

Three-dimensional geometric image analysis for interventional electrophysiology



John E McManigle Jr
Wolfson College
University of Oxford

A thesis submitted for the degree of
Doctor of Philosophy
Michaelmas 2014

Acknowledgements

Personal

I would first like to express my heartfelt appreciation for my Oxford advisor, Professor Alison Noble. From our first meeting through final thesis preparations, she encouraged my independent exploration, but was always ready with wisdom and support when it was most needed.

I would also like to thank my NIH and Johns Hopkins advisors, Dr. David Bluemke and Dr. Saman Nazarian. Their clinical perspective and enthusiastic counsel kept me from drifting too far afield. I greatly benefited from working with the whole Hopkins Electrophysiology team, especially Dr. Henry Halperin and Dr. Roy Beinart. Their early morning collegiality was refreshing, and their tutelage in electroanatomic mapping and MR physics never missed the mark.

My examiners, Prof. Paul Leeson and Dr. Kawal Rhode, have my sincere gratitude for a challenging *viva voce* and illuminating suggestions. Their input was invaluable in ensuring the rigour and quality of this thesis.

Work in the lab was made downright enjoyable by my exceptional comrades in the IBME, especially Susan Graham, Richard Stebbing, Phill Brown, Maryam Ahmed, Ana Namburete, and Michael Tee. We trudged through images and algorithms together, and are emerging a little wiser.

My enthusiastic thanks to everyone who made Wolfson College such a welcoming place, one cup of tea at a time, particularly Sonya Ziaja, Ingrid Cloete, Zoë Goodwin, Josh Curk, Laura Hawkins, and especially Sabine Parrish, who has truly mastered the ‘Oxford triathlon’ – rowing, croquet, and Foucault.

And I would be remiss not to acknowledge the friends who joined me on so many adventures from Baltimore: Eva Klinman, Jerome Molleston, Léa Savard-McNicoll, Rachel Poor, and Lisa Vienneau, whose marathon speed is matched only by the speed of her wit.

Last but by no means least, thanks to my amazing brother and sister, Bill and Lorne, and to my parents. Every year I see a little more how lucky I am to have this loving and supportive family. Their patience and encouragement – even when I’ve been too absorbed across the pond to call as often as I should – is humbling.

Institutional

I am extremely grateful for the generous support of the U.S. National Institutes of Health (NIH) Graduate Partnership Program (GPP). Various portions of my programme were funded by the National Institute of Biomedical Imaging and Bioengineering (NIBIB), the National Heart, Lung and Blood Institute (NHLBI), and the National Institute of Neurological Disorders and Stroke (NINDS).

The University of Oxford, and in particular Wolfson College, the Institute of Biomedical Engineering, and the Nuffield Department of Medicine, provided the academic framework and exhaustive administrative assistance throughout my degree, in spite of the inevitable complexity of a trans-Atlantic programme.

The John Radcliffe Hospital Cardiac Physiology Clinical Unit provided clinical ultrasound data and advice. The Johns Hopkins University Electrophysiology and Arrhythmia Service provided critical guidance in clinical applicability and a compelling vision of the future of interventional electrophysiology in addition to clinical MR, CT, and electroanatomic map data.

Finally, I am hugely appreciative of the Duke University School of Medicine and Medical Scientist Training Program, especially Prof. Chris Kontos, Prof. Philip Goodman, and Ms. Andrea Lanahan, who were unwavering in their moral and logistical support.

Technical

My algorithmic efforts were eased by open source programming toolkits, including:

- **scikit-learn**, an open-source Python machine learning library primarily supported by *Inventeurs du monde numérique*. It is available at www.scikit-learn.org.
- The Insight Segmentation and Registration Toolkit (ITK), an open source library developed as an initiative of the U.S. National Library of Medicine and available at www.itk.org.
- The Visualization Toolkit (VTK), an open source project developed by Will Schroeder, Ken Martin, and Bill Lorensen and now supported by Kitware, Inc. It is available at www.vtk.org.
- The multi-platform configuration tool CMake, used to configure ITK and VTK. Partially funded by the U.S. National Library of Medicine, CMake is available at www.cmake.org.

Abstract

Improving imaging hardware, computational power, and algorithmic design are driving advances in interventional medical imaging. We lay the groundwork here for more effective use of machine learning and image registration in clinical electrophysiology.

To achieve identification of atrial fibrosis using image data, we registered the electroanatomic map (EAM) data of atrial fibrillation (AF) patients undergoing pulmonary vein isolation (PVI) with MR ($n = 16$) or CT ($n = 18$) images. The relationship between image features and bipolar voltage was evaluated using single-parameter regression and random forest models. Random forest performed significantly better than regression, identifying fibrosis with area under the receiver operating characteristic curve (AUC) 0.746 (MR) and 0.977 (CT). This is the first evaluation of voltage prediction using image data.

Next, we compared the character of native atrial fibrosis with ablation scar in MR images. Fourteen AF patients undergoing repeat PVI were recruited. EAM data from their first PVI was registered to the MR images acquired before the first PVI ('pre-operative') and before the second PVI ('post-operative' with respect to the first PVI). Non-ablation map points had similar characteristics in the two images, while ablation points exhibited higher intensity and more heterogeneity in post-operative images. Ablation scar is more strongly enhancing and more heterogeneous than native fibrosis.

Finally, we addressed myocardial measurement in 3-D echocardiograms. The circular Hough transform was modified with a feature asymmetry filter, epicardial edges, and a search constraint. Manual and Hough measurements were compared in 5641 slices from 3-D images. The enhanced Hough algorithm was more accurate than the unmodified version (Dice coefficient 0.77 vs. 0.58). This method promises utility in segmentation-assisted cross-modality registration.

By improving the information that can be extracted from medical images and the ease with which that information can be accessed, this progress will contribute to the advancing integration of imaging in electrophysiology.

Contents

List of Figures	ix
List of Abbreviations	xii
1 Introduction	1
1.1 Motivation	1
1.2 Contribution	3
1.3 Thesis Structure	4
1.4 Originality and Individual Role	6
1.5 Prior Publication	8
2 Background	9
2.1 Introduction	9
2.2 Cardiac Imaging	10
2.2.1 Diagnostic Imaging	10
2.2.2 Interventional Imaging	18
2.3 Interventional Cardiology	22
2.4 Atrial Fibrillation	23
2.4.1 Atrial Remodelling	24
2.4.2 Sequelae	25
2.4.3 Treatment	25
3 Registration of Electroanatomic Map for Image Analysis	30
3.1 Introduction	31
3.1.1 Commercial Procedural Registration	32
3.1.2 Registration for Research	34
3.2 Experimental Data	35
3.3 Methods	36
3.3.1 Algorithm Input	38
3.3.2 Landmark Orientation	39
3.3.3 Iterative Similarity Registration	40
3.3.4 Iterative Bézier Registration	41
3.3.5 Evaluation	41

3.4	Results	42
3.5	Conclusions	43
4	Identification of Left Atrial Fibrosis Using Image Data	45
4.1	Introduction	47
4.1.1	Electroanatomic Mapping for Atrial Fibrosis	47
4.1.2	Correlating Image Data with Atrial Fibrosis	48
4.1.3	Random Forests	52
4.1.4	Receiver Operating Characteristic Curves	54
4.1.5	Principal Component Analysis	55
4.2	Experimental Data	56
4.2.1	Patients	56
4.2.2	Magnetic Resonance Imaging	56
4.2.3	Computed Tomography	57
4.2.4	Manual Segmentation	57
4.2.5	Electroanatomic Mapping	59
4.3	Methods	59
4.3.1	Calculating Image Features	59
4.3.2	Framework for Classifier Evaluation	61
4.3.3	Single-Parameter Regression	63
4.3.4	Random Forests	65
4.3.5	Principal Component Analysis	66
4.4	Results – MR Cohort	67
4.4.1	Electroanatomic Map and Image Parameters	67
4.4.2	Single-Parameter Regression	68
4.4.3	Random Forest	71
4.4.4	Principal Component Analysis	73
4.5	Results – CT Cohort	78
4.5.1	Electroanatomic Map and Image Parameters	78
4.5.2	Single-Parameter Regression	79
4.5.3	Random Forest	81
4.5.4	Principal Component Analysis	83
4.6	Conclusions	88
4.6.1	Single-Parameter Models	88
4.6.2	Random Forest	88
4.6.3	Principal Components Analysis	89
4.6.4	Future Directions	90

5	Comparing Post-ablation Lesions to Pre-ablation Fibrosis	92
5.1	Introduction	93
5.1.1	Pathological Fibrosis and AF Severity	94
5.1.2	Therapeutic Lesion and Treatment Efficacy	95
5.2	Experimental Data	98
5.2.1	Patients	98
5.2.2	Magnetic Resonance Imaging	98
5.2.3	Electroanatomic Mapping	99
5.2.4	Pulmonary Vein Isolation	99
5.3	Methods	99
5.3.1	Registration and Classification of EAM Points	99
5.3.2	Comparing Image Features	100
5.4	Results	101
5.4.1	Pre-/Post-Ablation Differences	101
5.4.2	Relationship of Fibrosis to Ablation	103
5.4.3	Image-Predicted Bipolar Voltage	106
5.5	Conclusions	107
5.5.1	Future Directions	108
6	Annular Hough Transform for LV in 3-D Echocardiogram	110
6.1	Introduction	111
6.1.1	Survey of Segmentation Approaches	112
6.1.2	The Hough Transform	117
6.1.3	The Hough Transform in Cardiac Segmentation	121
6.2	Experimental Data	122
6.3	Methods	123
6.3.1	Gradient-Based Circular Hough Transform	123
6.3.2	Feature Asymmetry Edge Detector	124
6.3.3	Annular Accumulator	126
6.3.4	Radial Constraint	127
6.3.5	Comparison Criteria	128
6.4	Results	129
6.5	Conclusions	133
6.5.1	Future Directions	134

7	Future Work	136
7.1	Introduction	136
7.2	MR-Compatible EP Lab Paradigm	137
7.2.1	Calibration	139
7.2.2	Projection	140
7.2.3	Conclusion	141
7.3	Pre-/Intra-Op Image Registration	141
7.3.1	3-D Segmentation-Based Registration	143
7.3.2	Temporal Alignment for 3D+t Registration	145
7.3.3	Conclusion	146
7.4	Conclusion	146
8	Conclusion	148
Appendices		
A	Review of Cardiac Physiology and Electrophysiology	153
A.1	Anatomy	153
A.2	Mechanical Cycle	160
A.3	Electrical Cycle	161
A.4	Cellular Electromechanical Coupling	164
B	Heart Disease and Percutaneous Treatment	168
B.1	Coronary Heart Disease	169
B.2	Valvular Heart Disease	169
B.3	Cardiac Arrhythmia	170
C	Machine Learning and Signal Processing by Example	175
C.1	Random Forests	175
C.1.1	Data	176
C.1.2	Training	177
C.1.3	Testing	179
C.1.4	Conclusion	180
C.2	Principal Component Analysis	181
C.2.1	Data	181
C.2.2	Transformation	184
C.2.3	Results and Insight	185
C.2.4	Conclusion	188
D	Supplemental Data for Chapter 4	189
	References	195

List of Figures

2.1	Schematic of echocardiography modes	11
2.2	Cardiac magnetic resonance image	13
2.3	Cardiac X-ray computed tomography	16
2.4	World War I interventional fluoroscopy apparatus	19
2.5	Electroanatomic map of the left atrium	21
3.1	Illustration of local minimum trap in iterative segmentation	36
3.2	Schematic of three-step registration algorithm	37
3.3	Registration accuracy in MR cohort	42
3.4	Registration accuracy in CT cohort	43
4.1	Schematic of the random forest method	53
4.2	Example receiver operating characteristic curve.	54
4.3	Schematic of manual left atrial segmentation technique	58
4.4	Schematic view of image feature region	60
4.5	Confidence scheme for ROC curve	64
4.6	Representative segmented MR images of left atrium	68
4.7	Bipolar voltage distribution from MR cohort	69
4.8	Distribution of image parameters for MR cohort	70
4.9	Median ROC curve and relative importance measures for random forest classifier of MR features	72
4.10	Median ROC curve and relative importance measures for random forest regressor of MR features	73
4.11	Principal component feature definitions for MR cohort	74
4.12	Cumulative variance explained by MR feature principal components.	75
4.13	ROC curve and importance metrics for random forest classifier trained on all MR PCA components	76
4.14	Area under ROC curve for random forest classifier vs number of MR PCA components in feature vector	77
4.15	ROC curve and importance metrics for random forest classifier trained on seventeen MR PCA components	77
4.16	Area under ROC curve and R^2 for random forest regressor vs number of MR PCA components in feature vector	78

4.17	Representative segmented CT images of left atrium	79
4.18	Bipolar voltage distribution from CT cohort	80
4.19	Distribution of image parameters for CT cohort	81
4.20	Median ROC curve and relative importance measures for random forest classifier of CT features	82
4.21	Median ROC curve and relative importance measures for random forest regressor of CT features	83
4.22	Principal component feature definitions for CT cohort	84
4.23	Cumulative variance explained by CT feature principal components.	85
4.24	ROC curve and importance metrics for random forest regressor trained on all CT PCA components	86
4.25	Area under ROC curve for random forest classifier vs number of CT PCA components in feature vector	86
4.26	ROC curve and importance metrics for random forest classifier trained on four CT PCA components	87
4.27	Area under ROC curve and R^2 for random forest regressor vs number of CT PCA components in feature vector	87
4.28	Example slices from LGE-MRI and CE-MDCT images.	90
5.1	Study design comparing pre- and post-ablation image features	100
5.2	Histogram of selected image feature values recorded at ablation points, pre- and post-procedure.	102
5.3	Histogram of selected image feature values recorded at non-ablation map points, pre- and post-procedure.	104
5.4	Box plot of matched point image feature changes (z-score) between pre- and post-ablation images.	105
5.5	Comparison of pre-ablation normal tissue, pre-ablation scar, and ablated tissue by image feature.	106
5.6	Image-predicted bipolar voltage of pre-ablation normal tissue, pre-ablation scar, and ablated tissue.	107
6.1	Illustration of the level set construction	114
6.2	Electronic plane transform circuits in Hough transform patent	117
6.3	Schematic of the Hough transform for 2-D line detection	119
6.4	Comparison of accumulator votes for edge vs gradient data	120
6.5	Schematic of conical vote for directional circle Hough transform	124
6.6	Comparison of gradient vs feature asymmetry edge detectors	126
6.7	Plot of centre likelihood scores $s(c_x, c_y)$	127
6.8	Manual and Hough transform segmentations and evaluation criteria	129
6.9	Performance of gradient endocardial Hough transformation	131

6.10	Performance of feature asymmetry endocardial Hough transform . . .	131
6.11	Performance of annular, feature asymmetry Hough transform	132
6.12	Performance of annular, feature asymmetry Hough transform with radial limits	132
6.13	Feature asymmetry edge map, filtered for endocardial and epicardial edge information	133
6.14	Preliminary alignment of 3-D echocardiogram and MR image based on left ventricular segmentation	134
7.1	Pipeline approach to intraoperative guidance using preoperative advanced imaging.	142
A.1	Four-chamber illustration of the human heart.	154
A.2	Left atrium in posterior cardiac illustrations	155
A.3	Left atrium in anatomic context	157
A.4	Atrioventricular valves, chordae tendonae, and papillary muscles . .	158
A.5	Coronary circulation	159
A.6	Wiggers diagram of the cardiac cycle	160
A.7	Electrical conduction system of the heart	163
A.8	Cellular structure of the heart with intercalated discs	165
A.9	Membrane potential of the myocardiocyte	166
B.1	Schematic overview of classic arrhythmias	172
C.1	Detailed schematic of training a random forest tree	178
C.2	Three-tree random forest model for example data	180
C.3	Women's 400 metre dash competitor height, weight, and speed . . .	183
C.4	PCA component definitions for runner data	186
D.1	Area under random forest MR feature classifier ROC curves for fibrosis voltage thresholds	189
D.2	Accuracy of random forest MR feature classifiers vs number of trees for normalisation strategies and fibrosis thresholds	194
D.3	Accuracy of random forest MR feature regressor vs number of trees and fibrosis voltage threshold	194

List of Abbreviations

1-D, 2-D, 3-D	One-, two-, or three-dimensional, referring in this thesis to spatial dimensions in an image.
2D+t, 3D+t	Two or three spatial dimensions plus a time dimension, i.e. a temporal sequence of 2-D or 3-D images.
AF	Atrial fibrillation, a common arrhythmia (Section 2.4).
AUC	Area under the curve, used in this thesis to indicate the area under the receiver operating characteristic curve, a measure of classifier accuracy (Section 4.1.4).
AV	Atrioventricular, i.e. between the atria and ventricles, e.g. AV valves (Section A.1), or block (an arrhythmia, Section B.3).
AV node	Atrioventricular node, part of the cardiac electrical system responsible for providing a conduction delay (Section A.3).
AVNRT	Atrioventricular nodal reentrant tachycardia, a cardiac arrhythmia (Section B.3).
AVRT	Atrioventricular reentrant tachycardia, a cardiac arrhythmia (Section B.3).
bpm	Beats per minute, unit of measurement for heart rate.
CABG	Coronary artery bypass graft, commonly ‘bypass surgery’, an open surgical treatment for coronary artery disease.
CCB	Calcium channel blocker, a class of medication used to lower blood pressure, decrease heart rate, and/or inhibit vasospasm depending on the particular medication within the class.
CE-MDCT	Contrast-enhanced multidetector computed tomography, a specific CT protocol that involves acquiring an image several minutes after contrast injection (Section 4.1.2). Also DE-MSCT for ‘delayed enhancement multislice computed tomography’.
CHD	Coronary heart disease, a disease of the coronary arteries that supply blood to the myocardium (Section B.1).
CNR	Contrast-to-noise ratio, a measure of image quality.

CT	Computed tomography, an imaging technique whereby a 3-D volume is constructed from multiple 2-D X-ray images (Section 2.2.1). Also CAT for ‘computed axial tomography’.
EAM	Electroanatomic mapping, a computerised system for tracking and guiding percutaneous procedures (Section 2.2.2).
ECG	Electrocardiogram, a recording of cardiac electrical activity using surface (typical) or intracardiac (‘local ECG’) electrodes. Occasionally EKG from the German ‘Elektrokardiogramm’.
ECGI	Electrocardiographic imaging, a method of deducing myocardial electrical potentials using surface electrogram data and advanced computer models (Section 2.2.1).
EP	Electrophysiology, a branch of cardiology concerned with disorders of the cardiac conduction system.
ICD	Implantable cardioverter defibrillator, implanted in patients at risk of sudden cardiac death due to arrhythmia (Section B.3).
ICE	Intracardiac echocardiography, use of an ultrasound transducer mounted on a catheter within the heart (Section 2.2.2).
LA	Left atrium, one of the four cardiac chambers (Section A.1).
LGE-MRI . .	Late gadolinium enhancement magnetic resonance imaging, an MRI protocol that involves acquiring an image minutes after contrast injection (Section 4.1.2).
LV	Left ventricle, one of the four cardiac chambers (Section A.1).
LVH	Left ventricular hypertrophy, the pathological thickening of the muscular wall of the left ventricle.
MI	Myocardial infarction, the dreaded sequela of coronary artery disease, colloquially known as a ‘heart attack’ (Section B.1).
MRI	Magnetic resonance imaging, an imaging modality using magnetic fields and radio waves to generate images (Section 2.2.1).
NMR	Nuclear magnetic resonance, the phenomenon by which nuclei in a magnetic field absorb and emit radiofrequency energy. MRI (above) relies on this effect (Section 2.2.1).
PCI	Percutaneous coronary intervention, a variety of percutaneous treatments for coronary artery disease, including balloon angioplasty and stenting (Sections 2.3 and B.1).

PET	Positron emission tomography, a nuclear medicine imaging modality that relies on detection of gamma ray pairs emitted upon the decay of injected radiotracers.
PV	Pulmonary vein, any of the vessels that carry oxygenated blood from the lungs to the cardiac left atrium (Section A.1).
PVI	Pulmonary vein isolation, a common percutaneous treatment for atrial fibrillation (Section 2.4.3).
RA	Right atrium, one of the four cardiac chambers (Section A.1).
RF	Random forests, a machine learning technique (Sections 4.1.3 and C.1). Also used in the electrophysiology and imaging literature for ‘radiofrequency’, a use avoided in this thesis.
ROC	Receiver operating characteristic curve, a representation of classifier performance over a sensitivity range (Section 4.1.4).
RV	Right ventricle, one of the four cardiac chambers (Section A.1).
SA node	Sinoatrial node, the natural pacemaker of the heart located in the right atrium (Section A.3).
SPECT	Single-photon emission computed tomography, an imaging modality in nuclear medicine that relies on detection of gamma rays emitted upon the decay of injected radiotracers.
SVT	Supraventricular tachycardia, a class of cardiac arrhythmia originating above the bundle of His (Section B.3).
TOE	Transoesophageal echocardiography, the use of an ultrasound transducer placed in the oesophagus (Section 2.2.1). Also ‘TEE’ from the American spelling, ‘transesophageal’.
TTE	Transthoracic echocardiography, the use of an ultrasound transducer placed on the external chest wall (Section 2.2.1).
VF	Ventricular fibrillation, a cardiac arrhythmia (Section B.3).
VT	Ventricular tachycardia, a cardiac arrhythmia (Section B.3).

