85**, 398–409.
- Geman, D. and Geman, S. (1984) Stochastic relaxation, Gibbs distributions, and the Bayesian restoration of images. *IEEE Trans. on Pattern Analysis and Machine Intelligence* **6**, 721–741.
- Germano, G., Kiat, H., Kavanagh, P. B., Moriel, M., Mazzantin, M., Su, H. T., Vantraine, K. F. and Berman, D. S. (1995) Automatic quantification of ejection fraction from gated myocardial perfusion spect. *Journal of Nuclear Medicine* **36**(11), 2138–2147.
- Giachetti, A. (1998) On-line analysis of echocardiographic image sequence. *Medical Image Analysis* **2**(3), 261–284.
- Giachetti, A., Gigli, G. and Torre, V. (1995) Computer assisted analysis of echocardiographic image sequences. In *CVRMed*, pp. 267–271. Nice, France.
- Godoy, I. E., Mor-Avi, V., Spencer, K. T. and Lang, R. M. (1997) Objective echocardiographic evaluation of the cardiovascular system: state of the art. *Current opinion in Cardiology* **12**, 553–560.
- Goldberg, S. J. (1984) Analysis and interpretation of thickening and thinning phases of left ventricular wall dynamics. *Ultrasound in Medicine and Biology* **10**(6), 797–802.
- Goldberg, S. J., Feldman, L., Reinecke, C., Stern, L. Z., Sahn, D. J. and Allen, H. D. (1980) Echocardiographic determination of contraction and relaxation of the left ventricular wall in normal subjects and patients with muscular dystrophy. *Circulation* **62**(5), 1061–1069.
- Gorce, J., Friboulet, D. and Magnin, I. E. (1997) Estimation of three-dimensional cardiac velocity fields: assessment of a differential method and application to three-dimensional CT data. *Medical Image Analysis* **1**(3), 245–261.
- Gordan, E. P., Schnittger, I., Fitzgerald, P. J., Williams, P. and Popp, R. L. (1983) Reproducibility of left ventricular volumes by two-dimensional echocardiography. *Journal of the American College of Cardiology* **2**(3), 506–513.

- Graffigne, C. (1986) *Application des statistiques au traitement d'images*. Ph.D. thesis, Universite de Paris-Sud, Centre d'Orsay.
- Green, P. J. and Silverman, B. W. (1994) *Nonparametric regression and generalized linear models. A roughness penalty approach*. Chapman & Hall.
- Grenander, U. (1983) Tutorial in pattern theory. Technical report, Division of Applied Mathematics, Brown University, Providence.
- Grenander, U., Chow, Y. and Keenan, D. M. (1991) *HANDS: A Pattern theoretic study of biological shapes*. Springer, New York.
- Grenander, U. and Miller, M. (1994) Representations of knowledge in complex systems (with discussion). *Journal Royal Stat. Soc. Ser. B* **56**, 549–603.
- Grover, M. and Schelbert, H. R. (1985) Positron emission computed tomography. In *Digital cardiac imaging*, eds A. J. Buda and E. J. Delp, Chapter 14, pp. 240–270. Boston, Lancaster: Nijhoff.
- Haber, E., Metaxas, D. N. and Axel, L. (1998) Motion analysis of the right ventricle from MRI images. In *1st International Conference on Medical Image Computing and Computer Assisted Intervention*, pp. 177–188. Cambridge, MA: Springer.
- Han, C. Y., Kwun, N. L., Wee, W. G., Minitz, R. M. and Porembka, D. T. (1991) Knowledge based image analysis for automated boundary enhancement of transoesophageal echocardiographic left ventricular images. *IEEE Trans. Medical Imaging* **10**(4), 602–610.
- Haque, T., Furukawa, T., Takahashi, M. and Kinoshita, M. (1995) Identification of hibernating myocardium by dobutamine stress echocardiography: comparison with thallium-201 reinjection imaging. *American Heart Journal* **130**(3), 553–563.
- Harbert, J. and Da Rocha, A. F. G. (1984) *Textbook of nuclear medicine Volume II. Clinical Applications*. Volume 2. Philadelphia: Lea & Febiger.
- Härdle, W. (1990) *Applied nonparametric regression*. CUP.
- Harris, C. (1992) Geometry from visual motion. In *Active Vision*, eds A. Blake and A. L. Yuille, pp. 263–284. MIT.
- Harvey, A. C. (1989) *Forecasting, structural time series models and the Kalman filter*. CUP.
- Heap, A. and Hogg, D. (1996) Extending the point distribution model using polar coordinates. *Journal of Image and Vision Computing* **14**(8), 589–599.
- Heap, A. and Hogg, D. C. (1997) Improving specificity in PDMs using a hierarchical approach. In *Proc. British Machine Vision Conf.*, ed. A. F. Clark, volume 1, pp. 80–89. Colchester, Essex: BMVA Press.
- Heap, T. and Hogg, D. (1998) Wormholes in shape space: tracking through discontinuous change in space. In *Proc. 6th Int. Conf. on Computer Vision*, pp. 344–349. Bombay, India.
- Hecht, H. S., DeBord, L., Shaw, R., Chin, H., Dunlap, R., Ryan, C. and Myler, R. K. (1993) Supine bicycle stress echocardiography versus tomographic thallium-201 exercise imaging for the detection of coronary artery disease. *Journal of the American Society of Cardiology* **6**(2), 177–185.
- Heikkilä, J., Tabakin, B. S. and Hugenholtz, P. G. (1972) Quantification of function in normal and infarcted regions of the left ventricle. *Cardiovascular research* **6**, 516–531.
- Hennessy, T. G., Codd, M. B., Hennessy, M. S., Kane, G., McCarthy, C., McCann, H. A. and Sugrue, D. D. (1997) Comparison of dobutamine stress echocardiography and treadmill exercise electrocardiography for detection of coronary artery disease. *Coronary artery disease* **8**(11/12), 689–695.
- Herlin, I. and Ayache, N. (1992) Feature extraction and analysis methods for sequences of ultrasound images. In *Proc. 2nd European Conf. Computer Vision*, pp. 43–57.

- Herlin, I., Bereziat, D., Giraudon, G., Nguyen, C. and Graffigne, C. (1994) Segmentation of echocardiographic images with Markov random fields. In *Proc. European Conf. on Computer Vision*, ed. J. Eklundh, volume 801, pp. 201–206. Springer-Verlag.