*Οἱ ἐκλυόμενοι πολλάκις καὶ ἰσχυρῶς, ἄνευ φανερῆς
προφάσιος, ἐξαπίνης τελευτῶσιν.*

*Those who suffer from frequent and extreme prostra-
tion without manifest cause die suddenly.*

— Hippocratic Aphorism^[1]

1

Introduction

Contents

1.1	Motivation	1
1.2	Contribution	3
1.3	Thesis Structure	4
1.4	Originality and Individual Role	6
1.5	Prior Publication	8

1.1 Motivation

The rapid advance of minimally-invasive cardiac procedures promises improvements in patient safety, procedure efficacy, and access to treatment. While percutaneous coronary intervention (PCI) has become routine and highly effective^[2], catheter procedures in areas such as electrophysiology (EP) and valve replacement are still coming of age. This progress is driven by demographics and the improvement in general cardiac care, as patients surviving initial cardiac events go on to require treatment for sequelae^[3]. The growing need for advanced treatment is being answered by developments in catheter technology and procedures. These tools are continually advancing to access and manipulate an ever-broader range of anatomy^[4].

The most obvious limitation of percutaneous procedures compared to an open

surgical approach is the difficulty of visualisation and guidance. This limitation has governed which percutaneous procedures have become standard of care. Contributing to the efficacy and ubiquity of PCI, for example, is that upon accessing a coronary artery, diagnostic and interventional manoeuvres take place in a one-dimensional branching environment under fluoroscopy, a century-old visualisation tool.

Electrophysiology procedures are not so straightforward. Even the simplest involve three-dimensional catheter manipulation in expansive cardiac chambers, requiring more complex guidance techniques. This need has been met by electroanatomic mapping (EAM) systems, computer systems which use probes on the catheter to determine its position in three-dimensional space^[5] (Section 2.2.2).

In addition to displaying catheter position to the cardiologist for procedural guidance, the EAM system allows systematic electrophysiological evaluation of the myocardial substrate. As the cardiologist manipulates a catheter electrode into position to record local electrograms, the EAM system builds a map of local electrograms and their locations^[5,6]. The resulting electroanatomic map can be evaluated on the basis of voltage level, activation timing, and wave morphology.

The wealth of information acquired in advanced cardiac imaging protocols raises the possibility that diagnostic imaging could supplant EAM for some purposes. Combined with the introduction of electrocardiographic imaging^[7,8] and advanced electrophysiologic modelling^[9,10], this would enable comprehensive preoperative procedure planning, even for complex arrhythmia ablations, before the patient is prepped.

In addition to simplifying substrate mapping, advanced imaging algorithms are a gateway to real-time image guidance and intra-operative treatment evaluation. Currently, EAM position data is used to guide and record mapping and ablation procedures. While commercial systems now allow manual pre-operative fusion with clinical images^[11], there is currently no provision for either automated post-hoc registration – useful for determining the limits of image-based scar and fibrosis characterisation – or live image registration – needed to monitor the effect of treatments in near-real-time.

Considerable ongoing research in electrophysiology aims to enable and enhance image-based myocardial substrate characterisation^[12,13], live procedure guidance^[14], and real-time treatment evaluation. This work applies and develops the tools of image analysis to address several of these aims.

1.2 Contribution

The contribution of this thesis is a series of algorithms to advance the utility of sophisticated imaging techniques for interventional electrophysiology. First, and as described in Chapter 3, a registration algorithm was developed to align 3-D clinical images, including magnetic resonance (MR) and X-ray computed tomography (CT) images with electroanatomic maps. This approach was validated in human clinical data in order to establish its utility as a post-hoc analysis tool.

This registration tool was used to confront the problem of prediction of bipolar voltage value from MR and CT images, as described in Chapter 4. While there is a body of work associating image intensity with local electrograms and histology data, this thesis is the first attempt to generate and evaluate a predictive model of bipolar voltage from image features alone. Successful models were found and analysed.

With an understanding of the relationship between fibrosis and image features, we compared the appearance natural (pathological) fibrosis and post-ablation scar in MR images in Chapter 5. We found that while natural scar has higher image enhancement and local variation than normal tissue, these features in ablation scar are still higher. Though these distinctions between native scar and ablation scar are significant, they are small in absolute terms. We hypothesize the difference in physiological implications of natural and ablation scar is due as much to lesion location as character.

Finally, in Chapter 6, we approached a related technological question. Rapid segmentation techniques have a number of uses in interventional procedures. We addressed segmentation by developing a modified Hough transform for the delineation of the left ventricular (LV) myocardium in 3-D echocardiograms. The algorithm improvements are evaluated sequentially and found to greatly improve accuracy.

1.3 Thesis Structure

The work described in this thesis is a cross-disciplinary effort to investigate and advance the role of computerised image analysis in clinical electrophysiology. As such, the thesis document is written to be read and understood not only by those with a background in bioengineering and image analysis, but also by those with a background in clinical cardiology. Where possible, language and explanation throughout the thesis is for both audiences. Further explanation and review of topics that may be unfamiliar to engineers or physicians are provided in the appendices.

The first chapter motivates and contextualises the work to follow, and outlines the structure of the remainder of the thesis. In Chapter 2, we review the historical development and current practice in cardiac imaging and interventional cardiology, with a focus on electrophysiology. We also summarise the natural history and standard treatment options for atrial fibrillation, a disease which drives much of the following work. Review of specific image analysis and machine learning methods is reserved to the chapters in which they are employed.

Chapters 3, 4, 5, and 6 contain original results, and share a common structure. Each chapter begins with an **Introduction**, which includes not only a clinical framing of the specific problems addressed by that chapter’s work, but also a full description – including review of the pertinent literature and methodological summary – of established image analysis methods adopted or adapted therein. This is followed by the **Experimental Data** section, detailing the acquisition of clinical data used for the analysis. In all cases, our image analysis work was performed in concert with established clinical studies, and anonymised data from those studies were used.

The **Methods** section of each chapter details original contributions to image analysis algorithms at a level sufficient to reproduce the work. Where novel techniques or modifications were employed, a full description is given. Where established methods were used, the parameters needed to achieve these results using publicly available tools are supplied. Finally, the **Results** and **Conclusions**

sections describe the experimental outcomes and the implications and future directions suggested by the work.

Chapter 3 lays out the geometric, semi-automated registration algorithm used to correlate electroanatomic map data with image data in subsequent chapters. The algorithm consists of applying a rigid landmark transform followed by iterative determination of affine and non-rigid transformation parameters that minimise mean square point-to-surface registration distance. Although it is a pipeline formed from established methods rather than a novel technique, the registration algorithm is described and validated because of its central role in the methods of Chapters 4 and 5. Validation was achieved in clinical left atrium data.

Chapter 4 details a thorough investigation into the use of clinical imaging to predict bipolar voltage values measured using electroanatomic mapping of the left atrium. While previous studies have established correlation between late gadolinium enhancement MR image (LGE-MRI) data and fibrosis, this work evaluates a full complement of image features on a point-wise basis. Furthermore, this is the first comprehensive investigation of contrast-enhanced multidetector computed tomography (CE-MDCT) images in this context. Features were evaluated in isolation using least-squares regression models, and in concert using the random forest technique. The relationship between image features was investigated with principal component analysis, and predictive accuracy reported for the first time.

Chapter 5 addresses the paradoxical observation that while pathological fibrosis is correlated with severity and persistence of atrial fibrillation, the extent of iatrogenic scar after ablation correlates with treatment efficacy. By comparing pre- and post-ablation MR images using the tools developed in Chapters 3 and 4, the work in Chapter 5 examines the relationship between ‘bad’ fibrosis from pathological remodelling of atrial fibrillation and ‘good’ fibrosis from ablation.

In Chapter 6, we put aside the clinical focus on atrial scarring to develop a more general segmentation tool. Accurate segmentation is a critical step in geometric registration, and a goal of near-real-time performance for procedural imaging requires a simple, parallelisable approach. This chapter describes modifications to

the circular Hough transform, including a feature asymmetry edge detector and a two-step, annular voting mechanism, to improve performance in segmentation of the left ventricle from 3-D echocardiograms. These improvements are evaluated in clinical echocardiograms from healthy volunteers.

Chapter 7 further contextualises the aforementioned work by presenting two clinical paradigms for the role of advanced image analysis in interventional electrophysiology. The first paradigm relies on current work in MR-compatible clinical EP laboratories to enable intra-operative MR monitoring of ablation delivery. The second paradigm consists of a combination of pre-operative MR imaging and intra-operative echocardiography, merged using geometric registration, to facilitate near-real-time, beating-heart image guidance of EP procedures. The contributions presented in previous chapters, in addition to proposed future work, is examined in the context of these developing clinical frameworks.

Finally, Chapter 8 consolidates the contribution of the thesis and its role in the future standard of clinical care in electrophysiology.

1.4 Originality and Individual Role

Generally, all image manipulation and registration, feature extraction from images, machine learning analysis, and data manipulation are my individual, original work. Clinical study design and patient recruitment, clinical image acquisition including MR pulse sequence design, and ablation procedures were performed by clinical collaborators. Manual image segmentation was performed with the help of collaborators. Specifically:

Chapter 3 Clinical collaborators acquired MR and CT images as well as electroanatomic maps. With collaborators, I used commercial software to manually segment MR and CT images.

Independently, I designed and wrote applications to extract EAM data from commercial output files; designed, implemented, and tuned the image registration

algorithm, and applied the algorithm to image data; designed the validation strategy; and performed all data analysis.

Chapter 4 Collaborators again acquired clinical images and maps, and assisted with manual segmentation. Techniques independently developed in Chapter 3 were used to register EAM points to images.

Independently, I designed and implemented custom applications to extract global and local image features, normalise those features, and associate them with EAM points; applied regression, machine learning techniques, and principal component analysis to generate predictive models from image data; evaluated those models, including writing an application to consolidate the output of repeated analyses into statistically meaningful results.

Chapter 5 Collaborators designed the clinical study protocol, acquired images and EAM data, performed the ablation procedures, and assisted with manual segmentation.

Independently, I registered EAM to image data; extracted and analysed image features; calculated predicted bipolar voltage values; and generated comparison between features of the tissue classes.

Chapter 6 Collaborators acquired 3-D echocardiograms. With collaborators, I manually segmented the images.

Independently, I created custom software in C++ to execute the Hough transform including all modifications described; developed the validation parameters; and wrote an application to evaluate the segmentation technique.

Authorship I am the sole author of this thesis document. Unless otherwise indicated, I produced all figures. Where figures are adapted from other sources, the provenance is indicated in the caption and detailed in the References. Critique, proofreading, and suggestions were provided by Prof. Alison Noble, Dr. Saman

Nazarian, Dr. Elizabeth Finn, Sabine Parrish, Zoë Goodwin, Maryam Ahmed, Michael Pontikos, and Ana Namburete. All errors and omissions are my own.

1.5 Prior Publication

Much of the work contained in this thesis is currently being prepared for publication. Preliminary presentations and publications have been accepted based on partial analyses of the same clinical data. Specifically:

A subset of the delayed enhancement multi-detector CT image data used in Chapter 4 was presented at the American Heart Association 2013 conference^[15]. The work presented there incorporated only a single-feature regression model. It did not incorporate multiple image features, machine learning algorithms, principal component analysis, or prediction-based performance metrics that comprise the original contribution described in this thesis.

The delayed enhancement MR image data used in Chapter 5 was presented at the 35th Heart Rhythm Society conference in May 2014^[16]. This preliminary presentation considered only a single mean intensity parameter. Unlike the more complete analysis described in this thesis, it did not incorporate multiple image features, principal component analysis, or region-based comparison of bipolar voltage data.

An earlier version of the Hough transform work presented in Chapter 6 was published at the 9th IEEE International Symposium on Biomedical Imaging (ISBI)^[17] in May 2012. That version incorporated the annular modification to the circular Hough transform and included the same clinical echocardiogram data for validation. However, it did not incorporate the feature asymmetry edge detector or radial constraint described in this thesis.

*Alles Gescheite ist schon gedacht worden.
 Man muss nur versuchen, es noch einmal zu denken.
 All intelligent thoughts have already been thought;
 what is necessary is only to try to think them again.*

— Johann Wolfgang von Goethe^[18]

2

Background

Contents

2.1	Introduction	9
2.2	Cardiac Imaging	10
2.2.1	Diagnostic Imaging	10
2.2.2	Interventional Imaging	18
2.3	Interventional Cardiology	22
2.4	Atrial Fibrillation	23
2.4.1	Atrial Remodelling	24
2.4.2	Sequelae	25
2.4.3	Treatment	25

2.1 Introduction

In this chapter, we review the historical development and current toolset of cardiac imaging and interventional cardiology, with a focus on electrophysiology, to establish a context for the original contributions presented in chapters that follow. We also summarise the natural history and treatment options for atrial fibrillation, a disease which motivates the clinical need for much of the work presented in this thesis.

Review of specific image analysis and machine learning methods is reserved to the chapters in which they are employed. An overview of cardiac anatomy and

physiology is available in Appendix A, and a summary of the impact of heart disease treatable by percutaneous techniques is found in Appendix B.

2.2 Cardiac Imaging

Within months of Röntgen’s discovery of the X-ray in 1895^[19], cardiac pathology was being investigated via non-invasive imaging^[20]. Over the intervening years, cardiac imaging modalities and techniques have advanced significantly. Clinically, cardiac imaging is used for two broad purposes: diagnosis of pathophysiology and guidance of interventional procedures. These applications impose different requirements on imaging equipment, image acquisition time, computational complexity, spatial and temporal resolution, and tissue discrimination. The common diagnostic and interventional cardiac imaging techniques in current clinical practice are reviewed below. An accessible introduction to the physics of medical imaging can be found in Webb’s *Introduction to Biomedical Imaging*^[21]. A comprehensive overview of the use of imaging in clinical cardiology is presented in Leeson’s *Cardiovascular Imaging*^[22].

2.2.1 Diagnostic Imaging

Beyond the chest X-ray (‘plain film’), the key non-invasive imaging modalities in diagnostic cardiology are echocardiography, magnetic resonance imaging, and X-ray computed tomography, which are reviewed below. Nuclear medicine, including positron emission tomography (PET) and single-photon emission computed tomography (SPECT), are not discussed here, as they do not play a role in the chapters to follow.

Echocardiography

The use of acoustic waves for medical diagnosis, inspired by naval sonar, was initially developed in the 1940s^[27]. By 1954, the first clinically useful cardiac ultrasound – examining motion of the mitral valve in stenosis – was reported^[28]. These early scans were one-dimensional images (‘A-mode’), sometimes repeated to generate a time axis (‘M-mode’; Figure 2.1, left). The sector-scanning probe was developed




Figure 2.1: Schematic of echocardiography modes. Left: 1-D M-mode. Centre: 2-D B-mode. Right: 3-D ultrasound. Figure includes images from *Wikimedia Commons* under permissive licenses^[23–26].

in the 1970s^[29,30], leading to the 2-D ‘B-mode’ that a modern cardiologist would recognise as an echocardiogram (Figure 2.1, centre).

Two-dimensional transthoracic echocardiography (TTE) is a well-validated technique for diagnosis of valvular heart disease^[31], structural disease including left ventricular hypertrophy^[32,33], and functional heart diseases such as dyssynchrony^[34] and heart failure^[35]. Stress echocardiography – with cardiac stress induced by exercise or pharmaceutical agents – is validated for the diagnosis of coronary heart disease and cardiac ischaemia and provides prognostic benefit beyond the more typical stress electrocardiography (ECG) exam^[36–38]. TTE is quick, non-invasive, inexpensive, and permits immediate review by the cardiologist and frequent follow-up imaging. Echocardiogram image quality is dependent on body habitus, requiring significant operator skill compared to CT and MR imaging in challenging patients.

Due to the shape of the ultrasound beam, the spatial resolution of images is anisotropic. Resolution is best in the axial direction – the direction of the beam – and depends on the ultrasound pulse size, which is typically between two and three times the wavelength. Higher-frequency ultrasound therefore attains better resolution, but at the expense of decreased tissue penetration. Typical TTE probes

emit 3 to 5 MHz ultrasound^[39], achieving axial resolution of 0.3 to 0.5 mm^[40]. Lateral resolution at the focal point of the ultrasound beam is similar to axial resolution, but drops off beyond the focal point.

Three-dimensional ultrasound probes feature two-dimensional raster scanning of the ultrasound beam to form a wedge-shaped 3-D image (Figure 2.1, right). Although this technology is already common in obstetrics, the transition from research to clinical use in cardiology is ongoing. Particular challenges for 3D-TTE include ECG gating, the necessity of breath-holding, and the artefacts and loss of temporal resolution introduced when multiple heart beats are needed to acquire one image^[41]. Nonetheless, 3D-TTE has been validated against conventional 2-D echocardiography in assessment of coronary artery disease in stress studies^[42] and compares favourably to SPECT in evaluation of myocardial ischaemia^[43].

Transoesophageal echocardiography (TOE) is performed using a probe positioned in the patient's oesophagus. This provides a more clear view of the atria, valves, and major vessels as compared to the transthoracic approach. TOE is well validated in evaluating left atrial thrombus in atrial fibrillation patients^[44]. TOE probes typically emit 7 MHz ultrasound, achieving axial resolution near 0.2 mm^[40].

Doppler mode scans rely on the Doppler effect – the frequency shift of waves reflected from a moving interface – to determine the direction and velocity of blood flow relative to the probe. This is invaluable in diagnosing valve pathology^[45,46], arterial stenosis^[47], and structural defects with flow^[48].

Magnetic Resonance Imaging

Nuclear magnetic resonance (NMR) is a phenomenon whereby atomic nuclei with an odd total number of nucleons (protons and neutrons) absorb and re-emit quantised radiation in a magnetic field. The radiation frequency is dependent on the local magnetic field strength. This observation was seized upon by chemists and physicists in the late 1940s to develop NMR spectroscopy, a tool to tease out molecular structure using atoms' local electromagnetic environments^[49].




Figure 2.2: Slices from a delayed enhancement cardiac magnetic resonance image (bottom). Anatomic slice positions in left lateral view (top) are approximate. Heart illustration from Gray^[50].

Researchers soon realised that just as the magnetic environment within a molecule alters NMR peaks to reveal atomic structure, an externally applied magnetic gradient could couple local magnetic field strength with position, allowing images to be generated. In the 1980s, this technique was refined to visualise the hydrogen nuclei in tissue water molecules, resulting in a magnetic resonance image (MRI)^[49,51]. MRI is extremely versatile: modern machines offer almost limitless customisation of pulse sequence, field of view, ECG and respiratory gating, and countless other options. A fuller treatment of MR physics is beyond the scope of this chapter, but an accessible introduction is found in Webb^[21].

MRI can deliver high resolution and high CNR technique for measuring anatomical and functional cardiac parameters^[52] (Figure 2.2). The variety of pulse sequences and gating technology available allows for a range of applications, including discrimination between viable and non-viable tissue after myocardial infarction (MI)^[53,54], measurement of coronary artery bloodflow^[55,56], and tracking myocardial tissue and stress-strain relationships^[57].

Exogenous contrast material was developed for MRI and first used in MR angiography, blood-brain barrier imaging, and bowel studies^[49,58,59]. The most

common agents are chelates of gadolinium(III), which lowers the T_1 relaxation time of nearby protons. Angiography using early contrast agents required first-pass imaging to acquire images before extravasation of contrast, while newer chelates have a higher intravascular lifetime^[60]. Late gadolinium enhancement MRI (LGE-MRI, also DE-MRI for ‘delayed enhancement’) was developed to identify latent cardiac injury after MI. LGE images are acquired several minutes after the injection of a gadolinium-based contrast agent^[61–63]. Over a time course of minutes, the contrast agent has a higher rate of extravasation and lower rate of clearance in fibrotic tissue^[64], an observation which allows lesion discrimination in a wide variety of clinical settings^[13,65,66].

Cardiac gating is the method used to mitigate the interference of cardiac motion during image acquisition. Gating can be classified broadly as prospective or retrospective. In prospective gating, image data is only collected during a desirable time period – usually a defined number of milliseconds after the electrocardiogram T wave. If the imaging system detects the subsequent QRS complex while image data is still being acquired, that set of data is discarded. By contrast, retrospective gating is a technique in which image data is continuously acquired and ‘binned’ based on its point in the cardiac cycle as determined by ECG. After acquisition, the data is rearranged to account for its temporal position. This method is commonly used to generate cine images, but is also helpful for still images of patients with arrhythmia.

Respiratory gating is the analogous approach to compensate for respiratory motion. While 2-D images may be acquired during a single breath-hold, this is impractical for 3-D stacks and for patients with difficulty complying with breath-holding. Instead, a respiratory navigator – a continuously acquired, 1-D MR image of the right hemi-diaphragm – is used to track respiratory phase. Like cardiac gating, respiratory gating may be performed prospectively or retrospectively^[67,68].

The number of choices possible during MR pulse sequence design, acquisition gating, and signal processing allow a prodigious range of trade-offs between spatial resolution, temporal resolution, and signal quality in MR images. Typically, clinical cardiac imaging achieves in-plane spatial resolution between 0.8 and 1.5 mm. 3-D

images are reconstructed from stacks of multiple 2-D planes, so in-plane and between-plane resolutions are often different. For cine images, temporal resolution is typically between 35 and 120 ms^[69–71].

While the versatility of MRI makes it ideal to establish diagnoses and generate management decisions, the prohibitive time and expense of the scan hinder its use in routine follow-up. Additionally, certain patients, including those with ferromagnetic implants or who are clinically unstable, are poor candidates for MR imaging. Gadolinium contrast agents are contraindicated in patients with severe kidney disease^[72].

Computed Tomography

Shortly after the discovery of the X-ray, work began to use the rays to produce 3-D images. Early efforts to extract a single cross-sectional plane resulted in ‘conventional tomography’, in which the X-ray source and film were moved in concert during imaging, resulting in the blurring of shadows outside the focal plane^[73]. The ambition of researchers quickly turned to generation of a full 3-D representation. In 1917, Johann Radon published his eponymous transform to mathematically reconstruct a 3-D volume from multiple 2-D projections^[74,75]. Radon’s foundations bore fruit in the early 1970s when Hounsfield constructed a computerised device to reconstruct a phantom image from multiple X-ray images. The original device required nine days of scanning time and 2.5 hours of computation time to process a single image^[76].

This technique was alternatively known as computed tomography (CT) and computed axial tomography (CAT). Initially its application was limited to neuroradiology due to the lengthy scan times, difficulties computing larger volumes, and the lack of other neuroimaging technology available at the time. However, by the late 1970s, full-body scanners were developed, and by the 1980s clinical use was becoming commonplace^[76].

CT provides high spatial resolution and CNR (Figure 2.3). Modern multi-detector CT (MDCT) devices, in which an X-ray tube is mounted opposite a fixed




Figure 2.3: Slices from an X-ray computed tomography image (bottom). Anatomic slice positions in left lateral view (top) are approximate. Heart illustration from Gray^[50].

bank of 64 to 320 detectors, can achieve spatial resolution of 0.3 to 0.6 mm. The temporal resolution is half the gantry rotation time, typically 167 to 250 milliseconds (ms). Dual source CT (DSCT) devices utilise two X-ray tubes and detector banks mounted at 90°, doubling temporal resolution (75 to 83 ms) without sacrificing spatial resolution^[22].

Like MRI, most CT imaging techniques make use of the surface electrocardiogram for cardiac gating. However, CT images can generally be acquired during a single breath-hold, obviating the need for respiratory gating.

Due to its resolution and the absorption of X-rays by heavy atoms such as calcium, CT is especially well-suited for study of the coronary arteries. Modern scanners have excellent sensitivity and specificity in detecting coronary artery stenosis^[77] and calcification^[78,79]. Iodinated contrast agents – designed around a tri-iodinated benzene ring – are used to enhance the tissue discrimination of CT images^[80]. In addition to specific use in CT angiography, contrast agents are used in most CT scans to provide parenchymal tissue enhancement. Iodinated agents carry a risk of nephrotoxicity, especially in patients with pre-existing kidney injury. While this risk can be mitigated by pre-hydration and selection of non-ionic and iso-

osmolar agents, use in patients with severe kidney disease remains controversial^[80,81]. The other major risk of CT is the burden of ionising radiation. A simple chest CT delivers approximately 70 times the radiation burden as a chest X-ray, and a coronary CT angiogram may deliver up to 160 times the dose^[82].

Unlike MRI, CT is compatible with metal devices including implanted pacemakers and stents, and is often used in place of MRI for these patients. Although CT is generally less well suited for soft-tissue discrimination, as permeability to X-rays is similar among soft tissue, contrast-enhanced multidetector computed tomography (CE-MDCT) has been developed as an alternative to LGE-MRI^[12,83,84] and has the same pharmacokinetic rationale, allowing detection of flow defects and late enhancement^[84].

Electrocardiographic Imaging and Modelling

In the last decade, advances in computer modelling of electrical propagation have opened the door to a wide range of advanced tools that stretch the boundaries of traditional definitions of ‘imaging’. Electrocardiographic imaging (ECGI) is an iconic example. Using CT images to generate a 3-D map of the position of all tissues between the myocardium and the skin, ECGI back-projects from a dense array of surface electrograms to calculate the course of electrical activity on the myocardium^[7,8]. This gives cardiologists access to complex data such as activation timing and specific areas of conduction block without invasive procedures.

In addition to using these electrical models to build a picture of a patient’s heart as it is, today’s advanced modelling techniques allow *in silico* experiments that consider an individual’s cardiac substrate as deduced from ECG and advanced imaging data. Recent studies have been able to model endocardial as well as epicardial activity^[85]. Though this is a nascent technology, it holds enormous promise in interventional electrophysiology as a way to evaluate the outcome of ablation patterns or even medication in an individualised manner before subjecting the patient to the risks of intervention^[9,10]. Personalised myocardial models have been validated in predicting outcomes of cardiac resynchronisation therapy^[86].

2.2.2 Interventional Imaging

While cardiac procedures conducted via sternotomy or thoracotomy are guided by direct visualisation of the structures involved, percutaneous procedures are without this luxury. Basic requirements for procedural guidance imaging include near-real-time data acquisition, rapid image presentation in a clinically useful format, physical compatibility with the operative procedure, and minimal risk to the patient. Currently, fluoroscopy and echocardiography are the most common imaging modalities guiding electrophysiology procedures. Electroanatomic mapping is another operative guidance tool, often integrated with imaging systems, which is also outlined below.

While investigations are being conducted into MRI and CT for live guidance or ‘online review’ of percutaneous cardiac procedures^[57,87–91], the barriers to adoption are significant and include cost, geometric constraints, incompatibility with surgical instruments, and computation time required to generate an image. These hurdles, the progress being made to overcome them, and the contributions of this thesis and related work in doing so, are described in Chapter 7.

Fluoroscopy

X-ray fluoroscopy was the first technique for intraoperative image guidance and remains the standard method of guidance in interventional cardiology. Röntgen’s initial discovery was that invisible rays emitted from a cathode ray tube could pass through black cardboard and cause fluorescence of a barium platinocyanide screen^[19]. Thus, fluoroscopy – the real-time observation of transmitted X-rays using fluorescent media – preceded X-ray photography, and the two technologies developed together. While George Eastman’s Kodak Company supplied plates and paper to capture X-rays in the late 1890’s, Thomas Edison and his associates experimented with fluorescent chemicals for viewing screens^[93]. They arrived at calcium tungstate, which has a fluorescent intensity six times that of barium platinocyanide. Fluorescent screens were soon used for real-time diagnostics and procedural guidance (Figure 2.4).



Figure 2.4: World War I surgeon uses a fluoroscope to locate bullets. X-rays from the source under the table excite fluorescent salts in the screen worn on the surgeon's face. Credit to Dr. Hoguet, American Ambulance Hospital, Neuilly, France^[92].

While incremental advances in screen design continued into the early decades of the 1900's, and multi-layer screen-film techniques improved static X-ray images, the next leaps in fluoroscopy would occur in the 1940's. The development of photomultiplier tubes, X-ray image intensifiers, and analogue television technology allowed clinicians to use central cone vision, eliminated the need for strict pre-procedure dark adaptation, and allowed them to move away from the X-ray beam^[94]. In the 1970's and 80's, digital fluoroscopy was developed, first using analogue-to-digital converters and then with digital detectors. These advances allowed further reduction in radiation dose and recording of images for instant playback or storage^[94].

Today, interventional cardiologists and radiologists rely on fluoroscopy for essentially all procedures. Electrophysiologists find the method particularly critical

in guiding catheters to the correct cardiac chamber, including by transseptal puncture^[95]. Fluoroscopy provides excellent instantaneous 2-D visualisation of radiopaque surgical instruments but is less useful in identifying soft tissue structures or characterising three-dimensional motion. The ionising radiation dose associated with fluoroscopy, although significantly reduced from its origins, is still a concern, and carries short- and long-term risks dependent on the length of the procedure and the skill of the operator^[96,97].

Echocardiography

Echocardiography methods used in procedural guidance include TOE, intracardiac echo (ICE), and, to a lesser degree, TTE^[98,99]. TTE is widely available and used in a diagnostic capacity due to its ease of acquisition and rapid interpretation. The thoracic view provides particular advantages in procedures across the chest wall, such as myocardial biopsy and pericardiocentesis. However, the left lateral decubitus position that gives the best TTE viewing windows are often incompatible with procedural guidance.

TOE holds the more general advantages of superior visualisation of most cardiac structures and relative impunity to body habitus. Additionally, transoesophageal placement is ideal for intraoperative manipulation by a clinician controlling the airway and does not complicate surgical access to the thorax when that is required. Advances in 3D-TOE continue to improve its utility in both open and percutaneous procedures^[100].

In ICE, the ultrasound transducer is positioned on a catheter within the heart, usually in the right atrium. This allows exceptional views of cardiac structures and instruments. When it is incorporated into atrial fibrillation ablation procedures, ICE is used to guide transseptal puncture, establish lesion position, ensure good contact between the catheter and the myocardium during ablation, and monitor for cavitation and effusion^[101]. Vascular and intracardiac probes generally emit ultrasound at 10 MHz or higher^[22,40], achieving axial resolution better than 0.15 mm.




Figure 2.5: Electroanatomic map of the left atrium. Colours represent bipolar electrogram voltage (red, < 0.1 mV; purple, > 0.5 mV), white dots represent sampled map points, and gray spheres are manually placed landmarks. The 3-D surface is derived from manual registration with a CT image using commercial software.

Electroanatomic Mapping

Electroanatomic mapping is a relatively recent technique for localising interventional catheters in the heart^[6,102]. EAM uses electromagnetic detectors and specially designed catheters, which can be tracked in three-dimensional space. There are two primary commercial EAM systems. The CARTO product was designed by Biosense Webster, Inc. of Diamond Bar, California and uses magnetic sensors embedded in the intracardiac catheters, coupled with external magnets, to localise the catheter^[103]. The NavX system, developed by St. Jude Medical of Saint Paul, Minnesota, uses electrical impedance measurements between the catheter electrode and fixed electrodes on the patient's body to determine catheter position. The two systems are comparable in position accuracy^[104].

Heart chamber volumes are initialised in the EAM computer when the cardiologist sweeps the catheter through the chamber. The EAM computer builds a record of every point reached, and estimates the chamber boundary using those points. Local electrocardiograms are mapped onto this virtual myocardium to identify arrhythmia sites and guide ablation. Throughout a procedure, local electrograms

and ablation points are depicted by a combination of colour coding, spherical markers, and points (Figure 2.5).

Recently, the CARTO and NavX products have incorporated the ability to manually register pre-operative MR or CT images to the EAM system, allowing cardiologists to take advantage of known anatomy and judge enhancement or other features^[5] (Figure 2.5). Incorporating these images has been shown to reduce fluoroscopy time and procedure complexity^[11]. Current state-of-the-art systems also incorporate ICE probes and use catheter position data to align these 2-D scans into the 3-D volume^[105]. However, EAM systems do not currently use intraoperative whole-heart imaging to provide dynamic, beating-heart guidance^[106].

2.3 Interventional Cardiology

The earliest attempts at deep vascular – particularly arterial – access were experiments with arteriography and central delivery of noxious medications^[107]. In 1929, minimally invasive cardiac access was first popularised when Werner Forssmann, a German surgical trainee, passed a catheter through a vein in his arm to his right ventricle. The procedure was guided via a mirror held by a nurse, allowing Forssmann to view the fluoroscopy screen^[108,109]. Attempts at vascular cutdown access and direct puncture of vessels continued for decades, succeeding piecemeal until Seldinger published his method of exchanging needles and catheters over a guide wire in 1953^[110]. By the 1960s, coronary angiography was achieved^[111].

Vascular intervention first occurred serendipitously in 1964 with the accidental recannulisation of an occluded iliac artery during angiography^[107,112]. Shortly thereafter, intravascular stents^[113] and balloon catheters^[114] were introduced, and catheterisation of the coronary arteries became possible^[114,115]. Coronary angioplasty and stenting, collectively known as ‘percutaneous coronary intervention’ (PCI), remain the mainstays of interventional cardiology today. The role of the interventional cardiologist is expanding to include closure of atrial septal defects^[116], mitral valve repair^[117], and obliteration of the left atrial appendage^[118].

The history of interventional electrophysiology parallels that of coronary intervention^[119]. In 1929, the first cardiac electrical stimulation device was reported^[120], and by the 1950s, external pacemakers were used to sustain patients in complete heart block^[121]. By the 1960s, synchronised cardioversion^[122] and demand pacing^[123] were being developed. Intracardiac recording and ablation catheters were investigated in the 1980s^[124,125], paving the way for modern ablation techniques.

Today, electrophysiologists perform diagnostic electrophysiology studies ('EP studies'), using electroanatomic map technology to record local electrograms, challenge the myocardium with aggressive pacing, and characterise the cardiac electromechanical milieu^[126]. They implant subcutaneous devices including defibrillators, pacemakers, and rhythm monitors. And they complete ablation procedures which apply intracardiac radiofrequency energy, laser, or cryoablation to myocardial substrate in order to interrupt accessory pathways and disrupt arrhythmogenic foci. Some of these procedures, such as accessory pathway ablation, are tailored specifically to findings on the EP study^[127]. Others, including pulmonary vein isolation for atrial fibrillation, are relatively standardised (Section 2.4.3).

2.4 Atrial Fibrillation

Atrial fibrillation (AF) is an arrhythmia (Appendix A) characterised by disordered electrical activity in the cardiac atria. Instead of receiving regular stimulation from the sinoatrial node, the atrial myocardium in AF pulses 400 - 600 times per minute due to ectopic triggers and disordered conduction^[128]. Atrial fibrillation is the second most common arrhythmia after simple premature beats. The incidence is highly age-dependent, affecting 0.5% of people in their 50s and 10% of those in their 80s. AF increases the mortality rate by 50% in men and nearly doubles the rate in women^[129].

Atrial fibrillation presents both as a lone entity, usually in patients between 40 and 60 years old^[130], and as a comorbidity in patients with underlying cardiac disease^[131]. Lone AF classically begins as paroxysms before increasing in frequency and developing into a chronic, persistent arrhythmia. AF patients are also at risk of further cardiac morbidities, especially congestive heart failure^[132].

2.4.1 Atrial Remodelling

The progressive natural history of AF is classically described as ‘atrial fibrillation begets atrial fibrillation’^[133]. The pathological basis of this course is a threefold remodelling process^[134]. Animal studies have shown that inducing sustained AF in otherwise healthy hearts causes the atrial refractory period to shorten and rate adaptation to be lost. Over the course of days to weeks, these changes result in higher susceptibility to AF and lower spontaneous resolution of subsequent episodes^[133,135]. While this electrical remodelling is a factor in the repeated induction of paroxysmal AF and the maintenance of persistent AF, it is fully reversible. After sinus rhythm is restored, the refractory period and rate adaptation return to normal over the course of days, even after prolonged periods of sustained fibrillation^[136].

Contractile remodelling develops over a longer time course. Useful atrial contraction is of course absent during fibrillation, but the mechanical defect lingers after sinus rhythm is restored, in a time-dependent fashion. It takes about 24 hours for full mechanical recovery after AF lasting two weeks, while a full month is required after AF lasting six weeks^[137]. This effect was initially thought to be ‘atrial stunning’ secondary to the electrical cardioversion used to restore sinus rhythm in the original studies, but further work showed that stunning is also present after pharmacologic cardioversion or spontaneous reversion^[138].

Finally, structural remodelling is a physical change in the cellular and extracellular myocardial structure after persistent AF. Chronically fibrillating atrial cardiomyocytes exhibit loss of myofibrils, mitochondrial changes, glycogen accumulation, dispersion of chromatin, and fragmentation of the sarcoplasmic reticulum^[139]. These changes are similar to those shown by chronically ischaemic myocardium (so-called ‘hibernating myocardium’) and is not a prelude to apoptosis^[140]. In humans, increased fibrosis is observed in the myocardium of chronic AF patients, both between individual myocytes and between atrial bundles^[141,142]. It is yet unclear whether this fibrosis is a pure effect of AF or part of the interplay between AF and comorbidities^[143] or the ageing process^[144].

Electrical, contractile, and structural remodelling play synergistic roles in the maintenance of atrial fibrillation^[134]. Higher levels of fibrosis correlate with lower bipolar voltage on EAM and enhancement on LGE-MRI, and are associated with higher AF recurrence rates after pulmonary vein isolation (Section 4.1.1).

2.4.2 Sequelae

Stroke is the dreaded complication of atrial fibrillation^[145,146]. Loss of atrial contractility during AF leads to stasis of blood, particularly in the left atrial appendage, which is haemodynamically sheltered from the passive flow of blood from the pulmonary veins into the left ventricle. Stasis, in combination with fibrotic atrial myocardium^[141,142] and increased fibrin turnover^[147], contributes to a prothrombotic state^[146]. Portions of left atrial thrombus may then embolise and enter the arterial, including cranial, circulation.

Several clinical classification strategies have been proposed to stratify stroke risk for the purpose of choosing recipients of anticoagulation therapy, most recently the ‘Birmingham 2009’ point-based scoring system commonly referred to by the acronym CHA₂DS₂-VASc^[148]. These scores recognise increased risk correlates with age, male sex, heart failure, vascular disease, hypertension, and diabetes. Based on stroke risk, AF patients may be treated with anticoagulation medication^[149,150] or considered for surgical closure of the left atrial appendage^[151]. Anticoagulation must be weighed against increased bleeding risk^[152].

Mitral valve changes are another complication of atrial fibrillation. As the atrial myocardium becomes more compliant and fibrotic in AF, atrial size increases^[153]. This leads to a stretching of the mitral annulus, resulting in loss of sphincteric action and possibly mitral regurgitation^[154].

2.4.3 Treatment

The two goals of AF treatment are prevention of embolic events and limitation of the cardiovascular effects of AF.

Therapeutic approach and challenges in anticoagulation are tangential to the focus of this thesis. Briefly, anticoagulation is indicated acutely for any patient undergoing electrical or medical cardioversion – unless specific criteria are met indicating no atrial thrombus – and chronically for any patient who will remain in AF and meets CHA₂DS₂-VASc or similar risk stratification criteria^[148]. While warfarin is classically used, newer options include direct thrombin or factor Xa inhibitors, including dabigatran, rivaroxaban, and apixaban. These agents offer similar or decreased rates of ischaemic stroke and significant bleeding compared to warfarin, but obviate the need for INR checks and carry fewer dietary restrictions^[155].

Direct signs and symptoms of AF are largely caused by rapid ventricular response. Therefore, the two broad treatment strategies for AF are rate control and rhythm control. The rate control strategy aims to pharmacologically limit the ventricular rate while allowing AF to persist in the atria. The rhythm control strategy, by contrast, is to terminate AF and maintain sinus rhythm via medication, surgery, or ablation. The largest clinical study of its type found that rate control plus anticoagulation offers an identical survival advantage to rhythm control with less risk of drug toxicity^[156]. However, a role for rhythm control remains, especially for patients who cannot tolerate antiarrhythmic medications and for younger patients who wish to avoid long-term atrial remodelling^[157].

Rate Control

Limiting the ventricular rate – which can reach 150 beats per minute (bpm) in untreated AF – prevents symptomatic haemodynamic instability and tachycardia-induced cardiomyopathy^[158]. While an optimal rate goal has not been established, a target of <110 bpm has proven as effective as <80 bpm and required fewer clinical visits to achieve^[159].

Rate control is typically achieved using beta blockers or nondihydropyridine calcium channel blockers (CCBs), with digoxin as a second-line choice^[160,161]. Beta blockers and CCBs carry a risk of hypotension, can exacerbate heart failure, and may contribute to atrioventricular node block. Due to digoxin's inotropic effects,

it was once first-line therapy for patients with AF and heart failure. However, its narrow therapeutic window and low efficacy in controlling heart rate have relegated its use to patients who cannot tolerate a beta blocker or CCB^[162].

Rhythm Control

The goal of rhythm control is to interrupt atrial fibrillation and maintain the patient in a sinus rhythm indefinitely thereafter.