- Herlin, I. and Giraudon, G. (1993) Performing segmentation of ultrasound images using temporal information. In *Computer vision and Pattern recognition*, pp. 373–378.
- Herlin, I., Nguyen, C. and Graffigne, C. (1992a) A deformable region model using stochastic processes applied to echocardiographic images. In *Proc. Conf. Computer Vision and Pattern Recognition*.
- Herlin, I., Nguyen, C. and Graffigne, C. (1992b) Stochastic segmentation of ultrasound images. In *Proc. 11th Int. Conf. on Pattern Recognition*.
- Hill, A., Cootes, T. F., Taylor, C. J. and Lindley, K. (1994) Medical image interpretation: a generic approach using deformable templates. *Journal of Medical Informatics* **19**(1), 47–59.
- Hoffmann, R., Lethen, H., Kleinhans, E., Weiss, M., Flachskampf, F. and Hanrath, P. (1993) Comparative evaluation of bicycle and dobutamine stress echocardiography with perfusion scintigraphy and bicycle electrocardiogram for identification of coronary artery disease. *The American Journal of Cardiology* **72**, 555–559.
- Hole JR., J. W. and Koos, K. A. (1994) *Human Anatomy*. Second Edition. Wm. C. Brown.
- Horn, B. K. P. (1986) *Robot Vision*. NY: McGraw-Hill.
- Hotelling, H. (1933) Analysis of a complex of statistical variables into principal components. *Journal of educational psychology* **24**, 417–441, 498–520.
- Huang, P. J., Ho, Y. L., Wu, C. C., Chao, C. L., Chen, M. F., Chieng, P. U. and Lee, Y. (1997) Simultaneous dobutamine stress echocardiography and thallium-201 perfusion imaging for the detection of coronary artery disease. *Cardiology* **88**(6), 556–562.
- Intille, S. S. and Bobick, A. F. (1995) Visual tracking using closed-worlds. In *Proc. 5th Int. Conf. on Computer Vision*, pp. 672–678. MIT, Cambridge, MA.
- Irvins, J. and Porrill, J. (1995) Active regional models for segmenting textures and colours. *Image and Vision Computing* **13**(5), 431–438.
- Isard, M. (1998) *Visual Motion Analysis by Probabilistic Propagation of Conditional Density*. Ph.D. thesis, University of Oxford.
- Isard, M. and Blake, A. (1996) Visual tracking by stochastic propagation of conditional density. In *Proc. 4th European Conf. Computer Vision*, pp. 343–356. Cambridge, England.
- Isard, M. and Blake, A. (1998) A mixed state condensation tracker with automatic model-switching. In *Proc. 6th Int. Conf. on Computer Vision*, pp. 107–112. Bombay, India.
- Jacob, G., Noble, J. A. and Blake, A. (1998) Robust contour tracking in echocardiographic sequences. In *Proc. 6th Int. Conf. on Computer Vision*, pp. 408–413. Bombay, India.
- Jacob, G., Noble, J. A., Mulet-Parada, M. and Blake, A. (1999) Evaluating a robust contour tracker on echocardiographic sequences. *Medical Image Analysis* **3**(1), 63–75.
- Jacobs, J. J., Feigenbaum, H., Corya, B. C. and Phillips, J. R. (1973) Detection of left ventricular asynergy by echocardiography. *Circulation* **48**, 263–271.
- Jebara, T., Russell, K. and Pentland, A. (1998) Mixtures of eigenfeatures for real-time structure from texture. In *Proc. 6th Int. Conf. on Computer Vision*, pp. 128–135. Bombay, India.
- Jolliffe, I. T. (1986) *Principal component analysis*. New York: Springer-Verlag.
- Julian, D. G. and Cowan, J. C. (1992) *Cardiology*. 6th Edition. Baillière Tindall.
- Juni, J. E. and Buda, A. J. (1985) Radionuclide imaging of the heart. In *Digital cardiac imaging*, eds A. J. Buda and E. J. Delp, Chapter 13, pp. 207–239. Boston, Lancaster: Nijhoff.

- Kan, G., Visser, C. A., Lie, K. I. and Durrer, D. (1981) Left ventricular volumes and ejection fraction by single plane two-dimensional apex echocardiography. *European Heart Journal* **2**(4), 339–343.
- Karrer, H. E. and Dickey, A. M. (1983) Ultrasound imaging: an overview. In *Hewlett-Packard Journal*.
- Kass, M., Witkin, A. and Terzopoulos, D. (1987) Snakes: active contour models. In *Proc. 1st Int. Conf. on Computer Vision*, pp. 259–268.
- Katz, A. S., Force, T. L., Folland, T. L., Aebischer, N., Sharma, A. and Parisi, A. F. (1996) Echocardiographic assessment of ventricular systolic function. In *Marcus cardiac imaging: a companion to Braunwald's Heart disease*, eds M. L. Marcus, D. J. Skorton, G. L. Wolf and B. H. Brundage. Philadelphia, London: W.B. Saunders.
- Kaucic, R. and Blake, A. (1998) Accurate, real-time, unadorned lip tracking. In *Proc. 6th Int. Conf. on Computer Vision*, pp. 370–375. Bombay, India.
- Kaucic, R., Dalton, B. and Blake, A. (1996) Real-time liptracking for audio-visual speech recognition applications. In *Proc. 4th European Conf. Computer Vision*, pp. 376–387. Cambridge, England.
- Kaucic, R. A. (1997) *Lip Tracking for audio-visual speech recognition*. Ph.D. thesis, University of Oxford.
- Kerber, R. E. and Abboud, F. M. (1973) Echocardiographic detection of regional myocardial infarction. *Circulation* **47**, 997–1005.
- Kerber, R. E., Marcus, M. L., Ehrhardt, J., Wilson, R. and Abboud, F. M. (1975) Correlation between echocardiographically demonstrated dyskinesia and regional myocardial perfusion. *Circulation* **52**, 1097–1104.
- Kerber, R. E., Marcus, M. L., Ehrhardt, J., Wilson, R. and Abboud, F. M. (1976) Effect of acute coronary occlusion on the motion and perfusion of the normal and ischemic interventricular septum: an experimental echocardiographic study. *Circulation* **55**, 928–935.
- Kerber, R. E., Martins, J. B. and Marcus, M. L. (1979) Effects of acute ischemia, nitroglycerin and nitroprusside on regional myocardial thickening, stress and perfusion. *Circulation* **60**, 121–129.