Pharmacologic Rhythm Control Antiarrhythmic medical therapy is beyond the scope of this thesis. Briefly, first-line agents include amiodarone, dofetilide, flecainide, propafenone, and sotalol^[160,163]. Choice of medication depends on the presence of structural heart disease, coronary artery disease, and heart failure. Amiodarone has the highest efficacy^[164], but also the worst long-term side effect profile, including central nervous system effects, gastrointestinal and hepatic disruption, hyper- and hypothyroidism, and damage to the lung parenchyma.

Cox Maze Procedure This procedure aims to cut a ‘maze’ of conducting myocardium within the atria, creating obstacles of scar that will prevent an AF wavefront from propagating^[165]. Though the procedure has undergone many revisions since its development in the 1990s, the core strategy of extensive cuts and re-sewing through both left and right atria remains unchanged. Though this is an invasive procedure requiring cardiopulmonary bypass, success rates are 89.3% after 10 years^[166].

Several alternative approaches have been evaluated. The ‘corridor’ operation isolated a path of atrial tissue connecting the SA node to the AV node. While this prevented rapid ventricular response and eliminated the irregular pulse, the bulk of the atrial myocardium remained fibrillating. Thus, anticoagulation is still necessary, and this approach has been clinically abandoned^[167]. Alternatively, taking a cue from catheter ablation procedures (below), surgical en bloc isolation of the pulmonary veins has been evaluated. This approach was found more effective than the maze procedure in a study of 30 patients^[168].

Finally, radiofrequency and cryothermal energy have been used during open chest procedures to duplicate the maze procedure lesions. These methods are quicker and easier than cutting through the myocardium, and do not require cardiopulmonary bypass^[169,170]. The radiofrequency approach has similar outcomes to the maze procedure, exhibiting 91% effectiveness one year post-procedure^[171].

Catheter Ablation Catheter ablation for AF was originally developed as an alternative to the Cox maze procedure^[172], but replicating those lesions with radiofrequency ablation proved difficult and fraught with complications. When it was observed that 94% of AF foci occur in the pulmonary veins (PV)^[173], the goal of catheter ablation quickly shifted to circumferential ablation of the pulmonary vein ostia in order to electrically isolate the arrhythmogenic foci from the atrial myocardium^[174].

Pulmonary vein isolation (PVI) remains the mainstay of AF catheter ablation today. In addition to radiofrequency energy, which carries risks of perforation and stenosis^[175,176], cryoablation has been used in PVI procedures. Cryolesions are transmural and homogeneous but with reduced inflammation compared to radiofrequency^[177–179].

Another substantial challenge to standard ablation is the point-by-point nature of the lesion, where success requires an unbroken circumferential scar. Alternative techniques have been introduced to address this concern using ‘one-shot’, balloon-guided ablation. In the cryoballoon method, a balloon is inflated at the tip of the catheter, obstructing one pulmonary vein at a time. While it is inflated and in circumferential contact with the ostium, pressurised nitrus oxide is delivered to the balloon. The liquid-to-gas transition of the nitrus oxide reduces the temperature of the balloon to -70 °C, creating a circumferential lesion^[180,181]. The laser ablation balloon approach is substantially similar. With the pulmonary vein occluded, an endoscopic aiming system is used to align an intra-balloon laser with the pulmonary vein ostium in contact. The laser is then activated, generating a circumferential lesion^[182]. Finally, in through-the-balloon ultrasound, an high-frequency ultrasound

transducer is used to heat the saline-filled balloon until it forms a lesion where it is in contact with the pulmonary vein ostium^[183]. These balloon methods, in different stages of evaluation, have proven effective^[184].

In addition to pulmonary vein triggers, localised electrical rotors and focal impulses within the atrial myocardium contribute to AF^[185,186]. Some clinical centres incorporate targeted ablation of these regions in addition to pulmonary vein isolation in their ablation strategies. Others continue to employ arbitrary ‘roof lines’ for substrate modification. In current practice, a single ablation procedure achieves sustained success in about 40% of cases. Long-term success considering clinically indicated repeat procedures approaches 80%^[187,188].

Risks of AF catheter ablation include pulmonary vein stenosis (1 - 3%), cardiac perforation with tamponade (1%), clinically evident embolism (0.4 - 2%), phrenic nerve injury (< 1%), and atrioesophageal fistula (0.3 - 0.5%)^[175,176].

Licensing Note

Due to its provenance, Figure 2.1 is hereby licensed under the Creative Commons Attribution-ShareAlike 4.0 International License. This applies to the image only, not its caption or other accompanying text. The license is available at <http://creativecommons.org/licenses/by-sa/4.0/>.

Τῶν δ' ὡς λόγου μόνου συμπεραينوμένων μὴ εἶη
ἐπαύρασθαι, τῶν δὲ ὡς ἔργου ἐνδείξιος· σφαιερὴ γὰρ
καὶ εὐπταιστος ἢ μετ' ἀδολεσχίης ἰσχύρισις.

*Conclusions which are merely verbal cannot bear fruit,
only those based on demonstrated fact.*

— Hippocratic Precept^[189]

3

Registration of Electroanatomic Map for Image Analysis

Contents

3.1	Introduction	31
3.1.1	Commercial Procedural Registration	32
3.1.2	Registration for Research	34
3.2	Experimental Data	35
3.3	Methods	36
3.3.1	Algorithm Input	38
3.3.2	Landmark Orientation	39
3.3.3	Iterative Similarity Registration	40
3.3.4	Iterative Bézier Registration	41
3.3.5	Evaluation	41
3.4	Results	42
3.5	Conclusions	43

Chapter 3 Abstract

Background Commercial solutions exist to register 3-D image data with electroanatomic map systems for procedural guidance. However, an offline, post-procedure registration solution would facilitate study of the statistical relationships between image and electrogram data. This information could elucidate the potential

role of LGE-MRI and other advanced imaging protocols in clinical characterisation of myocardial tissue.

Methods A three-step registration algorithm was developed, which aligns electroanatomic map points to a manually segmented 3-D image. The algorithm consists of a manual landmark-based alignment, followed by iterative similarity registration of the EAM point space, and finally an iterative, 3-D Bézier field (nonrigid) registration. The similarity and Bézier field steps minimise the mean square distance from each EAM point to the cardiac chamber wall using a standard trusted region approach.

The algorithm was validated in atrial fibrillation patients undergoing pulmonary vein isolation. Each patient underwent cardiac MR ($n = 17$) or CT ($n = 20$) imaging using standard clinical protocols. Dense electroanatomic maps were acquired prior to PVI using the CARTO EAM system.

Results and Conclusions Each EAM was registered to the 3-D clinical image and the quality of registration was evaluated quantitatively by point-to-surface distance and qualitatively by visual inspection. The algorithm achieved a point-to-surface distance of 3.52 ± 3.56 mm in MR patients and 2.20 ± 3.93 mm in CT patients. The rigid and Bézier steps improved the fraction of points within 5 mm of the surface from 65% to 80% and from 56% to 85% in MR and CT cohorts, respectively. Qualitatively, the registration was clinically successful in 34 of the 37 images, with ablation points agreeing well with clinical targets. These registered EAM-image pairs were found suitable for use in the more detailed investigations in the following chapters.

3.1 Introduction

Cardiac magnetic resonance imaging is increasingly integrated with clinical and research electrophysiology. MRI is routinely used to evaluate anatomic structure, dynamic function, and scar burden in the context of arrhythmia^[190]. Integration of MRI with EP procedures is advancing rapidly. While LGE-MRI is now a clinical

standard for cardiac scar detection, development of methods to automatically quantify fibrosis for research purposes is ongoing^[191].

The utility of EAM and advanced imaging in specific clinical situations is the focus of Chapters 4 and 5, and associated background information is reserved to those chapters. Here, we consider current frameworks for registration of electroanatomic maps to clinical imagery. While commercially available tools for pre- and intra-procedural registration for interventional guidance exist and are well-studied (Section 3.1.1), the literature in post-procedure registration for research purposes is considerably more sparse (Section 3.1.2). In this chapter, we introduce a practical algorithm that contributes to this latter category.

3.1.1 Commercial Procedural Registration

Procedure

Both the CARTO product by Biosense Webster and the NavX product by St. Jude Medical (Section 2.2.2) now have the ability to import segmented MR and CT images into the EAM system for procedural guidance. These options are called CARTOMERGE and NavX Fusion, respectively. The registration process in CARTOMERGE^[192] begins with supervised segmentation. Operators select upper and lower voxel value thresholds for the blood pool, and then manually place markers ('seeds') on the image representing individual anatomic structures. This segmentation procedure takes approximately 9 minutes^[192]. The commercial algorithm then extracts endocardial surfaces and imports them into the EAM software.

These surfaces are operatively registered to the EAM coordinate system in a two-step process. First, landmarks that were previously identified on the MR or CT image – pulmonary vein antra and the mitral valve annulus – are located in the EAM system using an intracardiac catheter under fluoroscopic guidance. These paired landmark points drive an initial rigid registration, which minimises distance between the image points and the corresponding locations in the EAM coordinate system. Second, the operator uses the catheter to identify three to seven points on each segment of the endocardial wall. The system incorporates these

points using a commercial algorithm to register the endocardial surfaces in EAM- and image-space. The registration process takes about 16 minutes^[192]. Anatomic structures in this paragraph refer to the left atrium, but analogous procedures may be followed in the other cardiac chambers.

Registration using the NavX Fusion system is broadly similar^[193]. Clinical images are segmented using proprietary software, and three landmark points are manually labeled in the image as well as using a catheter in the EAM space. After a rigid landmark registration, the operator has the option to manually tag misaligned points to trigger further non-rigid registration and fusion.

Accuracy

The accuracy of these commercial systems has been evaluated in a number of studies.

In 2006, Dong and colleagues^[5] investigated the accuracy of CARTOMERGE in a canine model. They surgically attached radiopaque fiducial markers to the epicardia of nine dogs and then acquired a high-resolution CT image of each. These images were registered to electroanatomic maps using the CARTOMERGE EAM package. The researchers then performed radiofrequency ablation targeting the fiducial markers, using only the registered CT image for guidance. At autopsy, the researchers found the lesions were separated from the markers by average distances ranging from 1.8 mm in the left atrium to 2.7 mm in the right ventricle.

The same group carried out the first investigation of CARTOMERGE in human atrial fibrillation patients^[192]. Sixteen patients who presented for pulmonary vein isolation underwent MR or CT imaging within 24 hours of their ablation procedure. These images were registered to the EAM system, and the registration was evaluated on three bases: First, ‘good alignment’ of the pulmonary veins was considered achieved if every EAM point sampled within the PVs corresponded to an image point within 5 mm of the segmented vein. Second, the ability to use the registered surface to guide the catheter into the PVs and left atrial appendage (LAA) was evaluated. Third, the distance between EAM points and the registered surface (‘point-to-surface’ distance) was measured. The researchers found good PV alignment in 57 of 60 total

PVs, and complete success in navigating to the PVs and LAA using the registered surface alone. They found a mean point-to-surface distance of approximately 3 mm.

In 2007, Zhong *et al.*^[194] raised questions about the accuracy of the CARTOMERGE system in a clinical setting. After registering pre-operative CT images with the EAM system, the researchers performed mock pulmonary vein isolation procedures in 16 patients. The cardiologist was blinded to the EAM display, instead guiding the catheter through a path around the pulmonary veins using only intra-cardiac echocardiography. As the catheter moved along this path, points were recorded in the CARTO system. When compared to the registered CT data, these points were inaccurate, differing from their true anatomic locations by more than 1 cm on average. Performance was worse (errors about 2 cm) when CT frames from the end of atrial systole were not selected. Zhong hypothesized that relying on point-to-surface distances may overestimate registration performance.

The first study of the competing NavX Fusion product in human atrial fibrillation was carried out by Brooks *et al.* in 2008^[193]. The researchers reported their experience with 55 patients undergoing PVI. Before instrumentation, patients underwent CT imaging. After registration, the group used custom software to measure the surface registration error at a mean 1.9 mm. They then stabilised the catheter on various anatomic structures using fluoroscopy and evaluated its reported position in the fused EAM system. They reported a mean distance of 3.4 mm from the CT representation of the anatomic structure.

As the inter-electrode distance on cardiac ablation catheters is generally about 2 mm, registration accuracy of 3 - 4 mm is generally considered clinically acceptable, while errors in the centimetre range are subject to scrutiny.

3.1.2 Registration for Research

Offline, post-procedure registration of EAM data to clinical images presents somewhat different goals and challenges compared to procedural registration. Post-procedure registration aims to facilitate automated analysis and statistical investigation of the relationship between image data and electrogram data. Because

this research is currently most applicable when it employs delayed enhancement imaging (Section 4.1.2), offline registration does not have the luxury of using the high-contrast, blood-pool-enhanced images preferred in procedural imaging. However, post-procedure imaging benefits from availability of the full complement of a procedure’s EAM data and is comparatively unconstrained by clinical pressures on segmentation and registration time.

To the best of our knowledge, the 2013 effort by Roujol and colleagues^[195] is the only published algorithm designed explicitly for offline registration of post-procedure EAM and clinical imaging data. Their approach began with a landmark registration of manually denoted feature points. Second, the operators manually segmented both the endocardium and areas of fibrosis they identified on the LGE-MRI. A 3-D distance map was generated from the segmented endocardium and scar tissue, and the distance between the voltage map points and the endocardium distance on the one hand and low-voltage points and the scar regions on the other hand were minimised. While the distance between the registered EAM and MR endocardial surfaces was only 3.57 mm on average, individually identified landmarks had an error of 14.6 mm. This result is similar to the commercial systems, and raises questions of its utility for voxel-to-point statistical purposes.

In this chapter, we present and validate a semi-automated, three-step registration algorithm that aligns EAM data with 3-D clinical images. The algorithm was developed for the utilitarian purpose of enabling the work presented in Chapters 4 and 5. Though the procedure presented below consists of standard registration manoeuvres, it is described here not for its novelty, but instead to establish the groundwork for the chapters to come.

3.2 Experimental Data

Human left atrial data was collected in patients undergoing first-time pulmonary vein isolation at the Johns Hopkins Hospital (Baltimore, MD, USA). Each patient underwent cardiac MR or CT imaging. These images were manually segmented and distance maps were calculated. As these maps are the only image features




Figure 3.1: Illustration of the local minimum trap in segmentation. Top: Iterative alignment of blue (moving) shape to black (fixed) shape. Due to unfortunate starting position, registration settles in a local minimum of the energy function. Bottom: Initialisation close to the global minimum prevents this error.

that contributed to the registration (Section 3.3.3), imaging parameters were not dispositive to the outcomes. The imaging protocols were identical to those described in Sections 4.2.2 and 4.2.3. Electroanatomic maps were acquired using the commercial CARTO EAM system. During the ablation procedure, pulmonary vein ostia were identified and marked on the EAM. The ostia were also labeled manually on the MR or CT image for use as landmarks.

3.3 Methods

A three-step registration algorithm was developed to register EAM coordinates (moving ‘image’) to MR or CT coordinates (fixed): First, landmark alignment was used to initialise the registration. Second, the optimal similarity (rigid) transform was identified iteratively. Finally, an iterative, Bézier curve-based (nonrigid) deformation of the EAM point space was performed.




Figure 3.2: Schematic of three-step registration algorithm. Red, green, and blue diamonds: Manual image landmarks. Red, green, and blue circles: EAM landmarks. Yellow circles: EAM points. Yellow grid: deformation field of the transform.

Fundamentally, registration is an exercise in the minimisation of an energy function which represents the distance from perfect registration. Our algorithm was designed to minimise the root mean square distance (RMSD) between EAM map points and the manually segmented endocardial surface. Especially as it relied on shape information alone, our methodology was inherently vulnerable to local minima in the RMSD function, i.e. settling on areas of near-alignment (Figure 3.1). This classic pitfall is especially profound when the granularity of the transformation iterations is small compared to noise in the energy function. Our three-step process (Figure 3.2), including an initial, non-iterative step, was designed to overcome this general weakness through stepwise increases in the flexibility of the transformation calculated.

3.3.1 Algorithm Input

The registration algorithm's input consisted of three sets of data: First, paired landmarks were supplied for initial gross orientation. Three to four landmarks, which must be identifiable both procedurally on EAM as well as on the image, were manually indicated by the operator. For LA use, the pulmonary vein ostia worked well. Note that the landmark registration step (Section 3.3.2) was used only for initial alignment. If the EAM and imaging coordinate systems were already orientated by virtue of harmonious coordinate systems, this step would be rendered unnecessary.

Second, a manual blood pool segmentation of the MR or CT image was provided to the algorithm. The segmented surface was converted to a signed distance map (Section 3.3.3), and the registration calculations were based only on that distance information. This rendered the algorithm independent of voxel-level image data. The experts' manual segmentations encompassed all potential catheter locations, including adjacent vessels or chambers when the EAM data had traversed those points. Manual segmentation was chosen in order to retain maximum flexibility over the choice of clinical image modality and parameters.

Finally, of course, the algorithm required EAM point geometry. Both diagnostic mapping and ablation points were included, and points indicating the positions of

pulmonary veins and other structures were included when those structures could also be identified and segmented in the MR or CT image. Custom applications were developed in Python to extract this data from the export files produced by the commercial EAM system.

3.3.2 Landmark Orientation

An initial landmark registration was employed to increase the likelihood of the subsequent iterative registration converging to the global minimum. This was especially important in the LV, which generally demonstrates a high degree of rotational symmetry. A similarity transform was designed using manually identified, paired landmark points.

Following Umeyama^[196], we analytically computed the transform parameters – rotation matrix $\mathbf{R}$, translation vector $\mathbf{t}$, and scale constant c – which minimised the mean squared landmark error,

$$e^2(\mathbf{R}, \mathbf{t}, c) = \frac{1}{n} \sum_{i=1}^n \|\mathbf{y}_i - (c\mathbf{R}\mathbf{x}_i + \mathbf{t})\|^2, \quad (3.1)$$

between the n paired landmark coordinate vectors $\mathbf{x}_i$ and $\mathbf{y}_i$ in the EAM space and image space, respectively. The closed-form solution is fully derived by Umeyama. Briefly, consider the mean coordinates $\boldsymbol{\mu}$, variances σ^2 , and covariance matrix $\boldsymbol{\Sigma}_{xy}$ of the coordinate vectors:

$$\boldsymbol{\mu}_x = \frac{1}{n} \sum_{i=1}^n \mathbf{x}_i \quad \boldsymbol{\mu}_y = \frac{1}{n} \sum_{i=1}^n \mathbf{y}_i \quad (3.2)$$

$$\sigma_x^2 = \frac{1}{n} \sum_{i=1}^n \|\mathbf{x}_i - \boldsymbol{\mu}_x\|^2 \quad \sigma_y^2 = \frac{1}{n} \sum_{i=1}^n \|\mathbf{y}_i - \boldsymbol{\mu}_y\|^2 \quad (3.3)$$

$$\boldsymbol{\Sigma}_{xy} = \frac{1}{n} \sum_{i=1}^n (\mathbf{y}_i - \boldsymbol{\mu}_y)(\mathbf{x}_i - \boldsymbol{\mu}_x)^T \quad (3.4)$$

Computing the singular value decomposition of the covariance matrix, $\boldsymbol{\Sigma}_{xy} = \mathbf{U}\mathbf{D}\mathbf{V}^T$, and selecting

$$\mathbf{S} = \begin{cases} \mathbf{I} & \text{if } \det(\mathbf{U}) \det(\mathbf{V}) = 1 \\ \text{diag}(1, 1, \dots, 1, -1) & \text{if } \det(\mathbf{U}) \det(\mathbf{V}) = -1 \end{cases}, \quad (3.5)$$

we find the similarity transform parameters are

$$\mathbf{R} = \mathbf{U}\mathbf{S}\mathbf{V}^T, \quad \mathbf{t} = \boldsymbol{\mu}_y - c\mathbf{R}\boldsymbol{\mu}_x, \quad \text{and} \quad c = \frac{1}{\sigma_x^2} \text{tr}(\mathbf{D}\mathbf{S}). \quad (3.6)$$

As noted above, the landmark registration may be modified when the EAM and imaging apparatus use physical-unit, anatomic-direction coordinate systems. As the rotation and scale would be roughly concordant without calculation in this case, the translation parameter alone could be determined by aligning the average EAM point coordinate with the most interior segmentation point (i.e. the most ‘internal’ point on the signed distance map).

3.3.3 Iterative Similarity Registration

The blood pool was segmented manually from each MR or CT image by experts using Seg3D software (Version 2.1.4, University of Utah, Salt Lake City, USA). A custom application written in C++ was used to convert this segmentation mask into a signed distance map using Danielsson’s method^[197] as implemented in the open-source Insight Toolkit^[198].

A similarity transform with iteratively evolving parameters was applied to the EAM point coordinates to minimise mean squared distance between each of the n map points and the blood pool boundary. For simplicity, the EAM point coordinates $\mathbf{p}_i$ in this section were taken to have already been oriented using the landmark transformation from Section 3.3.2.

The multi-level, iterative algorithm sought similarity transform parameters $\mathbf{R}, \mathbf{t}, c$ to minimise the mean square surface error

$$e^2(\mathbf{R}, \mathbf{t}, c) = \frac{1}{n} \sum_{i=1}^n [D(c\mathbf{R}\mathbf{p}_i + \mathbf{t})]^2, \quad (3.7)$$

where $D(\mathbf{p})$ represents the signed distance map. Rotations $\mathbf{R}$ were applied about the mean EAM coordinate to prevent conflation of translation and rotation.

Every iteration, uniform variations in each of the seven similarity parameters (three translation, three rotation, and one scale) were independently evaluated for their effect on $e^2(\mathbf{r}, \mathbf{t}, c)$. The transformation yielding a minimum e^2 was applied,

and the next iteration began. When a local minimum was reached, the scale of the variations were decreased by half. The registration was considered to have converged when the scale of parameter perturbations decreased to half of the image voxel size.

3.3.4 Iterative Bézier Registration

The third and final registration step was another iterative procedure to minimise the mean square distance from EAM points to the endocardial surface. A Bézier field $\mathbf{B}$ was defined using regularly spaced control points $\mathbf{b}_{i,j,k}$. The field vector at any point $\mathbf{p}$ within the field was computed:

$$\mathbf{B}(\mathbf{p}) = \sum_{i=0}^{n_i} \sum_{j=0}^{n_j} \sum_{k=0}^{n_k} B_i^{n_i}(\mathbf{p}_u) B_j^{n_j}(\mathbf{p}_v) B_k^{n_k}(\mathbf{p}_w) \cdot \mathbf{b}_{i,j,k}. \quad (3.8)$$

Here n_i , n_j , and n_k represent the order of the control point grid in each dimension; and $\mathbf{p}_u$, $\mathbf{p}_v$, and $\mathbf{p}_w \in [0, 1]$ are the normalised coordinates of $\mathbf{p}$. $B_i^n(u)$ is the Bernstein polynomial,

$$B_i^n(u) = \binom{n}{i} u^i (1-u)^{n-i}, \quad (3.9)$$

where $\binom{n}{i}$ represents the binomial coefficient. For every iteration, the effect of translating each control point $\mathbf{b}_{i,j,k}$ within a local trust region was evaluated in turn on the mean square error $e^2(\mathbf{B})$:

$$e^2(\mathbf{B}) = \frac{1}{n} \sum_{i=1}^n [D(\mathbf{p}_i + \mathbf{B}(\mathbf{p}_i))]^2 \quad (3.10)$$

and the position effecting the lowest $e^2(\mathbf{B})$ is applied. When an iteration fails to decrease $e^2(\mathbf{B})$, the translation scale is reduced by half. When the minimum scale size is reached, the algorithm ends.

3.3.5 Evaluation

With the aim of demonstrating its utility for Chapters 4 and 5, the algorithm was validated quantitatively by point-to-surface distance and qualitatively by agreement of ablation points with the clinical procedure in the left atrium of AF patients undergoing first PVI procedure.




Figure 3.3: Registration accuracy in MR cohort (median as well as 10th, 25th, 75th, and 90th percentiles).

The distance from the registered coordinates of each EAM point to the manually segmented endocardial contour was calculated by reference to the Danielsson distance map after landmark registration, iterative similarity registration, and iterative Bézier registration.

3.4 Results

In the MR cohort, the mean point-to-surface registration error was 4.7 ± 4.4 mm for landmark registration alone (mean $\pm$ standard deviation). This improved to 4.0 ± 4.4 mm after similarity registration and 3.5 ± 4.2 mm after Bézier registration. This modest decrease in mean error from 4.7 mm to 3.5 mm is not clinically significant on its face, as the uncertainty of the electroanatomic mapping system itself is often estimated at about 5 mm.

However, inspecting the distribution of errors (Figure 3.3) reveals a somewhat more substantial improvement. Though landmark registration alone achieves reasonable success on average, a significant number of points remain distant from the surface – the 90th percentile is 9.3 mm. After similarity and Bézier registration, this decreases to 8.0 mm. These registration steps bring the fraction of points within 5 mm of the surface from 65% with landmark registration alone, to 74% after similarity registration, and finally to 80% after non-rigid registration.

The improvement is somewhat more pronounced in the CT cohort. Mean point-to-surface registration distance improved from 6.2 ± 6.7 mm for landmark registration to 3.5 ± 4.7 mm and 3.0 ± 4.7 mm after similarity and Bézier registration,




Figure 3.4: Registration accuracy in CT cohort (median as well as 10th, 25th, 75th, and 90th percentiles).

respectively. The 90th percentile distance was 13.6 mm after landmark registration alone, compared with 6.8 mm after rigid and non-rigid registrations (Figure 3.4). The fraction of points closer than 5 mm improved from 56% to 78% and 85% after rigid and non-rigid transformations.

Qualitatively, the algorithm performed adequately. Ablation points were well aligned with their clinical targets and the registration results deemed clinically acceptable for use in the investigations in Chapters 4 and 5 in sixteen of seventeen MR images and eighteen of twenty CT images.

Overt registration failure occurred in one patient from the MR cohort and two from the CT cohort. The major observed failure mode (two patients) of this method occurred in patients with particularly flat EAM geometry due to sparse electroanatomic maps. With little curvature to provide a geometric reference to the cardiac chamber surface, these points had no geometric ‘anchor’ to their positions on the surface.

3.5 Conclusions

The registration algorithm presented here, while not fundamentally novel in its approach, is an effective, semi-supervised method to register EAM points to 3-D image data. Although the improvement in mean point-to-surface distance through the registration pipeline was modest, there was a significant improvement in the fraction of points brought within 5 mm of the manually segmented endocardial surface.

Offline registration tools are a necessary step in the point-by-point study of the relationship between EAM measurements and image features, and the current literature in this area is sparse. The registration algorithm achieved qualitative clinical success, based on agreement of ablation and map point positions to the physician's intent during the procedure, in 34 of 37 images. These successful registrations were used as the basis for the investigations in Chapters 4 and 5.

*The goal is to turn data into information, and
information into insight.*

— Carly Fiorina

4

Identification of Left Atrial Fibrosis Using Image Data

Contents

4.1	Introduction	47
4.1.1	Electroanatomic Mapping for Atrial Fibrosis	47
4.1.2	Correlating Image Data with Atrial Fibrosis	48
4.1.3	Random Forests	52
4.1.4	Receiver Operating Characteristic Curves	54
4.1.5	Principal Component Analysis	55
4.2	Experimental Data	56
4.2.1	Patients	56
4.2.2	Magnetic Resonance Imaging	56
4.2.3	Computed Tomography	57
4.2.4	Manual Segmentation	57
4.2.5	Electroanatomic Mapping	59
4.3	Methods	59
4.3.1	Calculating Image Features	59
4.3.2	Framework for Classifier Evaluation	61
4.3.3	Single-Parameter Regression	63
4.3.4	Random Forests	65
4.3.5	Principal Component Analysis	66
4.4	Results – MR Cohort	67
4.4.1	Electroanatomic Map and Image Parameters	67
4.4.2	Single-Parameter Regression	68
4.4.3	Random Forest	71
4.4.4	Principal Component Analysis	73
4.5	Results – CT Cohort	78
4.5.1	Electroanatomic Map and Image Parameters	78

4.5.2	Single-Parameter Regression	79
4.5.3	Random Forest	81
4.5.4	Principal Component Analysis	83
4.6	Conclusions	88
4.6.1	Single-Parameter Models	88
4.6.2	Random Forest	88
4.6.3	Principal Components Analysis	89
4.6.4	Future Directions	90

Chapter 4 Abstract

Background In atrial fibrillation patients, left atrial fibrosis correlates with disease that is more severe and refractory to treatment. Fibrosis is traditionally evaluated via intracardiac electroanatomic mapping. Recent studies have established correlation between atrial fibrosis and enhancement on magnetic resonance images. Delayed enhancement multi-detector computed tomography also merits evaluation in the left atrium, as it has recently been implicated in injury assessment in the ventricular myocardium. In this chapter, we analyse the relationship between MR and CT image features and bipolar voltage using single-parameter regression and random forest modelling.

Methods Sixteen AF patients underwent MR and eighteen underwent CT before EAM and pulmonary vein isolation. Images and EAMs were registered, and the LA myocardium was manually segmented from the images. Neighbourhoods were defined around each registered EAM point. The mean, maximum, and minimum voxel intensities, as well as voxel variation across the neighbourhood, were extracted from each neighbourhood and combined to form the feature vector for that point. Single-parameter regression and random forests were used to devise models of atrial fibrosis and bipolar voltage. Principal component analysis was used to characterise the influence of combinations of features on the random forest approach.

Results and Conclusions Random forest performed significantly better than the classifier based on single-parameter regression at identifying fibrosis (area under ROC curve (AUC) 0.746 vs. 0.621) and modelling bipolar voltage (R^2 0.190 vs. 0.034) using MR data. This relationship was more profound using CT data, with AUC 0.977 vs. 0.609 and R^2 0.770 vs. 0.043. The improved prediction of atrial fibrosis and bipolar voltage using image data raises the prospect of improved patient selection for invasive treatment and more comprehensive evaluation of fibrosis in groups not undergoing cardiac instrumentation.

4.1 Introduction

Left atrial fibrosis correlates with the severity of atrial fibrillation and recurrence after transcatheter pulmonary vein isolation (PVI). The gold standard for identification of fibrosis is intracardiac electroanatomic mapping (EAM, Section 4.1.1). However, the role of noninvasive imaging (particularly MR and CT) for this purpose is under active investigation (Section 4.1.2).

4.1.1 Electroanatomic Mapping for Atrial Fibrosis

The progressive natural history of atrial fibrillation has long raised suspicions about the role of the atrial extracellular matrix. In 2004, Boldt *et al.*^[199] showed clear differences in the histology of patients in sinus rhythm as compared with those with AF. The atrial cardiomyocytes of AF patients were surrounded by thick connective tissue fibres, with a clear increase in collagen by Western blot. Boldt concluded that this fibrosis results in physical separation and decreased conduction between atrial cardiomyocytes, leading to the observed decrease in bipolar peak-to-peak voltage on local electrograms during electroanatomic mapping and constituting a major factor in the observed self-perpetuating pathogenesis of AF.

Establishment of a ‘normal’ LA bipolar voltage threshold for the purpose of fibrosis identification in AF patients is ongoing^[200]. Verma *et al.*, in their work^[201] demonstrating that low left atrial voltage predicts PVI treatment failure (Section 5.1.1), defined ‘scar’ as bipolar voltage $BV \leq 0.05$ mV and ‘low-voltage

regions’ as $BV \leq 0.5$ mV. The group found that their ‘left atrial scar (LAS)’ cohort had extensive areas of scar associated with large low-voltage regions. In contrast, only 5% of their ‘non-LAS’ group showed scar, and even those regions covered $< 5\%$ of the LA area.

While other groups also use Verma’s 0.5 mV bipolar voltage threshold for ‘normal’ LA myocardium^[202,203], work by Kapa *et al.* suggests a 95% threshold between 0.27 mV and 0.45 mV depending on the region of the LA myocardium^[204]. They hypothesize that tissue thickness and adjoining conductive structures (pulmonary veins and heart valves) may affect normal voltages.

4.1.2 Correlating Image Data with Atrial Fibrosis

Non-invasive imaging is a tempting alternative to catheter electroanatomic mapping for evaluation of atrial fibrosis and scar^[205]. Although the mapping step is not thought to be an independent contributor to the risks of the PVI procedure^[206–208], the ability to more accurately risk-stratify patients in a non-invasive manner could improve the overall PVI benefit equation^[205]. Furthermore, effective non-invasive mapping of atrial fibrosis and scar would allow comparison of AF patients with groups that do not warrant the risks of catheter instrumentation.

Magnetic Resonance Imaging

Late gadolinium enhancement MRI (LGE-MRI) is the primary modality under investigation for identification of atrial fibrosis. In this technique, a gadolinium-based (T_1 -shortening) contrast agent is injected into the patient, and a T_1 -weighted MR image is acquired several minutes later. While clinical protocols vary, the unifying result is that the contrast agent is retained in areas – such as fibrotic tissue – with reduced bloodflow and increased extracellular matrix.

In 2007, Peters and colleagues^[209] at Beth Israel Deaconess Medical Center and Harvard Medical School first applied LGE-MRI to detect atrial scar in AF patients. Their work focused on recognising scar caused by ablation, rather than native fibrosis. The group imaged 15 patients before their ablation and 18 patients at least

30 days after pulmonary vein isolation (23 total patients). Of the 14 interpretable pre-ablation images, none demonstrated atrial wall or PV enhancement, whereas all of the 13 interpretable post-ablation images demonstrated enhancement of the PV ostia or LA wall. The images were compared to recorded ablation sites in one patient, and agreed well. In a circumferential view of the PV ostia, an average 88% of each ostium, by sector, was enhanced. This study marked the first use of LGE-MRI to identify atrial scar.

Using LGE-MRI to evaluate post-ablation scar remains a valuable technique. In 2014, Sohns and colleagues^[210] at King's College London characterised the relationship between catheter contact force and quality of scar following PVI for paroxysmal AF. Experts manually segmented enhanced areas from 3-month post-ablation LGE images in six patients. Quality of scar was determined quantitatively by the ratio of voxels identified as enhancement within 5 mm of the ablation point and qualitatively by expert impression. The researchers found significant correlation between catheter contact force and enhancement.

In 2009, Marrouche and colleagues^[211] at the University of Utah transitioned from identifying ablation scar to pathogenic fibrosis. They analysed 3-D LGE-MRI of 81 patients presenting for PVI in addition to six healthy volunteers. The researchers segmented the LA from the MR images manually, and a threshold method was used to determine the extent of fibrosis. The group defined 'normal' tissue by MRI as voxels with intensity between 2% and 40% of the maximum voxel intensity in the LA myocardium. The 'enhancement' threshold is then selected by an expert between two and four standard deviations above the mean 'normal' tissue intensity. The most commonly selected threshold was three standard deviations^[211].

The researchers stratified patients into three categories on the basis of the extent of enhancement by tissue area: minimal enhancement ($n = 43$, area $8.0 \pm 4.2\%$), moderate enhancement ($n = 30$, area $21.3 \pm 5.8\%$), and extensive enhancement ($n = 8$, area $50.5 \pm 15.4\%$). The rate of AF recurrence after PVI was 14.0% for those with minimal enhancement, 43.3% for the moderate enhancement group, and 75% for extensive enhancement. Correlation between fibrosis as identified by MR

and on EAM was also established in a descriptive, region-based approach. While this study launched the use of LGE-MRI in PVI stratification, its generality is limited by the use of expert knowledge to determine a threshold for ‘enhancement’. Questions have also arisen regarding the added value of imaging data compared to clinical history for risk stratification^[200,212].

Marrouche and colleagues validated their technique in the DECAAF study, a multicentre, prospective, observational cohort study of 329 AF patients undergoing their first PVI procedure published in 2014^[213]. In the 272 patients with adequate images, image quantification was performed as above, and patients were categorised by off-site experts as enhancement stage 1 ($n = 49$, less than 10% of the atrial wall), stage 2 ($n = 107$, 10 - 20%), stage 3 ($n = 80$, 20 - 30%), or stage 4 ($n = 94$, 30% or more). Treating physicians were blinded to this evaluation. For the 260 patients who followed up, there was a 1.06 hazard ratio for recurrence for each additional 1% of atrial fibrosis by MRI. The absolute recurrence risk after 325 days was 0.15 for those in stage 1 compared to 0.51 for those in stage 4. Similar trends persisted through all groups and after covariate adjustment.

In 2012, Nazarian and colleagues^[214] at Johns Hopkins University introduced the image intensity ratio (IIR), an alternative method of quantification of left atrial enhancement on LGE-MRI. The IIR is the quotient of the local image intensity and the mean intensity of the LA blood pool. The researchers acquired LGE-MRI for 30 patients before ablation procedures for AF. After manual segmentation of the LA wall, the IIR was computed in 20 sectors of each image. These sectors were correlated with EAM points, and local IIR was found to associate with bipolar voltage on EAM (slope -1.4 mV per IIR unit, $p < 0.001$). The same researchers further validated the IIR technique in a group of 75 patients using essentially the same methods and with similar results in 2014^[215]. Although this method avoids the need for expert input, the work establishes only significance of correlation, not amount of variance explained or predictive capacity between patients.

While these studies have laid the groundwork in establishing correlation between LGE-MRI enhancement, low bipolar voltage by EAM, and AF recurrence after PVI,

none have yet attempted to establish efficacy of image information in predicting bipolar voltage or fibrosis. Further, work to date has been limited to regional analyses; point-by-point correlation has not been deeply investigated. The work presented in this chapter evaluates the performance of modern machine learning algorithms to predict bipolar voltage data and fibrosis using image features.

X-ray Computed Tomography

Unlike MR imaging, CT is not under widespread investigation for detection of atrial fibrosis^[216]. Its primary role in the AF treatment pipeline is to assess atrial anatomy and volume, especially in patients with contraindications to MRI. However, contrast-enhanced multidetector computed tomography (CE-MDCT; also in some literature DE-MSCT for ‘delayed enhancement multislice’) has been proposed as an alternative to LGE-MRI for detection of injury to the ventricular myocardium^[83]. A growing body of literature supports CE-MDCT for assessment of myocardial infarction^[12,83], nonischemic cardiomyopathy^[217], and hypertrophic obstructive cardiomyopathy^[218].

In 2010, Tian and colleagues at the University of Maryland^[219] applied CE-MDCT to electrophysiology in a study on ventricular fibrillation (VF) ablation. They studied 11 patients presenting for VF ablation. CT images were acquired and registered to electroanatomic maps using a 17-segment myocardial model. The researchers only observed hyperenhancement in a minority of patients, instead finding that hypoperfusion correlated well with abnormal ventricular voltage (< 1.5 mV) on EAM.

The work below is, to the best of our knowledge, the first comprehensive evaluation of CE-MDCT in the left atrium. A preliminary related study was presented at the American Heart Association 2013 conference^[15]. That study incorporated a simple, single-feature regression model on a subset of the clinical data used below. It did not incorporate multiple image features, machine learning algorithms, principal component analysis, or prediction-based performance metrics.

4.1.3 Random Forests

Random forests is an ensemble machine learning technique for classification and regression. The random forests method constructs a large collection of individual decision trees, using bootstrap aggregation and random feature subspace splitting (Figure 4.1). The approach is described fully by Breiman *et al.*^[220] and is well known in the machine learning field due to its accuracy, robustness to noise, and computational efficiency. An example-based illustration of random forests is found in Section C.1, page 175.

In general terms, a random forest is constructed using a training dataset which consists of feature vectors of arbitrary length X_i and responses (i.e. class memberships for a classifier, dependent variable value for a regressor) y_i . In an experimental context, the entire experimental dataset is usually split into the aforementioned training set, used to construct the random forest, and a testing set, used to evaluate its performance. This avoids the problem of an aggressive machine learning algorithm overfitting its training data and boasting artificially impressive results.

Bootstrap aggregation (‘bagging’) is performed on a per-tree basis. For each tree, a bootstrap sample is selected randomly from the training dataset with replacement. Random feature selection is incorporated into the tree learning algorithm. At each candidate split, a random subset of features is selected, and the feature from that subset which provides an optimal (by Gini impurity or another metric) binary split is used to define the node. Each tree is built until all leaves are pure. The random forest technique is relatively robust to the number of features selected at each node. Breiman’s original paper uses $\log_2 f$, where f is the total number of features. Most modern implementations use $\sqrt{f}$. In our application ($f = 38$), the difference between these values is insignificant.

To apply a trained random forest for prediction (or evaluation), the feature vector for a novel data point (or a point from the reserved testing dataset) is passed through each tree of the forest. The response of each tree is averaged,




Figure 4.1: Schematic of the random forest method as used in this work.




Figure 4.2: Example receiver operating characteristic curve.

and the average response is the output of the random forest (for a classifier, the highest-probability class is selected).

A final advantage of the random forest methodology is that it allows evaluation of the importance of each feature in the feature vector. During training, the ‘out-of-bag error’ for each data point is computed. To measure the importance of a given feature, that feature is permuted among the training data, and the error is again computed. Those features which produce substantially higher error upon permutation are ranked more important.

4.1.4 Receiver Operating Characteristic Curves

The receiver operating characteristic (ROC) curve is a standard means of evaluating a classifier^[221,222]. It is especially popular in biomedical applications, where the complex interplay between sensitivity, specificity, and prevalence confound quick comparisons of tests. Fundamentally, the ROC curve is a descriptor of a classifier drawn on axes of true positive rate (sensitivity, y-axis) and false positive rate ($1 - \text{specificity}$, x-axis).

A static classifier on these axes is a point, but in most relevant cases some threshold value (blood chemistry level, antibody titre, millimeters of ST elevation,

etc.) can be varied continuously between a value that always registers a ‘negative’ result – the (0,0) point – and one that always registers a ‘positive’ result – the (1,1) point. The diagonal line connecting these points represents tests which are uncorrelated with the gold standard. Higher correlation shifts the curve into the upper left region of the ROC space (Figure 4.2). The area under this curve (AUC) is a standard measure of classifier performance. An AUC of 1.0 represents 100% accuracy, while an AUC of 0.5 represents guesswork.