- Keren, A., DeAnda, A., Komeda, M., Tye, T., Handen, C. R., Daughters, G. T., Ingels, N. B., Miller, C., Popp, R. L. and Nikolic, S. D. (1995) Pitfalls in creation of left atrial pressure-area relationships with automatic border detection. *Journal of American Society of Echocardiography* **8**, 669–678.
- Kerwin, W. S. and Prince, J. L. (1998) Cardiac material markers from tagged MR images. *Medical Image Analysis* **2**(4), 339–353.
- Kimball, T. R., Witt, S. A., Daniels, S. R., Khoury, P. R. and Meyer, R. A. (1996) The role of automatic endocardial edge detection in the evaluation of left ventricular diastolic function in children. *Journal of the American Society of Echocardiography* **9**(1), 18–26.
- Koch, R., Mor-Avi, V., Garcia, M. J., Bednarz, J., Weinert, L., Korcarz, C., Sandelski, J. and Lang, R. M. (1996) Automated interpretation of dobutamine stress echocardiography using segmental analysis of color kinesis images. *Circulation* **94**(8 Supplement 1), 680.
- Koenderink, J. and van Doorn, A. (1975) Invariant properties of the motion parallax field due to the movement of rigid bodies relative to an observer. *Optica Acta* **22**(9), 773–791.
- Koenderink, J. J. and van Doorn, A. J. (1991) Affine structure from motion. *J. Optical Soc. of America A* **8**(2), 337–385.
- Kotler, M. N. and Steiner, R. M. (1986) *Cardiac imaging: new technologies and clinical applications*. Philadelphia: F.A. Davis.

- Kovesi, P. (1996) *Invariant measures of image features from phase information*. Ph.D. thesis, Department of Psychology, U. West. Australia.
- Krahwinkel, W., Ketteler, T., Gödke, J., Wolfertz, J., Ulbricht, L. J., Krakau, I. and Gülker, H. (1997) Dobutamine stress echocardiography. *European Heart Journal* **18**(Supplement D), 9-15.
- Krauss, D., Feldman, T., Marcus, R., Resnekov, L. and Land, R. (1994) Paradoxical imaging in regional wall motion during high dose dobutamine stress echocardiography. A marker for hibernating myocardium? *Chest* **106**, 291-293.
- Kremkau, F. W. (1993) *Diagnostic Ultrasound: Principles and Instruments*. 4th edition. W. B. Saunders Company, Philadelphia, PA, USA.
- Krzanowski, W. (1988) *Principles of Multivariate analysis: A users perspective*. OUP.
- Lang, R. M., Vignon, P., Weinert, L., Bednarz, J., Korcarz, C., Sandelski, J., Koch, R., Prater, D. and Mor-Avi, V. (1996) Echocardiographic quantification of regional left ventricular wall motion with color kinesis. *Circulation* **93**(10), 1877-1885.
- Lange, A., Moran, C. M., Palka, P., Fenn, L., Sutherland, G. and McDicken, W. N. (1995) The variation in integrated backscatter in human hearts in differing ultrasonic transthoracic views. *Journal of American Society of Echocardiography* **8**(6), 830-838.
- Leavitt, S. C., Hunt, B. F. and Larsen, H. G. (1983) A scan conversion algorithm for displaying ultrasound images. In *Hewlett-Packard Journal*.
- Lee, D., Kent, J. T., Mardia, K. V. and de Souza, K. M. A. (1998) Using a Bayesian approach on a network of quads to track tagged cardiac MR images. In *Second British Conference on Medical Image Understanding and Analysis*, pp. 1-4. Leeds, UK.
- Levick, J. R. (1995) *An introduction to cardiovascular physiology*. Volume Second. Butterworth Heinemann.
- Lieberman, A. N., Weiss, J. L., Jugdutt, B. I., Becker, L. C., Bulkley, B. H., Garrison, J. G., Hutchins, G. M., Kallman, C. A. and Weisfeldt, M. L. (1981) Two-dimensional echocardiography and infarct size: relationship of regional wall motion and thickening to the extent of myocardial infarction in the dog. *Circulation* **63**(4), 739-746.
- Limacher, M. C., Quinones, M. A., Poliner, L. R., Nelson, J. G., Winters, W. L. and Waggoner, A. D. (1983) Detection of coronary artery disease with exercise two-dimensional echocardiography. *Circulation* **67**(6), 1211-1218.
- Lipson, P., Yuille, A. L., O'Keefe, D., Cavanaugh, J., Taaffe, J. and Rosenthal, D. (1990a) Deformable templates for feature extraction from medical images. Technical Report 89-14, Harvard Robotics Laboratory.
- Lipson, P., Yuille, A. L., O'Keefe, D., Cavanaugh, J., Taaffe, J. and Rosenthal, D. (1990b) Deformable templates for feature extraction from medical images. In *Proc. 1st European Conf. Computer Vision*, ed. O. F. Ed., pp. 413-417. Springer-Verlag.
- Lowe, D. (1992) Robust model-based motion tracking through the integration of search and estimation. *Int. J. Computer Vision* **8**(2), 113-122.
- Magnus, J. R. and Neudecker, H. (1988) *Matrix differential calculus with applications in statistics and economics*. Wiley.
- Mailloux, G. E., Bleau, A., Bertrand, M. and Petitclerc, R. (1987) Computer analysis of heart motion from two-dimensional echocardiograms. *IEEE Trans. on Biomedical Engineering* **34**(5), 356-364.
- Mann, D. L., Gillam, L. D. and Weyman, A. E. (1986) Cross-sectional echocardiographic assessment of regional ventricular performance and myocardial perfusion. *Progress in cardiovascular diseases* **29**, 1-52.
- Marais, P. C. (1998) *The segmentation of sparse MR images*. Ph.D. thesis, Department of Engineering Science, Oxford, UK.

- Marcovitz, P. A. and Armstrong, W. F. (1992) Accuracy of dobutamine stress echocardiography in detecting coronary artery disease. *The American Journal of Cardiology* **69**, 1269–1273.
- Marcus, M. L. and Braunwald, E. (1991) *Cardiac imaging: a companion to Braunwald's Heart disease*. Philadelphia, London: Saunders.
- Marriott, F. H. C. (1974) *The interpretation of multiple observations*. London: Academic Press.
- Marwick, T. H. (1996) *Cardiac stress testing & imaging: a clinician's guide*. New York, Edinburgh: Churchill Livingstone.