A variety of methods of generating confidence intervals around ROC curves have been proposed. Macskassy and Provost^[223] compared a number of these methods and found that empirically constructing a fixed-width band around an average ROC curve, as proposed by Campbell^[224], resulted in a confidence interval that was consistent and robust to the number of samples used to generate the band. The fixed-width band method is used in this work, and is more fully described in Section 4.3.2.

4.1.5 Principal Component Analysis

Principal component analysis (PCA) is a statistical technique which transforms a set of (possibly) correlated variables into a set of linearly uncorrelated values using an orthogonal transformation^[225]. It is defined such that the components are ordered in decreasing order of variance explained in the data. An example-based illustration of PCA is found in Section C.2, page 181.

Formally, given a data matrix $\mathbf{X}$, with each of n rows representing a data point and each of f columns representing a (zero-mean) feature of that point, the PCA transformation is defined by the matrix $\mathbf{W}$ and maps $\mathbf{X}$ to a transformed data matrix $\mathbf{T}$:

$$\mathbf{T} = \mathbf{X}\mathbf{W} . \quad (4.1)$$

$\mathbf{W}$ is defined by unit loading vectors $\mathbf{w}_{(k)}$ which individually define the weights of each feature in component k . The first loading vector is

$$\mathbf{w}_{(1)} = \arg \max_{\|\mathbf{w}\|=1} \left(\|\mathbf{X}\mathbf{w}\|^2 \right) . \quad (4.2)$$

Subsequent components are calculated by first subtracting higher component contributions from $\mathbf{X}$,

$$\hat{\mathbf{X}}_k = \mathbf{X} - \sum_{s=1}^{k-1} \mathbf{X} \mathbf{w}_{(s)} \mathbf{w}_{(s)}^T, \quad (4.3)$$

and then calculating the (necessarily orthogonal) component weighting that corresponds to the most remaining variance,

$$\mathbf{w}_{(k)} = \arg \max_{\|\mathbf{w}\|=1} \left(\|\hat{\mathbf{X}}_{k-1} \mathbf{w}\|^2 \right). \quad (4.4)$$

As an unsupervised technique, PCA follows the structure inherent in the data, and so avoids operator bias in determining feature relationships. This work uses PCA to account for the independent contribution of image features in predicting bipolar voltage values in the left atrial myocardium.

4.2 Experimental Data

4.2.1 Patients

Patients undergoing first-time pulmonary vein isolation procedures for atrial fibrillation at the Johns Hopkins Electrophysiology and Arrhythmia Service (Baltimore, MD, USA) were offered enrolment in this study. All patients consented to the investigation, and the ablation and imaging protocols were approved by the Johns Hopkins Institutional Review Board (IRB study identifiers NA-00019442, NA-00021383, and NA-00070577).

4.2.2 Magnetic Resonance Imaging

In patients who underwent MR analysis ($n = 16$), late gadolinium enhancement MR imaging (LGE-MRI) was performed 3-5 days prior to catheter ablation. A commercially available 1.5-Tesla MR scanner (Avanto, Siemens, Erlangen, Germany) was used with a 6-channel phased-array torso coil in combination with the spine matrix coil for a total of 12 independent receiver channels. A contrast-enhanced 3-D fast low-angle MR angiography sequence was obtained immediately following administration of 0.1 mmol/kg of intravenous gadopentetate dimeglumine (Bayer

Healthcare Pharmaceuticals, Montville, NJ). A further 0.1 mmol/kg of contrast was injected (total 0.2 mmol/kg) and LGE-MRI scans were acquired 15 minutes after injection.

The LGE-MRI scans were acquired via a 3-D inversion recovery prepared fast spoiled gradient recalled, respiratory navigated, ECG-gated, fat-suppressed sequence (repetition time 500-700 ms, echo time 1.5 ms, field of view 340 mm, flip angle 101, in-plane resolution 1.3×1.3 mm, slice thickness 2.0 mm). The respiratory navigator was positioned on the right hemidiaphragm and trigger time was selected to correspond with atrial diastole as evaluated in cine images. Inversion time was determined for optimal nulling of the myocardium (280-300 ms).

4.2.3 Computed Tomography

In patients who underwent CT analysis ($n = 18$), contrast-enhanced multidetector computed tomography (CE-MDCT) was performed with a commercially available 320-detector CT scanner (One Aquilion, Toshiba Medical Systems, Otawara, Japan) prior to the ablation procedure. Prior to the CE-MDCT study, heart rate was optimised to less than 80 bpm using a beta-blockade protocol^[226]. Noniodinated intravenous contrast (Isovue 370; Bracco Diagnostics, Princeton, NJ) was administered following a triphase protocol at a rate of 5-6 mL/sec in the following sequence: 20 mL saline test injection, 40 mL (100% contrast), 20 mL (50% saline, 50% contrast), 65 mL (30% saline, 70% contrast), and followed by 30 mL saline flush. Image acquisition was gated to 40% of the R-R interval during a breath-hold. Tube voltage was 120 kV, and tube amperage ranged from 400 to 580 mA depending on body habitus.

4.2.4 Manual Segmentation

Clinical experts using Seg3D software (Version 2.1.4, University of Utah, Salt Lake City, USA) manually contoured the endocardial and epicardial borders in all LGE-MRI and CE-MDCT images. First, the whole heart blood pool was segmented using manually supervised region-growing. The left atrial region, including pulmonary




Figure 4.3: Left atrial segmentation in a CE-MDCT image. (a) Axial CT slice. (b) Supervised region-growing segmentation of the blood pool (red). (c) Manual delineation of the left atrial epicardial surface (blue). (d) Myocardial mask generated by excluding blood volume from epicardial segmentation (green).

veins, was manually reviewed and corrected as necessary. A separate, manual delineation of the epicardial surface was then performed, carefully excluding epicardial fat. The epicardial segmentation, with the blood pool volume excluded, was taken as the myocardial mask (Figure 4.3). Although every effort was taken to ensure accurate segmentation, the epicardial boundary can be difficult to discern, and at times the operator's expert knowledge of the LA shape and myocardial thickness was used to delineate that boundary. As the atrial myocardium is only 1.5 to 2 mm thick^[227], even minor inaccuracies could hamper analysis.

4.2.5 Electroanatomic Mapping

LA endocardial electroanatomic mapping was performed in atrial fibrillation patients in conjunction with PVI ablation procedures. At the time of ablation, patients underwent double transseptal puncture using standard techniques. The CARTO 3 commercial electroanatomic mapping system was used in conjunction with an irrigated ablation catheter (Thermo-Cool, Biosense Webster) with a 3.5 mm tip and 2 mm inter-electrode spacing to record LA endocardial voltage maps. Patients with persistent AF were referred for external cardioversion 3–4 weeks prior to the procedure so all maps were obtained in sinus rhythm. Bipolar electrograms were filtered at 30 to 400 Hz.

The electroanatomic map coordinate space was registered to the MR or CT image coordinate space as appropriate using the technique described in Chapter 3. Subsequently, the image coordinates of each EAM point were taken as the nearest point within the manually-segmented LA myocardium. This step facilitates consistent evaluation of local image parameters and compensates for beat-to-beat variation which necessarily introduces small discrepancies between EAM catheter position and image myocardial position, even after registration.

4.3 Methods

4.3.1 Calculating Image Features

For each EAM point, corresponding image features were calculated as follows. Each 3-D image, along with the binary mask M representing its manual LA myocardial segmentation, was loaded into custom software written in Python. For each EAM point p , ten spherical regions $R_{p,d}$ with diameter $d \in [1, 2, \dots, 10 \text{ mm}]$ were defined centered about p . Each of the n voxels $v \in M \cap R$ located within the myocardial mask and with midpoints falling inside the spherical region (Figure 4.4) were included in the calculation of image features for that point and scale (i.e. region diameter).

For each region $R_{p,d}$, the mean, maximum, and minimum voxel values were calculated. Local voxel variation was computed as the standard deviation without




Figure 4.4: Schematic view of image feature region. Image features for a given point p and diameter d are calculated using the image voxels (blue) within the intersection of the myocardial contour M (green) and the spherical neighbourhood $R_{p,d}$ (red).

Bessel's correction. To exclude cases in which few voxels contribute local variation, $v_{p,d}^\sigma$ was only calculated for regions with $d \geq 3$ mm. The image features were

$$\begin{aligned}
 v_{p,d}^{avg} &= \frac{1}{n} \sum_{v \in M \cup R_{p,d}} v, & v_{p,d}^{min} &= \min_{v \in M \cup R_{p,d}} v, \\
 v_{p,d}^\sigma &= \sqrt{\frac{1}{n} \sum_{v \in M \cup R_{p,d}} (v - v_{p,d}^{avg})^2}, \text{ and} & v_{p,d}^{max} &= \max_{v \in M \cup R_{p,d}} v.
 \end{aligned} \tag{4.5}$$

These features were selected for comparison with the prevailing research on image-based fibrosis identification (Section 4.1.2). By building both single-parameter regression models and machine learning models using this single set of features, the intrinsic value of a machine learning approach could be evaluated.

Normalisation

As detailed in Section 4.1.2, current research favours two methods of normalising image voxel values for the purpose of correlation with local electrogram data. This work adapts both the IIR method^[214] and a streamlined Utah method^[211] for comparison.

Image Intensity Ratio Method In the IIR method^[214], values are divided by the mean voxel value of the LA blood pool $\bar{b}$:

$$\hat{v}_{p,d}^{avg} = \frac{v_{p,d}^{avg}}{\bar{b}}, \quad \hat{v}_{p,d}^{min} = \frac{v_{p,d}^{min}}{\bar{b}}, \quad \hat{v}_{p,d}^{max} = \frac{v_{p,d}^{max}}{\bar{b}}, \text{ and } \hat{v}_{p,d}^{\sigma} = \frac{v_{p,d}^{\sigma}}{\bar{b}}. \quad (4.6)$$

To calculate $\bar{b}$, a Danielsson distance map^[197] was computed from each manually segmented myocardium using a custom application written in C++. A spherical blood pool sample region B , centred on the maximum-distance point and with radius 80% of the computed distance to the myocardium, was defined. The 80% factor accounts for artefacts at the periphery of the blood pool and slight inaccuracies possible in the Danielsson method. The mean voxel value within B was taken as $\bar{b}$.

Utah Method In the Utah method^[211] as adapted here, myocardial intensity values are normalised to their z-scores – standard deviations above the myocardial mean $\bar{m}$ on a per-image basis. Analogously, the variation $v_{p,d}^{\sigma}$ is normalised by the overall myocardial variation:

$$\begin{aligned} \hat{v}_{p,d}^{avg} &= \frac{v_{p,d}^{avg} - \bar{m}}{\sigma_m}, & \hat{v}_{p,d}^{min} &= \frac{v_{p,d}^{min} - \bar{m}}{\sigma_m}, \\ \hat{v}_{p,d}^{\sigma} &= \frac{v_{p,d}^{\sigma}}{\sigma_m}, \text{ and } & \hat{v}_{p,d}^{max} &= \frac{v_{p,d}^{max} - \bar{m}}{\sigma_m}. \end{aligned} \quad (4.7)$$

All image voxels falling within the manual LA myocardial segmentation mask were used to calculate $\bar{m}$ and σ_m for each image. Note that while the original Utah method calls for expert thresholding^[228], we use the z-score as a continuous feature.

In summary, for each normalisation method, there are 38 image parameters associated with each point p :

$$\begin{aligned} \hat{v}_{p,d}^{avg}, \quad d \in [1, 2, \dots, 10 \text{ mm}]; & \quad \hat{v}_{p,d}^{min}, \quad d \in [1, 2, \dots, 10 \text{ mm}]; \\ \hat{v}_{p,d}^{\sigma}, \quad d \in [3, 4, \dots, 10 \text{ mm}]; \text{ and } & \quad \hat{v}_{p,d}^{max}, \quad d \in [1, 2, \dots, 10 \text{ mm}]. \end{aligned} \quad (4.8)$$

4.3.2 Framework for Classifier Evaluation

In this section, we detail the strategies employed to predict fibrosis (i.e. low bipolar voltage) at a given EAM point on the basis of associated image features. Each strategy was evaluated in an analogous manner: The EAM data points (with

their associated image features) were randomly divided into training (75%) and testing (25%) sets. The classifier or regression model was fitted to the training set, and then evaluated using the testing set, and an ROC curve was generated. For probabilistic classification methods (e.g. random forest classifier), the threshold fibrosis class probability score was taken as the ROC curve parameter. For regression methods (e.g. linear model of log bipolar voltage), the threshold predicted bipolar voltage was the parameter.

This evaluation process, including randomisation of test and training sets, was repeated 100 times for each method being evaluated. Adapting Campbell^[224], a composite ROC plot was produced composed of a fixed-width band which entirely contained at least 68 of the 100 individual ROC curves. (The value 68 was chosen in analogy with the standard deviation error bars used elsewhere in this work.) The centre of the fixed-width band was taken as a composite ROC plot, and its area under the curve (AUC) was used as the primary evaluation metric for the method.

Calculating Fixed-Width Bands

As detailed in Section 4.1.4, the fixed-width band technique is effective in generating confidence bands around an averaged ROC curve. The width of a band is determined empirically, and the direction of the band is the slope $b = -\sqrt{m/n}$, where m and n are the total numbers of sampled condition positives and condition negatives, respectively^[223,224]. (Figure 4.5)

Each ROC curve $\mathcal{C}_i$ consisted of points $\mathbf{r}_i(t) = [TPR_i(t), FPR_i(t)]^T$, where t is the threshold for the classification metric and $TPR_i(t)$ and $FPR_i(t)$ are the true and false positive rates at that threshold for trial i . To simplify the calculation of the fixed-width band, a change of basis $\mathcal{T}(b) : \mathbf{r} \rightarrow \mathbf{u}$ was performed from the (TPR, FPR) basis ('ROC basis') to an orthogonal unit basis aligned with b (' u basis'):

$$\mathcal{T}(b)(\mathbf{r}(t)) = \mathbf{u}(t) = \begin{bmatrix} u_1(t) \\ u_2(t) \end{bmatrix} = \frac{1}{\sqrt{1+b^2}} \begin{bmatrix} -b & 1 \\ 1 & b \end{bmatrix} \mathbf{r}(t). \quad (4.9)$$

Given the geometric constraints inherent to an ROC curve – that it must traverse a path from (0,0) to (1,1) in the ROC basis and must always proceed 'north' and

‘east’ – and given that $[-b, 1]^T$ is, by construction, in the first quadrant (in the ROC basis), the ROC curve under $\mathcal{T}$ can be expressed by the functional relationship

$$\mathcal{C}'_i : u_2 = f_i(u_1), \quad \text{where} \quad u_1 \in \left[0, \frac{1-b}{\sqrt{1+b^2}}\right]. \quad (4.10)$$

The average ROC curve was defined as the median u_2 value over a resampled (500 steps) u_1 range:

$$\mathcal{C}'_{\text{avg}} : u_2 = f_{\text{avg}}(u_1) = \text{median}_i [f_i(u_1)], \quad u_1 \in \left[0, \frac{1-b}{\sqrt{1+b^2}}\right]. \quad (4.11)$$

The maximum distance d_i in the u_2 direction (perpendicular to b in the ROC basis) from individual curves $\mathcal{C}'_i$ to $\mathcal{C}'_{\text{avg}}$ is

$$d_i = \max_{u_1} (|f_i(u_1) - f_{\text{avg}}(u_1)|). \quad (4.12)$$

The semi-width of the band, d , was chosen as the 68.2-percentile of all d_i . This percentile was chosen such that visually, the band is analogous to the standard deviation error bars used elsewhere in this work. After transformation back into the ROC space ($\mathcal{T}^{-1}(b)$), the median curve and band were rendered (Figure 4.5). The area under the curve (AUC) for the median curve was calculated and taken as the primary metric for the classifier. (Note that Figure 4.5 depicts 20 curves rather than 100 for illustrative clarity.)

4.3.3 Single-Parameter Regression

Most prior work on image-based LA fibrosis quantification has relied on the MR voxel brightness distribution (throughout the myocardium or within myocardial regions) either to estimate bipolar voltage by linear regression or to estimate scar burden by threshold (Section 4.1.2). To establish a baseline for comparison against that work and to set a benchmark for the subsequent machine learning approaches, a simple linear model was employed to estimate bipolar voltage (‘linear models’) and log bipolar voltage (‘exponential models’) separately for each of the 38 normalised image parameters in (4.8).

EAM points and their associated image parameters were loaded into a Python application. Following Section 4.3.2, the points were divided into training and




Figure 4.5: Confidence scheme for ROC curve. Twenty individual curves (purple) generate the median curve (blue) and fixed-width band (cross-sections, yellow).

testing sets, and for each of the 38 image parameters (individually, $\hat{v}_p$), a single-parameter linear model,

$$BV_p = a_1 \cdot \hat{v}_p + a_0, \quad (4.13)$$

for the bipolar voltage BV_p was fitted by least squares.

Due to the approximately log-normal distribution of the BV_p in the sampled points compared to the approximately normal distribution of the image parameters (Section 4.4.1), an exponential model,

$$\ln(BV_p) = a_1 \cdot \hat{v}_p + a_0, \quad (4.14)$$

was also fitted. Linear least-squares was employed, with the log-transformed BV_p as the dependent variable.

For each model, the coefficient of determination (R^2), p value, and area under the median ROC curve (AUC) were calculated and reported.

4.3.4 Random Forests

Random forest is a well-validated machine learning technique which performs classification and regression using noisy feature sets, without prior knowledge of the relationships at play (Section 4.1.3). Random forest was chosen to evaluate non-linear and complex relationships between the image parameters.

Both classifier and regressor approaches to random forest were evaluated, as each would be expected to excel under different biological conditions: A classification approach would be expected to be more effective if there is a discrete structural (and therefore imaging) cutoff between ‘fibrotic’ and ‘normal’ tissue. A regression approach would be more accurate if the tissue was on a biological spectrum between ‘more fibrotic’ and ‘less fibrotic’, which correlated in a continuous way with image data and bipolar values. Both approaches were pursued and evaluated.

Classifiers

For each image feature set (i.e. normalisation method), preliminary tests were undertaken to determine the number of trees needed to achieve maximum classifier accuracy and the bipolar voltage threshold most strongly supported by image data. EAM points and their associated image parameters were loaded into a custom application written in Python. Following Section 4.3.2, for each of 100 trials, points were divided into training (75%) and testing (25%) datasets. The points in a given training set were classified as fibrosis or normal based on the appropriate ground truth bipolar voltage threshold.

The `scikit-learn`^[229] random forest classifier class was used to train a classifier on the training set. Robust parameters were chosen in accordance with Breiman’s original implementation^[220]: trees were constructed until all leaves were pure, bootstrap samples were used, and the number of randomly selected features was the square root of the total number of features available.

Classifiers were evaluated using the testing set, and ROC curves were generated. The area under the median ROC curve was used as the primary metric for the classifier.

Regressors

For the most effective image feature set as determined above, EAM points and image parameters were loaded into a custom Python application. Following Section 4.3.2, for each of 100 trials, points were divided into training (75%) and testing (25%) datasets.

A log-bipolar voltage regressor was trained using the `scikit-learn` random forest regressor class. Again, trees were constructed until leaves were pure, bootstrap samples were used. Unlike the procedure for classifiers, Breiman^[220] calls for all features to be considered in each decision step when training regressors, and we followed that procedure.

Regressors were evaluated using the testing set, and ROC curves were generated using threshold classification on the predicted bipolar value. The area under the median ROC curve was used as the primary metric for the classifier, and the coefficient of determination R^2 was taken as a secondary metric.

4.3.5 Principal Component Analysis

A high degree of intra-feature correlation is expected in the image feature data. For example, neighbourhoods of similar radii about the same point will yield similar measures across all features, mean voxel values will be highly correlated with minimum and maximum values, and local voxel variation will correlate with higher maximum and lower minimum values even under liberal assumptions. Although there are formally 38 image features for each point, the number of meaningful degrees of freedom represented is likely far lower.

Principal component analysis was employed to reduce the dimensionality of the image feature vectors. Image parameters were read into a Python database. In order to harmonise the variances between local intensity measures ($\hat{v}_{p,d}^{avg}$, $\hat{v}_{p,d}^{min}$, and $\hat{v}_{p,d}^{max}$) and the local variation measure ($\hat{v}_{p,d}^\sigma$), the latter was normalised by the constant $\frac{1}{|p|} \sum_p \left(\hat{v}_{p,5mm}^{avg} / \hat{v}_{p,5mm}^\sigma \right)$. The `scikit-learn` PCA decomposition class was used to transform the image feature vectors into their principal component space. All components were retained, and further whitening was not performed.