- Marwick, T. H. (1997) Adenosine echocardiography in the diagnosis of coronary artery disease. *European Heart Journal Supplement D*, 31–36.
- Marwick, T. H., Nemecek, J. J., Pashkow, F. J., Stewart, W. J. and Salcedo, E. E. (1992) Accuracy and limitations of exercise echocardiography in a routine clinical setting. *Journal of the American College of Cardiology* **19**(1), 74–81.
- Mason, J. R., Palac, R. T., Freeman, M. L., Virupannavar, S., Loeb, H. S., Kaplan, E. and Gunnar, R. M. (1984) Thallium scintigraphy during dobutamine infusion: nonexercise-dependent screening test for coronary disease. *American Heart Journal* **107**(3), 481–485.
- Mazeika, P. K., Nadazdin, A. and Oakley, C. (1992) Dobutamine stress echocardiography for detection and assessment of coronary artery disease. *Journal of the American College of Cardiology* **19**(6), 1203–1211.
- McEachen II, J. C. and Duncan, J. S. (1997) Shape-based tracking of left ventricular wall motion. *IEEE Trans. Medical Imaging* **16**(3), 271–283.
- McEachen II, J. C., Meyer, G. G., Constable, R., Nehorai, A. and Duncan, J. S. (1995) A recursive filter for phase velocity assisted shape-based tracking of cardiac non-rigid motion. In *Proc. 5th Int. Conf. on Computer Vision*. MIT, Cambridge, MA.
- McEachen II, J. C., Nehorai, A. and Duncan, J. S. (1994) Analysis of cardiac motion with recursive comb filtering. In *Proc. 1994 SPIE Mathematical Methods in Medical Imaging III*, pp. 46–57. San Diego, CA.
- McHenry, P., Phillips, J. and Knoebel, S. (1972) Correlation of computer-quantitated treadmill exercise electrocardiogram with arteriographic location of coronary artery disease. *The American Journal of Cardiology* **30**, 747–52.
- McInerney, T. and Terzopoulos, D. (1996) Deformable models in medical image analysis. In *IEEE/SLAM Workshop on Mathematical Methods in Biomedical Image Analysis*, pp. 171–180.
- McKiddie, F. I., Gemmell, H. G., Redpath, T. W., Trent, R. J., Philip, W. J. U., Norton, M. Y. and Smith, F. W. (1996) A variable threshold edge-detector for improved quantification of gated imaging of the left ventricular pool. *Nuclear Medicine Communications* **17**, 410–417.
- McVeigh, E. R. (1996) MRI of myocardial function: motion tracking techniques. *Magnetic Resonance Imaging* **14**(2), 137–150.
- Menet, S., Saint-Marc, P. and Medioni, G. (1990) B-snakes: implementation and application to stereo. In *Proc. DARPA*, pp. 720–726.
- Mertes, H., Sawada, S. G., Ryan, T., Segar, D. S., Kovacs, R. and Feigenbaum, H. (1993) Symptoms, adverse effects and complications associated with dobutamine stress echocardiography — experience in 118 patients. *Circulation* **88**(1), 15–19.
- Meyer, F., Constable, R., Sinusas, A. and Duncan, J. (1996) Tracking myocardial deformation using phase contrast MR velocity fields: a stochastic approach. *IEEE Trans. Medical Imaging* **15**(4), 453–465.
- Mor-Avi, V., Gillesberg, I. E., Jorcarz, C., Sandelski, J. and Lang, R. M. (1995) Improved quantification of left ventricular function by applying signal averaging to echocardiographic acoustic quantification. *Journal of American Society of Echocardiography* **8**, 679–689.

- Mor-Avi, V., Vignon, P., Koch, R., Weinert, L., Garcia, M. J., Spencer, K. T. and Lang, R. M. (1997) Segmental analysis of color kinesis images. *Circulation* **95**(8), 2082–2097.
- Moran, B. A., Leonard, J. J. and Chrysostomidis, C. (1997) Curved shape reconstruction using multiple hypothesis tracking. *IEEE Journal of Oceanic Engineering* **22**(4), 625–638.
- Moran, C. M., Sutherland, G. R., Anderson, T., Riemersma, R. A. and McDicken, W. N. (1994) A comparison of methods used to calculate ultrasonic myocardial backscatter in the time domain. *Ultrasound in Medicine and Biology* **20**(6), 543–550.
- Morganroth, J., Parisi, Alfred, F. and Pohost, G. M. (1983) *Noninvasive cardiac imaging*. Chicago: Year Book Medical Publishers.
- Mueller, X., Stauffer, J. C., Jaussi, A., Goy, J. J. and Kappenberger, L. (1991) Subjective visual echocardiographic of left ventricular ejection fraction as an alternative to conventional echocardiographic methods. *Clinical cardiology* **14**(11), 898–907.
- Mulet-Parada, M. (1999) *Phase-based intensity invariant feature detection for echocardiography (in preparation)*. Ph.D. thesis, Department of Engineering Science, University of Oxford, Oxford.
- Mulet-Parada, M. and Noble, J. A. (1998) Intensity-invariant 2D+T acoustic boundary detection. In *IEEE Workshop in Biomedical Image Analysis*, pp. 133–142. Santa Barbara, California, USA.
- Mulet-Parada, M. and Noble, J. A. (1999) 2D+T acoustic boundary detection in echocardiography. *Medical Image Analysis to appear*.
- Mundy, J. L. and Zisserman, A. (1992) *Geometric invariance in computer vision*. Cambridge, Mass; London: MIT.
- Nagel, H. H. and Enkelmann, W. (1986) An investigation of smoothness constraints for the estimation of displacement vector fields from image sequences. *IEEE Trans. on Pattern Analysis and Machine Intelligence* **8**(5), 565–593.
- Nalwa, V. (1993) *A Guided Tour of Computer Vision*, Chapter 1, pp. 3–25. Addison-Wesley.
- Nastar, C. and Ayache, N. (1993) Fast segmentation, tracking and analysis of deformable objects. In *Proc. 4th Int. Conf. on Computer Vision*, pp. 275–279. Berlin, Germany.
- Nastar, C., Moghaddam, B. and Pentland, A. (1996) Generalized image matching: statistical learning of physically-based deformations. In *Proc. 4th European Conf. Computer Vision*, pp. 589–598.
- Nastar, C. and Pentland, A. (1995) Matching and recognition using deformable intensity surfaces. In *IEEE International Symposium on Computer Vision*, pp. 223–228. Miami Beach, Florida.
- Navaratnam, V. (1975) *The human heart and circulation*. London: Academic press.
- Noble, J. (1988) Finding corners. *J. Image and Vision Computing* **6**(2), 121–128.
- North, B. (1998) *Learning Dynamical Models for Visual Tracking*. Ph.D. thesis, University of Oxford.
- North, B. and Blake, A. (1998) Learning dynamical models using expectation-maximisation. In *Proc. 6th Int. Conf. on Computer Vision*, pp. 384–389. Bombay, India: .
- Owen, A. N., Melucci, M., Adeleke, K., Oliner, C. M., Kotler, M. N. and Jacobs, L. E. (1998) Use of atropine early during dobutamine stress echocardiography. *Journal of the American College of Cardiology* **31**(2 SA), A41. Meeting Abstract.
- Panidis, J. P. (1996) *Cardiac Ultrasound*. Blackwell Science.
- Paragios and Deriche, R. (1998) A PDE-based level set approach for detection and tracking of moving objects. In *Proc. 6th Int. Conf. on Computer Vision*, pp. 1139–1145. Bombay, India.
- Parisi, A. F., Moynihan, P. F. and Feldman, C. L. (1981) Quantitative detection of regional left ventricular contraction abnormalities by two-dimensional echocardiography. *Circulation* **63**, 761–767.

- Park, J., Metaxas, D. and Young, A. (1994) Deformable models with parameter functions: application to heart-wall modelling. In *Computer Vision and Pattern Recognition*, pp. 437–442. Seattle, Washington.
- Park, J., Metaxas, D. and Young, A. (1996a) Analysis of left ventricular wall motion based on deformable models and MRI-SPAMM. *Medical Image Analysis* **1**(1), 53–71.
- Park, J., Metaxas, D., Young, A. A. and Axel, L. (1996b) Deformable models with parameter functions for cardiac motion analysis from tagged MRI data. *IEEE Trans. Medical Imaging* **15**(3), 278–289.
- Parker, A. D., Hill, A., Taylor, C. J., Cootes, T. F., Jin, X. Y. and Gibson, D. G. (1994) Application of point distribution models to the automated analysis of echocardiograms. In *Proc. of Computers in Cardiology*, pp. 25–28. IEEE Computer Society Press.
- Parker, J. A., Secker-Walker, R. and R., H. (1972) A new technique for the calculation of left ventricular ejection fraction. *Nuclear Medicine* **13**, 649–651.
- Pearson, K. (1901) On lines and planes of closest fit to systems of points in space. *Phil Mag* **2**(6), 559–572.
- Pentland, A. and Horowitz, B. (1991) Recovery of nonrigid motion and structure. *IEEE Trans. on Pattern Analysis and Machine Intelligence* **13**(7), 730–742.
- Pentland, A. P. and Sclaroff, S. (1991) Closed-form solutions for physically-based shape modelling and recognition. *IEEE Trans. on Pattern Analysis and Machine Intelligence* **13**(7), 715–729.
- Peréz, J. E., Klein, S. C., Prater, D. M., Fraser, C. E., Cardona, H., Waggoner, A. D., Holland, M. R., Miller, J. G. and E., S. B. (1992a) Automated, on-line quantification of left ventricular dimensions and function by echocardiography with backscatter imaging and lateral gain compensation. *The American Journal of Cardiology* **70**, 1200–1205.
- Peréz, J. E., Lang, R. M. and Mor-Avi, V. (1997) *Echocardiography and cardiovascular function: Tools for the next decade*. Dordrecht: Kluwer Academic publishers.
- Peréz, J. E., Waggoner, A. D., B., B., Melton, . H. E., Miller, J. G. and Sobel, B. E. (1992b) On-line assessment of ventricular function by automatic boundary detection and ultrasonic backscatter imaging. *Journal of the American College of Cardiology* **19**(2), 313–320.
- Peréz, J. E., Waggoner, A. D., Dávila-Roman, V. G., Cardona, H. and Miller, J. G. (1992c) On-line quantification of ventricular function during dobutamine stress echocardiography. *European Heart Journal* **13**, 1669–1676.
- Perona, P. and Malik, J. (1990) Scale space and edge detection using anisotropic diffusion. *IEEE Trans. on Pattern Analysis and Machine Intelligence* **12**(7), 629–639.
- Phillips, D. B. and Smith, A. F. M. (1994) Bayesian faces via hierarchical template modelling. *Journal of American Statistical Association* **89**(428), 1151–1163.
- Picano, E. and Lattanzi, F. (1991) Dipyridamole echocardiography. *Circulation* **83**(5), III–19.
- Picano, E., Lattanzi, F., Masini, M. and Labbate, A. (1987) Comparison of two-dimensional echocardiography for diagnosis of coronary artery disease. *The American Journal of Cardiology* **59**(6), 539–542.
- Picano, E., Lattanzi, F., Orlandini, A., Marini, C. and L'abbate, A. (1991) Stress echocardiography and the human factor: the importance of being expert. *Journal of the American College of Cardiology* **17**(3), 666–669.
- Pozen, R. G., Dibianco, R., Katz, R. J., Bortz, R., Myerburg, R. J. and Fletcher, R. D. (1981) Myocardial metabolic and hemo-dynamic effects of dobutamine in heart failure complicating coronary artery disease. *Circulation* **63**(6), 1279–1284.

- Previtali, M., Lanzarini, L., Ferrario, M., Tortorici, M., Mussini, A. and Montemartini, C. (1991) Dobutamine versus dipyridamole echocardiography in coronary artery disease. *Circulation* **83**(5), III-27.
- Previtali, M., Lanzarini, L., Fétiqueau, R., Poli, A., Ferrario, M., Falcone, C. and Mussini, A. (1993) Comparison of dobutamine stress echocardiography, dipyridamole stress echocardiography and exercise stress testing for diagnosis of coronary artery disease. *Coronary artery disease* **72**(12), 865-870.
- Quiñones, M. A., Verani, M. S., M., H. R., Mahmarian, J. J., Suarez, J. and Zoghbi, W. A. (1992) Exercise echocardiography versus  $^{201}\text{TI}$  single-photon emission computed tomography in evaluation of coronary artery disease. *Circulation* **75**(3), 1026-1031.