Components were analysed to determine the proportion of relative variation explained by each component. The transformation from raw image features to PCA components was studied to identify the gestalt of the more fundamental components (e.g. ‘global brightness’, ‘global variation’, ‘near-field brightness’).

The principal components were evaluated using random forest analysis in a two-step methodology. First, random forest classifiers and regressors were trained using each of 38 restricted PCA feature sets containing the n primary principal components ($n \in [1, 38]$). For example, in the $n = 5$ case, the first, second, third, fourth, and fifth principal components are included in the feature vector.

In the second step, the principal components were ordered by their importance in the $n = 38$ case from the first step and the process is repeated. For example, in the $n = 5$ case, the first, fifth, sixth, thirteenth, and fourteenth components are included in the feature vector, if these components were the five most important when all 38 are considered.

For each step, each random forest is independently trained and tested 100 times, and the AUC for the median ROC is calculated. This allows an evaluation of which principal components contribute information to the classification task, and how many useful degrees of freedom are represented by the image features.

4.4 Results – MR Cohort

4.4.1 Electroanatomic Map and Image Parameters

Bipolar voltage data for a total of 1619 EAM points from 16 patients (mean 110 points per patient) were recorded. Voltage values exhibited an approximate log-normal distribution with $\mu = -0.58$ and $\sigma = 1.10 \log \text{ mV}$ (Figure 4.7). Because these points were sampled during clinical procedures, the operators preferentially sampled from regions of clinical interest within the LA. Thus, the sampling was not statistically representative of the whole LA myocardium, and its distribution does not reveal fundamental truth regarding the distribution of the myocardial tissue as a whole.




Figure 4.6: Representative MR images used in left atrium analysis. Coloured contours represent manual left atrial myocardial segmentation.

Image parameter data for each EAM point was calculated. Representative images, including manual LA segmentation, are shown in Figure 4.6. Mean, maximum, and minimum local voxel values were approximately normally distributed with slight negative skew. Voxel variation $\hat{v}_{p,d}^\sigma$ had positive skew. (Figure 4.8 shows the data for $d = 5$ mm.)

4.4.2 Single-Parameter Regression

Linear Models

The mean, minimum, and maximum voxel values under IIR normalisation were significantly associated with bipolar voltage ($p < 0.05$). All mean and minimum values, in addition to maximum values over neighbourhoods 5 mm or smaller




Figure 4.7: Bipolar voltage distribution from MR cohort (top). Log of the bipolar values (bottom) demonstrates log-normal distribution.

(together, ‘most significant parameters’), produced $p < 0.001$. However, R^2 was uniformly low, with range 0.7 - 1.3% for the most significant parameters.

The IIR normalisation method provided improved p values and determination coefficients (R^2) compared to the Utah method. Prediction values (AUC) were also uniformly higher over significantly correlated variables. As expected given the low explained variance, the prediction value of these linear models was modest, with range 0.577 - 0.618 over the most significant parameters.

Under no formulation of the linear model was the local variation of voxel values significantly correlated to bipolar voltage value ($p > 0.2$).

The full results of the single-parameter linear model analysis are listed in Table D.1, page 190.

Exponential Models

Though broadly similar to the linear model – significant p values with low coefficients of determination and modest AUC values – the exponential models showed modestly improved R^2 values (2.5 - 3.3%) and much improved statistical significance ($p < 10^{-9}$)




Figure 4.8: Distribution of image parameters at $d = 5$ mm for MR cohort, normalised via IIR method (left) and Utah method (right). Red curves represent normal probability density function with distribution μ and σ .

compared to the linear models over the most significant parameters. AUC values were similar between linear and exponential models (maximum difference $< 1.5\%$), with maximum AUC 0.621.

The exponential model also reached significance in local voxel variation over small neighbourhoods ($p < 0.01$ for $d \leq 5$ mm), though R^2 was about 0.5%. For voxel variation, the Utah method performed slightly better than the IIR method of normalisation.

The highest single-parameter classification accuracy was achieved using the exponential model for the 8 mm neighbourhood maximum voxel value. This model achieved an area under the median ROC curve of 0.621, which is low in absolute terms.

The results of the single-parameter exponential model analysis are listed in Table D.2, page 191.

4.4.3 Random Forest

Classifier

Based on the results of the single-parameter exponential models, three normalisation strategies were identified to generate random forest feature sets: the IIR and Utah methods, as well as a hybrid method in which IIR normalisation was used for voxel intensity parameters ($\hat{v}_{p,d}^{avg}$, $\hat{v}_{p,d}^{min}$, and $\hat{v}_{p,d}^{max}$), while Utah normalisation was employed for variation ($\hat{v}_{p,d}^{\sigma}$).

Preliminary tests were carried out to empirically determine the number of trees necessary to achieve maximum classification accuracy. For each of the normalisation methods and using the most extreme published bipolar voltage threshold values – the lower Penn criterion ($BV_p < 0.27$ mV) and the Melbourne criterion ($BV_p < 0.5$ mV) – a random forest classifier 100 times and evaluated for accuracy. This process was repeated with a range of trees per random forest between 1 tree and 250 trees. In all cases, there was no advantage to using more than 200 trees (Figure D.2, page 194).

An additional preliminary experiment determined the normalisation strategy and bipolar voltage threshold. Threshold values between 0.15 mV and 0.75 mV were sampled every 0.025 mV. For each threshold value and normalisation strategy, 100 random forest classifiers were trained and evaluated for area under the ROC curve (Figure D.1, page 189). The hybrid normalisation feature set showed the strongest predictive values, with ROC AUC at least 0.73 for threshold voltages between 0.35 mV and 0.50 mV. The highest classification accuracy was for $BV_p = 0.475$ mV. Because it has the strongest image support and fits well with published values, this threshold was used in subsequent analysis.

The median ROC curve and relative feature importances for the random forest classifier with hybrid image feature set and 0.475 mV bipolar voltage threshold are shown in Figure 4.9. The area under the median ROC curve is 0.746, significantly higher than the 0.621 achieved by the best single-parameter model. Feature importances demonstrate that both near-field and far-field intensity values (mean, minimum, and maximum intensities) are important, as are variance values throughout the feature diameter range. Because of the high correlation




Figure 4.9: Left: Median ROC curve for random forest classifier of MR features with hybrid-normalised image feature set and 0.475 mV bipolar voltage threshold. Right: Relative importance measures for the same ($\pm$ standard deviation).

between features, detailed investigation of feature importances is deferred until the PCA analysis (Section 4.3.5).

Regressor

Preliminary tests (Figure D.3, page 194) showed that as for classifiers, random forest regressors performed optimally with 200 trees, using the hybrid-normalised feature set, and with a 0.475 mV bipolar voltage threshold.

The median ROC curve and feature importances for the random forest regressor are shown in Figure 4.10. The area under the median ROC curve is 0.742, significantly higher than the 0.621 achieved by the best single-parameter model and similar to the 0.746 achieved by the classifier. The coefficient of determination (R^2) for the regressor is 0.190 ± 0.034 , significantly higher than the 0.034 achieved by the most correlated single-parameter exponential regression model. The feature importances are substantially similar to those of the classifier, with emphasis on




Figure 4.10: Left: Median ROC curve for random forest regressor of MR features with hybrid-normalised image feature set and 0.475 mV bipolar voltage threshold. Right: Relative importance measures for the same ($\pm$ standard deviation).

near- and far-field intensity features as well as variance throughout. As before, detailed investigation of importance is deferred for PCA analysis.

4.4.4 Principal Component Analysis

Principal component analysis was run on the full image feature set. The loading vectors for the first 30 principal components are depicted in Figure 4.11. The first component is a global brightness measure, with equal weights for minimum, maximum, and average voxel intensity across the spectrum of feature diameters. The second component is a measure of variation, responding to points with low voxel variation, low maximum, and high minimum intensities. Note that the sign of these components is inconsequential; the second principal component measures high or low global variation.

Further principal components capture differences between local neighbourhood sizes (i.e. higher-frequency data). The third component recognises a difference in variation between neighbourhoods below and above the 7 mm diameter size.




Figure 4.11: Principal component feature makeups for the first 30 principal components based on MR feature vectors. Components are mean voxel intensity (green), maximum local intensity (purple), minimum local intensity (blue), and local voxel variation (red).




Figure 4.12: Cumulative variance explained by MR feature principal components. Above: Additional detail.

The fourth component captures significant brightness changes at the 5 mm averaging window (e.g. a bright point in a darker field). Higher components capture correspondingly more complex observations, with higher spatial frequencies.

As expected, a large amount of the total variation is captured by the first few principal components. The first (global intensity) component accounts for 70.2% of variance, while the second (global variation) accounts for a further 19.9%. The first nine components together account for 99% of the variance, and the first 22 components account for 99.9%. The cumulative variation explained by each component is shown in Figure 4.12.

A random forest classifier trained with the full PCA component set has an area under the median ROC curve of 0.737, similar to the 0.746 of the full hybrid feature set. The median ROC curve and component importances are shown in Figure 4.13.

Because of the random feature subspace splitting integral to random forest (Section 4.1.3), performance of the weak learners is degraded when, as after PCA




Figure 4.13: Left: Median ROC curve for random forest classifier trained on full MR PCA data set. Right: Relative importance of the same ($\pm$ standard deviation).

transformation, relevant information is concentrated into a few features. This effect was compensated for by training 2000 trees (rather than 200) per classifier. Additionally, this effect is only expected when irrelevant features are included.

Components 1, 14, 6, and 13 are significantly more important for the classifier than the other components. As seen in Figure 4.11, component 1 is global brightness. Component 14 corresponds with a particularly extreme near-field (2 mm) maximum local voxel value. Components 6 and 13 identify higher-frequency voxel variation modes, potentially identifying heterogeneous local tissue structure.

Random forest classifiers were trained with ordered subsets of the PCA components, and their ROC AUCs are shown in Figure 4.14. Adding components in order of importance outperforms adding in order of feature variance explained. Good tradeoff between components (degrees of freedom) and accuracy is achieved with nine importance-ordered components (1st, 14th, 6th, 13th, 5th, 29th, 2nd, 21st, and 10th), which achieves AUC of 0.746. Maximum accuracy is achieved with seventeen importance-ordered components (adding the 3rd, 36th, 27th, 23rd, 11th, 12th, 7th, and 8th), which gives an AUC of 0.752 (Figure 4.15).

Random forest regressors (for log bipolar voltage) were trained with ordered subsets of the PCA components, and their ROC AUCs and coefficients of deter-




Figure 4.14: Area under the median ROC curve for random forest classifiers built with subsets of MR PCA components, compared to the AUC for all components (dashed line). Top: Components added in order of feature variance explained. Bottom: Components added in order of classifier importance.




Figure 4.15: Left: Median ROC curve for random forest classifier trained on first seventeen importance-ordered MR PCA components. Right: Relative importance measures for the same ($\pm$ standard deviation).




Figure 4.16: Area under the median ROC curve (blue) and coefficient of determination (R^2 , red) for random forest regressor built with subsets of MR PCA components, compared to the AUC for all components (dashed lines). Top: Components added in order of variance explained. Bottom: Added in order of regressor importance.

mination are shown in Figure 4.16. As in the classifier case, adding features in order of importance outperforms adding in order of feature variance explained. Good tradeoff between components included and accuracy is achieved with ten importance-ordered components (6th, 1st, 14th, 5th, 13th, 3rd, 32nd, 2nd, 29th, and 26th), which achieves AUC 0.728 and coefficient of determination (R^2) 0.164 ± 0.035 . Maximum performance is achieved with nineteen importance-ordered components (adding the 26th, 23rd, 31st, 21st, 27th, 11th, 36th, 4th, and 10th), which gives an AUC of 0.740 and R^2 of 0.175 ± 0.035 .

4.5 Results – CT Cohort

4.5.1 Electroanatomic Map and Image Parameters

Bipolar voltage data for 4320 EAM points was recorded in 18 patients (mean 220 points per patient) with $\mu = 0.64$ and $\sigma = 0.94$ log mV (Figure 4.18). Unlike the bipolar data in the MR cohort, these data do not form a clear log-normal distribution.




Figure 4.17: Representative CT images used in left atrium analysis. Coloured contours represent manual left atrial myocardial segmentation.

This is likely due to denser sampling in these patients. Voxel mean, maximum, minimum, and variation of voxel values were approximately normally distributed with the exception the IIR-normalised neighbourhood maximum intensity, which appears bimodal (Figure 4.8). Representative CT images, including manual LA myocardial segmentations, are shown in Figure 4.17.

4.5.2 Single-Parameter Regression

Linear Models

Similar to the findings in the MR cohort, almost all image features were significantly associated with bipolar voltage ($p < 0.05$) in single-parameter linear models. The exceptions were the maximum value over neighbourhoods with diameter 3 mm




Figure 4.18: Bipolar voltage distribution from CT cohort (top). Log of the bipolar values (bottom) demonstrates log-normal distribution.

to 5 mm with IIR normalisation and with diameters 3 mm and 4 mm with Utah normalisation. However, R^2 and AUC were uniformly low, with values 1.7% and 0.604, respectively, for the most significant parameters.

The full results of the single-parameter linear model analysis are listed in Table D.3, page 192.

Exponential Models

Again, the exponential models are similar to the linear models, but with marginally better correlation. All models reach statistical significance ($p < 0.05$) except the maximum value at 4 and 5 mm IIR-normalised and 4 mm Utah-normalised. While higher than for the linear models, the R^2 and AUC remain low, reaching values of 4.3% and 0.609, respectively.

The results of the single-parameter exponential model analysis are listed in Table D.4, page 193.




Figure 4.19: Distribution of image parameters at $d = 5$ mm for CT cohort, normalised via IIR method (left) and Utah method (right). Red curves represent normal probability density function with distribution μ and σ .

4.5.3 Random Forest

Classifier

A series of 100 random forests, each made up of 200 trees, were constructed to classify points as normal ($BV_p > 0.475$ mV) or fibrotic based on the CT image features (IIR-normalised $\hat{v}_{p,d}^{avg}$, $\hat{v}_{p,d}^{min}$, and $\hat{v}_{p,d}^{max}$; Utah-normalised $\hat{v}_{p,d}^{\sigma}$).

The median ROC curve and relative feature importances for the random forest classifier is shown in Figure 4.20. The area under the median ROC curve is 0.977, an extremely high performance compared to the single-feature regression models. While some improvement over the MR data may be explained by the more broad distribution of bipolar voltage data (i.e. fewer ‘borderline’ points, Figure 4.18), this is a very significant improvement.

The feature importance data indicates a strong preference for maximum local voxel value over an 8 - 9 mm neighbourhood. However, given the low performance




Figure 4.20: Left: Median ROC curve for random forest classifier with hybrid-normalised image feature set and 0.475 mV bipolar voltage threshold. Right: Relative importance measures for the same ($\pm$ standard deviation).

of the same features as bases for single-parameter regression models, relationship is unlikely to be simply enhancement-as-fibrosis. Further consideration of feature contribution is reserved to principal component analysis.

Regressor

The median ROC curve and feature importances for the random forest regressor are shown in Figure 4.21. The area under the mean ROC curve is 0.951, again greatly surpassing the performance of the regression models and slightly below that of the classifier. The coefficient of determination (R^2) for the regressor is 0.770 ± 0.018 , an extremely high value in this context. The feature importances are substantially similar to those of the classifier, with emphasis on the 8 - 9 mm neighbourhood maximum voxel intensity value.




Figure 4.21: Left: Median ROC curve for random forest regressor of CT features with hybrid-normalised image feature set and 0.475 mV bipolar voltage threshold. Right: Relative importance measures for the same ($\pm$ standard deviation).

4.5.4 Principal Component Analysis

The feature makeup of the most significant principal component definitions are somewhat more diverse for the CT image features compared to MR (loading vectors, Figure 4.22). The first principal component represents brightness opposed to variance (i.e. high intensity coupled with low variance or vice versa). The second component responds to far-field variance, while the third finds close-field extreme intensities. Higher-order components begin to represent more obscure, higher-spatial frequency modes.

Although there is somewhat more freedom in the CT features compared to MR, a large amount of the total variation is nonetheless captured by the first principal components. The first (uniform brightness) component accounts for 54.4% of variance, while the second (far-field variance) accounts for a further 23.4%. The first four components together account for 92% of variance, the first twelve account for 99%, and the first 26 components account for 99.9%. The cumulative variation




Figure 4.22: Principal component feature makeups for first 30 principal components based on CT feature vectors. Components are mean voxel intensity (green), maximum local intensity (purple), minimum local intensity (blue), and local voxel variation (red).




Figure 4.23: Cumulative variance explained by CT feature principal components. Above: Additional detail.

explained by each component is shown in Figure 4.23.

A random forest classifier trained with all PCA components set has an area under the median ROC curve of 0.963, similar to the 0.977 of the full hybrid feature set (Figure 4.24). Components 1, 2, and 4 are significantly more important for classification. As seen in Figure 4.22, these components respectively select for uniform enhancement, heterogeneity, and complex structure.

Random forest classifiers were trained with ordered subsets of the PCA components, and their ROC AUCs are shown in Figure 4.25. Good tradeoff between components included and classifier accuracy is achieved with four importance-ordered components (1st, 4th, 2nd, and 10th), which achieves AUC of 0.941 (Figure 4.26).

Random forest regressors (for log bipolar voltage) were trained with ordered subsets of the PCA components (Figure 4.27). Good tradeoff between components included and accuracy is achieved with four importance-ordered components (1st,




Figure 4.24: Left: Median ROC curve for random forest regressor trained on full CT PCA data set. Right: Relative importance of the same ($\pm$ standard deviation).




Figure 4.25: Area under the median ROC curve for random forest classifiers built with subsets of CT PCA components, compared to the AUC for all components (dashed line). Top: Components added in order of feature variance explained. Bottom: Components added in order of classifier importance.




Figure 4.26: Left: Median ROC curve for random forest classifier trained on first four importance-ordered CT PCA components. Right: Relative importance measures for the same ($\pm$ standard deviation).




Figure 4.27: Area under the median ROC curve (blue) and coefficient of determination (R^2 , red) for random forest regressor built with subsets of CT PCA components, compared to the AUC for all components (dashed lines). Top: Components added in order of variance explained. Bottom: Added in order of regressor importance.

4th, 2nd, and 3rd), which achieves AUC 0.918 and coefficient of determination (R^2) 0.661 ± 0.026 .

4.6 Conclusions

4.6.1 Single-Parameter Models

While local image intensity is very significantly correlated with bipolar voltage in both MR and CT images, that correlation only accounts for a tiny amount of the total variation in bipolar values, and is insufficient to predict voltage on the basis of image features.

In the MR dataset, the most significant correlation (by p value) with the highest explained variance (by R^2 value) and best predictive value (by area under a median ROC curve) is achieved by an exponential model on the IIR-normalised mean neighbourhood voxel intensity. There is not a dramatic difference in correlation between these parameters at varying neighbourhood sizes, but an 8-mm diameter neighbourhood provides a maximum (AUC 0.620, R^2 0.033).

Similar conclusions can be drawn from the CT dataset, where the strongest bipolar voltage correlation is with minimum neighbourhood voxel values in the 6 - 7 mm voxel range (R^2 0.043) and the best predictions are achieved using the minimum values over the 10 mm neighbourhood (AUC 0.609).

For both modalities, the broad uniformity of correlation at different neighbourhood sizes likely reflects that atrial fibrosis is a global rather than local disease, and the likelihood that an individual point will be identified as fibrotic by EAM will as much depend on its regional environment as its immediate one.

While there is strong correlation between individual image features and bipolar voltage did not translate into predictive accuracy, limiting its clinical utility.

4.6.2 Random Forest

The random forest machine learning algorithm yields a much more useful prediction of bipolar voltage using image features. The random forest model exists in two flavours: The classifier's success (and hence training) is determined solely by labelling

image points as fibrotic or normal, evaluated by the area under the ROC curve. Meanwhile, the regressor aims to predict bipolar voltage directly, a task best judged by the determinant of correlation R^2 . In order to compare these methods, an area under the ROC curve was calculated for the regressor as well, parameterised by varying the bipolar voltage threshold for fibrosis.

The classifier and the regressor each significantly outperformed the single-parameter regression models in both the MR and the CT data sets. Preliminary testing showed that there is no advantage to using more than 200 trees in constructing the forest; a fibrosis bipolar voltage threshold of 0.475 mV had the highest image support; and training on a hybrid feature set, with IIR normalisation for intensity features and the Utah method for variation, produced the best results.

On the MR data, the classifier achieved ROC AUC 0.746, while the regressor achieved R^2 of 0.190 ± 0.034 and ROC AUC 0.742. The similarity in performance between the classifier and regressor suggests that image features do vary continuously with bipolar voltage, rather than exhibiting a distinct image ‘fingerprint’ for fibrotic tissue. While there remains significant unexplained variation in the bipolar voltage data, the 19% of variation modelled by the random forest classifier is much improved over the 3.3% achieved by the best single-parameter model.