- Rallidis, L., Cokkinos, P., Tousoulis, D. and Nihoyannopoulos, P. (1997) Comparison of dobutamine and treadmill exercise echocardiography in inducing ischemia in patients with coronary artery disease. *Journal of the American College of Cardiology* **30**(7), 1660-1668.
- Rao, C. R. (1973) *Linear Statistical Inference and Its Applications*. John Wiley and Sons.
- Reynard, D., Blake, A., Azzawi, A., Styles, P. and Radda, G. K. (1994) Cardiac motion analysis using computer tracking of Tagged H images. In *Cardiovascular MRI, Present and Future*. New Mexico, USA.
- Reynard, D., Wildenberg, A., Blake, A. and Marchant, J. (1996) Learning dynamics of complex motions from image sequences. In *Proc. 4th European Conf. Computer Vision*, pp. 357-368. Cambridge, England.
- Rifkin, R. D. and Koito, H. (1990) Comparison with radionuclide angiography of two geometric and four nongeometric models of echocardiographic estimation of left ventricular ejection fraction using segmental wall scoring. *The American Journal of Cardiology* **65**, 1485-1490.
- Rijsterborgh, H., Mastik, F., T., L. C., Verdouw, P., Roelandt, J. and Bom, N. (1993) Ultrasound myocardial integrated backscatter signal processing: frequency domain versus time domain. *Ultrasound in Medicine and Biology* **19**(3), 211-219.
- Ripley, B. D. (1987) *Stochastic Simulation*. Wiley, New York.
- Ripley, B. D. (1988) *Statistical inference for spatial processes: an essay awarded the Adams Prize of the University of Cambridge*. Cambridge University Press, Cambridge.
- Ripley, B. D. (1992) Classification and clustering in spatial and image data. In *Analyzing and modelling data and knowledge*, ed. M. Schader, pp. 93-10. Springer.
- Ripley, B. D. and Sutherland, A. I. (1990) Finding spiral structures in images of galaxies. *Phil. Trans. Roy. Soc. London Ser. A* **332**, 377-384.
- Rogers, E. W., Feigenbaum, H. and Weyman, A. E. (1979) Echocardiography for quantification of cardiac chambers. In *Progress in cardiology*, eds P. N. Yu and J. F. Goodwin, volume 8. Lea & Febiger.
- Rogers, G. S. (1980) *Matrix derivatives*. New York: Marcel Dekker.
- Rohling, R. N., Gee, A. H. and Berman, L. (1997) Spatial compounding of 3D ultrasound images. *Medical Image Analysis* **1**(3), 177-193.
- Rose, J. H., Kaufmann, M. R., Wickline, S. A., Hall, C. S. and Miller, J. G. (1995) A proposed microscopic elastic wave theory for ultrasonic backscatter from myocardial tissue. *Journal of Acoustical Society America* **97**(1), 656-667.
- Rowe, S. (1996) *Robust feature search for active tracking*. Ph.D. thesis, University of Oxford.
- Rowe, S. and Blake, A. (1996) Statistical feature modelling for active contours. In *Proc. 4th European Conf. Computer Vision*, pp. 560-569. Cambridge, England.
- Rueckert, D., Burger, P., Forbat, S. M., Mahiaddin, R. D. and Yang, G. Z. (1997) Automatic tracking of the aorta in cardiovascular MR images using deformable models. *IEEE Trans. Medical Imaging* **16**(5), 581-590.

- Ruff, C. F., Highes, S. W. and Hawkes, D. J. (1997) Volume estimation from sparse images using deformable models. In *Proc. British Machine Vision Conf.*, pp. 90–99. Colchester, Essex: BMVA Press.
- Ryan, T. (1991) Dobutamine stress echocardiography. *Coronary Artery Disease* **2**(5), 552–558.
- Ryan, T., Vasey, C. G., Presti, C. F., O'Donnell, J. A., Feigenbaum, H. and Armstrong, W. F. (1988) Exercise echocardiography: detection of coronary artery disease in patients with normal left ventricular wall motion at rest. *Journal of the American College of Cardiology* **11**(5), 993–999.
- Sanchez-Ortiz, G. I., Noble, J. A., Wright, G. J. T., Feldmar, J. and Mulet-Parada, M. (1999a) Automated LV motion analysis from 3D echocardiography. In *Third British Conference on Medical Image Understanding and Analysis*, pp. 85–88. Oxford, UK.
- Sanchez-Ortiz, G. I., Rueckert, D. and Burger, P. (1999b) Knowledge-based tensor anisotropic diffusion of cardiac magnetic resonance images. *Medical Image Analysis* **3**(1), 77–101.
- Sasayama, S., Franklin, D., Ross, J. R., Kemper, W. S. and McKown, D. (1976) Dynamic changes in left ventricular wall thickness and their use in analyzing cardiac function in the conscious dog. *The American Journal of Cardiology* **38**, 870–879.
- Sawada, A. G., Segar, D. S., Ryan, T., Brown, S. E., Dohan, A. M., Williams, R., S., F. N., Armstrong, W. F. and H., F. (1991) Echocardiographic detection of coronary artery disease during dobutamine infusion. *Circulation* **83**(5), 1605–1614.
- Schiller, N. B. (1985) Echocardiography: current status and future prospectus. *Int. J. Cardiac Imaging* **1**, 41–60.
- Schiller, N. B., Shah, P. M., Crawford, M., DeMaria, A., Devereux, R., Feigenbaum, H., Guttesell, H., Reichek, N., Sahn, D., Schnittger, I., Silverman, N. H. and Tajik, A. J. (1989) American society of echocardiography committee on standards, subcommittee on quantitation of two-dimensional echocardiograms: recommendations for quantitation of the left ventricle by two-dimensional echocardiography. *Journal of the American Society of Echocardiography* **2**, 358–367.
- Schistad, A. and Taxt, T. (1991) Speckle reduction in ultrasound images using temporal or spatial context. In *Proc IEEE Nuclear Science symposium and Medical imaging conference*, pp. 2210–2219.
- Schnittger, I., Fitzgerald, P. J., Alderman, E. L. and Popp, R. L. (1984) Computerized quantitative analysis of left ventricular wall motion by two-dimensional echocardiography. *Circulation* **70**, 242–254.
- Scaroff, S. and Isidoro, J. (1998) Active blobs. In *Proc. 6th Int. Conf. on Computer Vision*, pp. 1146–1153. Bombay, India.
- Scaroff, S. and Pentland, A. P. (1994) On modal modelling for medical images: Underconstrained shape description and data compression. In *Proc. of the IEEE workshop on biomedical image analysis*, pp. 70–79. Seattle.
- Secknus, M. A. and Marwick, T. H. (1997) Evolution of dobutamine echocardiography protocols and indications: safety and side effects in 3,011 studies over 5 years. *Journal of the American College of Cardiology* **29**(6), 1234–1240.
- Segar, D. S., E., B. S., Sawada, S. G., Ryan, T. and Feigenbaum, H. (1992) Dobutamine stress echocardiography: correlation with coronary lesion severity as determined by quantitative angiography. *Journal of the American College of Cardiology* **19**(6), 1197–1202.
- Senior, R. and Lahiri, A. (1995) Enhanced detection of myocardial ischemia by stress dobutamine echocardiography utilizing the “biphasic” response of wall motion thickening during low and high dose dobutamine infusion. *Journal of the American College of Cardiology* **26**(1), 26–32.
- Sheehan, F. H., Bolson, E. L., Dodge, H. T., Mathey, D. G., Schoffer, J. and Woo, H. (1986) Advantages and applications of the centerline method for characterizing regional ventricular function. *Circulation* **74**(2), 293–305.

- Sheehan, F. H., Bolson, E. L., Martin, R. W., Bashein, G. and McDonald, J. (1998) Quantitative three-dimensional echocardiography: methodology, validation and clinical applications. In *1st International Conference on Medical Image Computing and Computer Assisted Intervention*, pp. 102–109. Cambridge, MA: Springer.
- Shen, X. and Hogg, D. (1994) Shape models from image sequences. In *Proc. 4th European Conf. Computer Vision*, ed. J. Eklundh, volume 1, pp. 225–230. Springer-Verlag.
- Shi, P., Robinson, G. and Duncan, J. (1994) Myocardial motion and function assessment using 4D images. In *Proceedings of IEEE conference on Visualization and Biomedical computing*. Rochester MN.
- Skorton, D. J., Brundage, B. H., Schelbert, H. R. and Wolf, G. L. (1996) Goals of cardiac imaging. In *Marcus cardiac imaging: a companion to Braunwald's Heart disease*, eds M. L. Marcus, D. J. Skorton, G. L. Wolf and B. H. Brundage. Philadelphia, London: W.B. Saunders.
- Smith, A. F. M. and Roberts, G. O. (1993) Bayesian computation via the Gibbs sampler and related Markov chain monte carlo methods. *Journal Royal Stat. Soc. Ser. B* **55**(1), 3–23.
- Sonka, M., Zhang, X., Siebes, M., Bissing, M., Dejong, S. C., Collins, S. M. and McKay, C. R. (1995) Segmentation of intravascular ultrasound images — a knowledge-based approach. *IEEE Trans. Medical Imaging* **14**(4), 719–732.
- Sozou, P., Cootes, T. F., Taylor, C. J. and DiMauro, E. (1995) Non-linear point distribution modelling using a multi-layer perceptron. In *Proc. British Machine Vision Conf.*, ed. D. Pycock. Birmingham: BMVA Press.
- Stamm, R. B., Gibson, R. S., Bishop, H. L., Carabello, B. A., Beller, G. A. and Martin, R. P. (1983) Echocardiographic detection of infarct-localized asynergy and remote asynergy during acute myocardial infarction: correlation with the extent of angiographic coronary disease. *Circulation* **67**(1), 233–244.
- Storvik, G. (1994) A Bayesian approach to dynamic contours through stochastic sampling and simulated annealing. *IEEE Trans. on Pattern Analysis and Machine Intelligence* **16**(10), 976–986.
- Stuber, M., Nagel, E., Fischer, S. E., Spiegel, M. A. and Scheidegger, M. B. (1998) Quantification of the local heartwall motion by magnetic resonance tagging. *Computerized Medical Imaging and Graphics* **22**(3), 217–228.
- Sullivan, G. D. (1992) Visual interpretation of known objects in constrained scenes. *Phil. Trans. R. Soc. Lond. B.* **337**, 361–370.
- Székely, G., Kelemen, A., Brechbüler, C. and Gerig, G. (1996) Segmentation of 2D and 3D objects from MRI volume data using constrained elastic deformations of flexible fourier contour and surface models. *Medical Image Analysis* **1**(1), 19–34. •
- Szeliski, R. and Terzopoulos, D. (1991) Physically-based and probabilistic models for computer vision. In *SPIE Procs. Geometric methods in computer vision*, ed. B. C. Vemuri, volume 1570, pp. 140–152. San Diego, CA: Society of Photo-Optical Instrumentation Engineers.
- Tained, M. C., Herment, A., Diebold, B. and Peronneau, P. (1994) Segmentation of cardiac and vascular ultrasound images with extension to border kinetics. In *Ultrasonics symposium*, pp. 1773–1776.
- Tanner, M. A. (1991) *Tools for statistical Inference*. Volume 67. Lecture notes in Statistics.
- Tennant, R. and Wiggers, C. J. (1935) The effect of coronary occlusion on myocardial contraction. *American Journal of Physiology* **112**, 351–361.
- Terzopoulos, D. and Metaxas, D. (1992) Tracking nonrigid 3D objects. In *Active Vision*, eds A. Blake and A. L. Yuille, pp. 75–89. MIT.
- Terzopoulos, D. and Szeliski, R. (1992) Tracking with Kalman snakes. In *Active Vision*, eds A. Blake and A. L. Yuille, pp. 3–20. MIT.

- Terzopoulos, D., Witkin, A. and Kass, M. (1991) Constraints on deformable models: recovering 3D shape and nonrigid motion. *Artificial Intelligence* **36**(1), 91–123.
- Thomas III, L. J., Barzilai, B., Pérez, J. E., Sobel, B. E., Wickline, S. A. and Miller, J. G. (1989) Quantitative real-time imaging of myocardium based on ultrasonic integrated backscatter. *IEEE Trans. on Ultrasonics, Ferroelectrics and frequency control* **36**(4), 466–470.
- Thomas III, L. J., Wickline, S. A., Pérez, J. E., Sobel, B. E. and Miller, J. G. (1986) A real-time integrated backscatter measurement system for quantitative cardiac tissue characterization. *IEEE Trans. on Ultrasonics, Ferroelectrics and frequency control* **33**(1), 27–32.
- Torry, R. J., Myers, J. H., Adler, A. L., Liut, C. L. and Gallagher, K. P. (1991) Effects of nontransmural ischemia on inner and outer wall thickening in the canine left ventricle. *American Heart Journal* **122**(5), 1292–1299.
- Turk, M. and Pentland, A. P. (1991) Eigenfaces for recognition. *Cognitive Neuroscience* **3**(1), 71–96.
- Tuttle, R. R. and Mills, J. (1975) Dobutamine: development of a new catecholamine to selectively increase cardiac response. *Circulation research* **36**, 185–196.
- Venkatesh, S. and Owens, R. (1990) On the classification of image features. *Pattern Recognition Letters* **11**, 339–349.
- Vered, M. D., Barzilai, B., Mohr, G. A., Thomas III, L. J., Genton, R., Sobel, B. E., Shroup, T. A., Melton, H. E., Miller, J. G. and Pérez, J. E. (1987) Quantitative ultrasonic tissue characterization with real-time integrated backscatter imaging in normal human subjects and in patients with dilated cardiomyopathy. *Circulation* **76**(5), 1067–1073.
- Vignon, P., Mor-Avi, V., Weinert, L., Koch, R., Spencer, K. T. and Lang, R. M. (1998) Quantitative evaluation of global and regional left ventricular diastolic function with color kinesis. *Circulation* **97**, 1053–1061.
- Vitarelli, A., Luzzi, M. F., Penco, M. and Fedele, F. Dagianti, A. (1997) On-line quantitative assessment of left ventricular filling during dobutamine stress echocardiography: a useful addition to conventional wall motion scoring. *International Journal of cardiology* **59**, 57–69.
- Vitarelli, A., Sciomer, S., Penco, M., Dagianti, D. and Pugliese, M. (1998) Assessment of left ventricular dyssynergy by color kinesis. *The American Journal of Cardiology* **81**(12A), 87G–90G. June.
- Wagner, R. F., Smith, S. W., Sandrik, J. M. and Lopez, H. (1983) Statistics of speckle in ultrasound B-scans. *IEEE Trans. on Sonics and Ultrasonics* **30**, 156–163.
- Wahba, G. (1990) *Spline models for observational data*. Volume 59 of *CBMS-NSF Regional Conference Series in Applied Mathematics*. Philadelphia, PA: SIAM.
- Weatherall, D. J., Ledingham, J. G. G. and Warrel, D. A. (1996) *Oxford textbook of medicine*. Third Edition, volume 2 of 11–17 and index. Oxford: Oxford medical publications. Oxford University press.
- Webb, S. (1988) *The physics of medical imaging*. Institute of Physics Publishing.
- Wells, P. N. T. (1993) The present status of ultrasonic imaging in medicine. *Ultrasonics* **31**, 345–352.
- Wells, P. N. T. and Halliwell, M. (1981) Speckle in ultrasonic imaging. *Ultrasonics* **19**, 225–229.
- Weng, J., Singh, A. and Chiu, M. Y. (1997) Learning-based ventricle detection from cardiac MR and CT images. *IEEE Trans. Medical Imaging* **16**(4), 378–391.
- Wildenberg, A. P. (1997) *Learning and initialisation for visual tracking*. Ph.D. thesis, University of Oxford.
- Wilson, D. C., Geiser, E. A. and Li, J. (1993) Feature extraction in two-dimensional short-axis echocardiographic images. *Journal of mathematical imaging and vision* **3**, 285–298.

- Winkler, G. (1995) *Image Analysis, random fields and dynamic Monte Carlo methods*. Springer.
- Wiske, P. S., Pearlman, J. D., Hogan, R. D., Franklin, T. D. and Weyman, A. E. (1990) Echocardiographic detection of the left ventricular centroid: I. echocardiographic definition of the left ventricular centroid. ii. determination of the optimal centroid during systole in normal and infarcted hearts. *Journal of the American College of Cardiology* **16**(4), 993–999.
- Wong, M., Bruce, S., Joseph, D. and Lively, H. (1991) Estimating left ventricular ejection fraction from two-dimensional echocardiograms: visual and computer-processed interpretations. *Echocardiography* **8**(1), 1–7.
- Woods, R. H. (1892) A few applications of a physical theorem to membranes in the human body in a state of tension. *J Anat. Physiol* **26**, 362–370.
- Yacoob, Y. and Black, M. J. (1998) Parametrized modelling and recognition of activities. In *Proc. 6th Int. Conf. on Computer Vision*, pp. 120–127. Bombay, India.
- Yacoob, Y. and Davis, L. (1998a) Learned temporal models of image motion. In *Proc. 6th Int. Conf. on Computer Vision*, pp. 446–453. Bombay, India.
- Yacoob, Y. and Davis, L. (1998b) Robust tracking with spatio velocity snakes: Kalman filtering approach. In *Proc. 6th Int. Conf. on Computer Vision*, pp. 433–439. Bombay, India.
- Young, A. A. (1998) Motion tags: direct 3D tracking of heart wall motion from tagged MR images. In *1st International Conference on Medical Image Computing and Computer Assisted Intervention*, pp. 92–101. Cambridge, MA: Springer.
- Yuille, A. L. and Hallinan, P. (1992) Deformable templates. In *Active Vision*, eds A. Blake and A. L. Yuille, pp. 20–38. MIT.
- Zeng, X., Staib, L. H., Schultz, R. T. and S., D. J. (1998) Segmentation and measurement of the cortex from 3D MR images. In *1st International Conference on Medical Image Computing and Computer Assisted Intervention*, pp. 519–530. Cambridge, MA, USA.
- Zhang, L. and Geiser, E. A. (1982) An approach to optimal threshold selection on a sequence of two-dimensional echocardiographic images. *IEEE Trans. on Biomedical Engineering* **29**(8), 577–581.
- Zhong, Y., Jain, A. K. and Dubuisson-Jolly, M. (1998) Object tracking using deformable templates. In *Proc. 6th Int. Conf. on Computer Vision*, pp. 440–445. Bombay, India.
- Zile, M. R., Tanaka, R., Lindroth, J. R., Spinale, F., Carabello, B. A. and Mirsky, I. (1992) Left ventricular volume determined echocardiography by assuming a constant left ventricular epicardial long-axis/short axis dimension ratio throughout the cardiac cycle. *Journal of the American College of Cardiology* **20**(4), 986–993.