On the CT data, the classifier achieved ROC AUC 0.977, while the regressor achieved R^2 of 0.770 ± 0.018 and ROC AUC 0.951. These values are very high by image analysis standards, and suggest that the health of the atrial myocardium is well defined by CE-MDCT image data alone. A partial explanation is the higher tissue consistency in CT data compared to MR data (Figure 4.28). This level of performance demands further investigation, and should be evaluated on a larger data set.

4.6.3 Principal Components Analysis

As expected, most of the overall variance in image features was captured in just a few principal components, although the CT image features showed more effective degrees of freedom than the MR features (99% variance explained with




Figure 4.28: Example slices from LGE-MRI (left) and CE-MDCT (right) images.

12 principal components and 9 principal components, respectively). Higher-order components are increasingly complex, and collectively, high-order components represent specific patterns of heterogeneity in the image data. By building random forest classifiers and regressors using PCA components sorted in order of importance, fewer components are needed to achieve performance comparable to the whole. Especially for MR features, relatively minor components still contribute significantly to bipolar voltage prediction.

4.6.4 Future Directions

The above work has shown that a standard machine learning method, by capturing more nuanced components of image feature data, can infer bipolar voltage using image data far more accurately than simple regression models.

Kapa *et al.* determined that normal bipolar voltage varies with region of LA myocardium^[204]. Separate analysis of points from different areas of the LA myocardium might lend further imaging support to this conclusion. It is also possible that due to different bioelectrical characteristics of surrounding tissues, image features and local electrograms yield different and complimentary information when LA region is taken into account.

The reason for the significantly higher performance of CT data compared to MR features should also be investigated further. A triple-modality data set, in

which each patient is evaluated with an MRI, CT, and EAM, would allow direct comparison of these modalities. It would also clarify whether LGE-MRI and CE-MDCT are windows to the same tissue characteristics, or alternatively whether these procedures reveal independent information and could be used in combination to achieve even higher accuracy.

Although this work relates image features to bipolar voltage, it does not make the leap to clinical outcome in terms of AF recurrence after PVI, the major conclusion of atrial fibrosis work to date in the context of AF. Given the increased correlation between image features and bipolar voltage using the machine learning method presented here, it is reasonable to conjecture that an image-based assessment of clinical outcome would be correspondingly more accurate, and that is worth investigation.

Finally, one of the major advantages of image-based atrial fibrosis evaluation compared to EAM is that it is appropriate for patients and volunteers who are not undergoing intracardiac instrumentation. The improved accuracy in this work would enable characterisation of changing LA fibrosis with age, prospective studies of AF incidence compared to atrial fibrosis, and development of fibrosis in AF patients receiving medical treatment rather than PVI.

At a cardiac arrest, the first procedure is to take your own pulse.

— Samuel Shem^[230]

5

Comparing Post-ablation Lesions to Pre-ablation Fibrosis

Contents

5.1	Introduction	93
5.1.1	Pathological Fibrosis and AF Severity	94
5.1.2	Therapeutic Lesion and Treatment Efficacy	95
5.2	Experimental Data	98
5.2.1	Patients	98
5.2.2	Magnetic Resonance Imaging	98
5.2.3	Electroanatomic Mapping	99
5.2.4	Pulmonary Vein Isolation	99
5.3	Methods	99
5.3.1	Registration and Classification of EAM Points	99
5.3.2	Comparing Image Features	100
5.4	Results	101
5.4.1	Pre-/Post-Ablation Differences	101
5.4.2	Relationship of Fibrosis to Ablation	103
5.4.3	Image-Predicted Bipolar Voltage	106
5.5	Conclusions	107
5.5.1	Future Directions	108

Chapter 5 Abstract

Background The effect of atrial fibrosis and scar on atrial fibrillation is paradoxical. Progressive fibrosis of the myocardium, associated with decreased bipolar

voltage and increased gadolinium enhancement, is a marker of the AF progression and negatively correlated with treatment efficacy. However, the degree of scar laid down during ablation correlates positively with treatment success. In this chapter, a comparison between the MR image characteristics of pathologically fibrotic tissue and ablated tissue is undertaken.

Methods Fourteen patients undergoing repeat pulmonary vein isolation (PVI) were recruited. Each patient had undergone late gadolinium enhancement MR imaging (LGE-MRI) before the first PVI procedure ('pre-ablation'), and the electroanatomic map and ablation data were recorded at that procedure. Patients then underwent a second ('post-ablation') LGE-MRI. The procedural EAM and ablation data were registered to both pre- and post-ablation MR images. Image feature data and image-predicted bipolar voltage values were calculated for each ablation point and non-ablation map point. The population distributions of these parameters were compared between pre- and post-procedure images and between pathogenic and ablation scar.

Results and Conclusions Non-ablation map points had similar characteristics in pre- and post-ablation images, while ablation points exhibited higher intensity and more heterogeneity in post-ablation images. Additionally, ablation scar is both more strongly enhancing and more heterogeneous than pre-ablation, periantral fibrosis. Future investigation should include post-ablation EAM data and treatment responders.

5.1 Introduction

A key conceptual question in AF research is the paradoxical relationship between pathological fibrosis and therapeutic scar. While atrial fibrosis and scar develop progressively and correlate with disease severity in the natural course of atrial fibrillation (Section 5.1.1), the amount of scarring left by ablation therapy correlates with the efficacy of treatment (Section 5.1.2). Conventional wisdom holds that

while pathological scar creates heterogeneous areas of slow and tortuous conduction, dense therapeutic scarring effectively ablates or isolates arrhythmogenic foci.

In this chapter, we apply the image analysis tools developed in Chapter 4 to investigate the relationship between pathological and post-ablation fibrotic tissue. The clinical data used in this work was also presented at the 35th Heart Rhythm Society conference in May 2014^[16]. Unlike the more complete analysis described below, that preliminary work did not incorporate multiple image features, principal component analysis, or local comparison of bipolar voltage data.

5.1.1 Pathological Fibrosis and AF Severity

Atrial fibrillation appears both as a lone entity, typically presenting in patients between 40 and 60 years old^[130], and as a comorbidity in patients with underlying cardiac disease, typically congestive heart failure^[131]. Lone AF may begin with paroxysms before increasing in frequency and developing into a persistent arrhythmia, leading to the maxim ‘atrial fibrillation begets atrial fibrillation’^[133].

Investigating the mechanism of this progression in a caprine (goat) model, Wijffels and colleagues^[133] showed that by inducing isolated incidents of AF in otherwise healthy hearts, electrical remodelling – a decrease in cardiomyocyte refractory period (Section A.4) – could be induced, which would decrease the likelihood of spontaneous abortion of the arrhythmia. As the arrhythmia persisted over time, this electrical remodelling was supplanted by physical remodelling – collagenisation and fibrosis – that would lead to persistent disease. This fibrosis can be identified as a low bipolar voltage on electroanatomic map^[199] or via analysis of MR or CT image data (Chapter 4).

Verma *et al.*^[201] demonstrated that low left atrial voltage is a powerful, independent predictor of PVI treatment failure using EAM. The researchers found that patients with left atrial scarring had a 57% recurrence rate after PVI, compared to a 19% rate in patients without scar. Building on this work, Marrouche and colleagues^[211] analysed 3-D LGE-MRI of 81 patients presenting for PVI. The researchers determined the extent of image enhancement (fibrosis) via a threshold

method, classifying patients as exhibiting ‘minimal’, ‘moderate’, or ‘extensive’ enhancement. Increasing levels of enhancement correlated with higher rates of treatment failure – 14%, 45%, and 75%, respectively, of patients in the minimal, moderate, and extensive enhancement groups. These papers are summarised from a data analysis perspective in Sections 4.1.1 and 4.1.2.

Marrouche *et al.*^[191] went on to investigate the atrial fibrosis burden – as evaluated by LGE-MRI – of patients with lone atrial fibrillation as compared to those with comorbidities. Considering 333 patients presenting for LGE-MRI evaluation prior to PVI, the researchers made two classifications. First, patients were staged based on the amount of MR atrial enhancement, a proxy for fibrosis. Second, patients under 60 years old without comorbidities were classified as ‘lone AF’, while those of any age with comorbidities were classified as ‘non-lone AF’. The researchers found that fibrosis correlated strongly with treatment failure: 96% recurrence was observed in the highest-enhancement group compared to no recurrence in the lowest-enhancement group. They also found an association between complicated AF and higher scar burden.

The same group later elaborated on these findings^[228], this time finding 21% recurrence in the lowest-enhancement group and 71% recurrence in the highest-enhancement group. Additionally, this study histologically validated the LGE-MRI fibrosis observation in six surgical samples. Taken together with previous studies, these findings support the notion that fibrillation induces remodelling which leads to further cardiac comorbidities in time. Atrial fibrosis, as identified by tissue enhancement on LGE-MRI, is a good proxy for the extent of extant damage, which plays a major role in treatment efficacy.

5.1.2 Therapeutic Lesion and Treatment Efficacy

Catheter ablation is the standard non-pharmacologic approach to atrial fibrillation and is growing in popularity^[163]. It is especially appropriate for young, otherwise healthy patients, in which the procedure has a higher success rate, and for patients with significant cardiac comorbidity, who might not tolerate pharmacologic

treatment. Initially developed as a minimally-invasive version of the Cox maze procedure^[165], catheter ablation for AF aims to eliminate or electrically isolate arrhythmogenic foci using radiofrequency energy, cryoablation, or laser energy. (Section 2.4.3)

Early efforts by cardiologists to replicate the maze procedure lesions using catheter ablation proved difficult and fraught with complications. In 1998, Haïssaguerre and colleagues^[173] reported that 94% of AF foci occurred in the pulmonary veins, and that ablating those foci abolished the arrhythmia. The goal of catheter ablation quickly shifted from reproducing the maze lesions to circumferentially ablating the pulmonary vein ostia^[174]. Establishing an atrial substrate that is inhospitable to arrhythmia, as is achieved by the maze procedure, becomes less necessary when the myocardium can be isolated from arrhythmogenic foci, as in pulmonary vein isolation.

In addition to pulmonary vein triggers, recent evidence supports localised electrical rotors and focal impulses within the atrial myocardium which act to sustain AF^[185,186]. Some clinical centres are incorporating targeted ablation of these regions in addition to pulmonary vein isolation in their ablation strategies. Others continue to use arbitrary ‘roof lines’ to disrupt AF wavelets in addition to isolating the veins. For simplicity, standard-of-care AF ablation procedures which feature confirmed electrical isolation of the pulmonary veins are referred to here as ‘pulmonary vein isolation’ (PVI). In current clinical practice, a single PVI procedure achieves a sustained success in approximately 40% of cases, while the long-term success rate when clinically indicated repeat procedures are included approaches 80%^[187,188].

It has long been observed that in PVI procedures, higher success rates are achieved when more tissue is ablated^[70,231,232]. In 2008, the University of Utah group lead by Marrouche^[70] examined enhancement in LGE-MR images acquired three months after PVI ablation in 46 patients with atrial fibrillation. Enhancement was observed following the expected pattern in all images, but a significant difference in scar burden was identified between responders and non-responders to PVI treatment. While treatment responders had an average 19% post-ablation left

atrial scar burden, the figure for non-responders was only 12%. On the contrary, stratifying patients simply by the clinical persistence of their atrial fibrillation did not achieve significance.

In a pilot study, Peters and colleagues^[232] at the Beth Israel Deaconess Medical Center in Boston also investigated the relationship between post-PVI scar burden and recurrence of AF. Thirty-five patients were imaged by LGE-MRI in the second month after circumferential ablation of the pulmonary vein ostia, and scar was identified by manual thresholding. Scar burden was assessed by calculating the percent scar volume as well as by qualitatively evaluating the circumferential extent of enhancement (from ‘minimal’ to ‘circumferential and extensive’).

Peters found that the post-ablation scar burden was significantly correlated with success of ablation. Patients were also significantly less likely to suffer recurrence if they had higher post-ablation enhancement in the right inferior pulmonary vein. Though a small pilot study, this work highlights the dose response effect of ablation lesion area in PVI treatment efficacy. More recently, the Utah group investigated the role of LGE-MRI immediately post-ablation (IPA)^[231]. Heterogeneous injury was observed in IPA MR images, which then consolidated over the course of three months to gadolinium-enhancing fibrosis. This consolidation resulted in a decrease in the visible lesion area from a mean of 72% of the LA wall IPA to residual scarring involving only 34% of the wall three months later. Although IPA images were not clearly helpful in predicting patient outcomes, the extent of enhancement on three-month images was again observed to correlate with treatment success.

Harrison and colleagues at King’s College London^[233] were the first to histologically validate the relationship between LGE-MRI and ablation scar. In a porcine model, the researchers acquired delayed enhancement MR images and high-density voltage maps of the right atrium before performing a linear radiofrequency ablation from the superior vena cava to the inferior vena cava along the posterior right atrial wall. A post-ablation LGE-MRI was then obtained. Eight pigs were sacrificed immediately for histological examination, while the other eight recovered from anaesthesia. Eight weeks later, the remaining pigs underwent another LGE-MRI and

a detailed EAM before sacrifice and histological examination. The myocardial wall was manually segmented from MR images, and voxels were evaluated with reference to the mean and standard deviation of a reference region within the atrial blood pool. Thresholds of 2.3 standard deviations above the reference region mean in immediate post-ablation LGE-MRI and 3.3 standard deviations in 8-week images were found to identify image regions that correspond with transmural scar on histology.

In this chapter, we present our investigation of the relationship between the image features of pre-ablation fibrosis and those of post-ablation lesion in an effort to resolve their seemingly paradoxical clinical relationship.

5.2 Experimental Data

5.2.1 Patients

Patients undergoing first-time pulmonary vein isolation procedures for atrial fibrillation at the Johns Hopkins Electrophysiology and Arrhythmia Service (Baltimore, MD, USA) were offered enrolment in this study. Patients who required a second PVI procedure were included. The MRI acquired in preparation for the second PVI procedure was used as the post-first-procedure MRI. All patients consented to the investigation, and the ablation and imaging protocols were approved by the Johns Hopkins Institutional Review Board (IRB study identifiers NA-00019442, NA-00021383, and NA-00070577).

5.2.2 Magnetic Resonance Imaging

The MRI protocol for both the pre- and post-ablation scan was identical to the one described in Section 4.2.2, page 56. Briefly, a 1.5-Tesla scanner was used with a torso coil and spine matrix coil. LGE-MRI scans were acquired 15 minutes after injection of a total 0.2 mmol/kg intravenous gadopentetate dimeglumine. The scans were acquired via a 3-D inversion recovery prepared fast spoiled gradient recalled, respiratory navigated, ECG-gated, fat-suppressed sequence. The respiratory navigator trigger time was selected to correspond with atrial diastole.

The LA myocardium, including pulmonary vein antra, was manually segmented from the LGE-MRI images using Seg3D software. Epicardial fat and blood pool were carefully excluded from the myocardial region.

5.2.3 Electroanatomic Mapping

LA endocardial electroanatomic mapping was performed in conjunction with PVI ablation procedures. At the time of ablation, patients underwent double transseptal puncture using standard techniques. The CARTO 3 mapping system was used in conjunction with the Thermo-Cool irrigated ablation catheter to record LA endocardial voltage maps. Patients with persistent AF were referred for external cardioversion 3–4 weeks prior to the procedure so all maps were obtained in sinus rhythm. Bipolar electrograms were filtered at 30 to 400 Hz.

5.2.4 Pulmonary Vein Isolation

Pulmonary vein isolation was performed using the standard clinical practice of the Johns Hopkins Electrophysiology and Arrhythmia Service^[234]. A wide circumferential linear ablation strategy has been employed with an endpoint of pulmonary vein isolation as assessed with a multipolar circular mapping catheter. The Thermo-Cool catheter was used, with target energy delivered to the anterior and posterior left atrial wall of 30 – 40 and 20 – 30 W, respectively. All patients had an esophageal temperature probe in place during ablation, and lesion application was temporarily ceased if a rapid rise in esophageal temperature was observed.

5.3 Methods

5.3.1 Registration and Classification of EAM Points

The electroanatomic map from the PVI procedure, including ablation and mapping points, was registered separately to the pre-procedure image and the post-procedure image using the technique described in Chapter 3 (Figure 5.1). Procedural ablation points as marked on the EAM were identified and classified as such. EAM points that were recorded before the beginning of ablation, and that were at least 2 mm




Figure 5.1: Schematic of study design comparing pre- and post-ablation image features.

away from any ablation point, were classified as ‘non-ablation’ points. EAM points that were recorded before the beginning of ablation, but which fell within 2 mm of an ablation point, were discarded for the purposes of this analysis.

5.3.2 Comparing Image Features

For each point, local neighbourhoods were defined and image features were extracted and normalised as described in Chapter 4. The mean, minimum, and maximum voxel values were recorded for neighbourhoods between 1 mm and 10 mm and normalised

using the IIR method. Local voxel variation was recorded for neighbourhoods between 3 mm and 10 mm and normalised using the simplified Utah method. (Section 4.3.1)

The principal component definitions found in Section 4.4.4 were applied to re-orient the image feature vectors. The bipolar voltage random forest regressor model (Section 4.4.3) was used to compute an ‘image-estimated bipolar voltage’ for each registered EAM point. The distributions of raw image parameters, principal components, and image-estimated bipolar voltage values were compared between ablation and non-ablation points in pre- and post-procedure images. Using each EAM point as its own matched control, the distribution in feature parameter and image-estimated bipolar voltage change between pre-procedure to post-procedure images was assessed in ablation and non-ablation points.

5.4 Results

Fourteen patients were recruited. A total of 1303 map points (mean 93 points per patient) and 6378 ablation points (mean 456 per patient) were collected. Map points with registered image coordinates within 2 mm of an ablation point were excluded, leaving 728 non-ablation map points (52 per patient) for analysis.

5.4.1 Pre-/Post-Ablation Differences

Distributions of selected image features measured pre- and post-procedure at ablation points are shown in Figure 5.2. It is visually apparent that the IIR-normalised mean image intensity is significantly higher post-ablation compared to pre-ablation. This fits with the consensus in the literature that enhancement is observed in ablated regions^[70,231]. The Utah-normalised image variance is also higher post-ablation, and this effect is more pronounced at higher neighbourhood sizes.

These visual relationships are borne out statistically. Recall from Section 4.4.4 that the first principal component feature for the MR data was a global measure of image intensity (i.e. enhancement), while the second feature corresponded to (inverse) heterogeneity (Figure 4.11, page 74). These features together account for over 90%




Figure 5.2: Histogram of selected image feature values recorded at ablation points, pre- and post-procedure (blue and red, respectively). Top row: IIR-normalised mean image intensity at 3, 6, and 9-mm radius neighbourhoods. Second row: Utah-normalised image variance over 3, 6, and 9-mm radius neighbourhoods. Bottom half: principal component features.

of variation in the raw image features. As expected, the first principal component ('PCA 1') is manifestly shifted positively after ablation (0.87 standard deviations higher in post-procedure images, $p < 0.0001$ by t-test), while PCA 2 is slightly lower, i.e. higher heterogeneity on average (0.47 standard deviations, $p < 0.0001$).

The corresponding data for non-ablation map points are shown in Figure 5.3. As expected, there is far less difference between pre- and post-ablation feature values in these non-ablation points. The most obvious difference is in PCA 2, which is still shifted lower (higher heterogeneity), despite little difference in the individual variation values. The likely explanation is that one component of PCA 2 is far-field (up to 10 mm) maximum voxel intensity (Figure 4.11). Given the human left atrium has a reference diameter of 27 - 40 mm^[235], the probability of a given map point falling within 10 mm of an ablation point is high, and excluding all such points would be prohibitive.

Given the single EAM that is registered to both pre- and post-ablation images for each patient, image feature sets were inherently matched between these images. Figure 5.4 depicts the matched-point difference in image features from pre- to post-ablation. In accordance with the observations above, the most significant differences are in the mean intensity values (including PCA 1) at the ablation points, which increase by 0.8 - 0.9 standard deviations.

5.4.2 Relationship of Fibrosis to Ablation

The fundamental question posed in this chapter is whether there is a structural difference between pathologically occurring fibrosis and therapeutic scar. To address this question, we evaluated the first two PCA components for three groups of points: pre-procedure map points with EAM bipolar voltage less than 0.475 mV ('pre-ablation normal'), pre-procedure map points with bipolar voltages greater than 0.475 mV ('pre-ablation scar'), and post-procedure ablation points ('ablation'). To reduce confounding by myocardial region, and in contrast to the procedure above, only pre-procedure map points within 2 mm of an ablation point were considered.




Figure 5.3: Histogram of selected image feature values recorded at non-ablation map points, pre- and post-procedure (blue and red, respectively). Top row: IIR-normalised mean image intensity at 3, 6, and 9-mm radius neighbourhoods. Second row: Utah-normalised image variance over 3, 6, and 9-mm radius neighbourhoods. Bottom half: principal component features.




Figure 5.4: Box plot of matched point image feature changes (z-score) between pre- and post-ablation images. Left: Non-ablation map points. Right: Ablation points.




Figure 5.5: Comparison of pre-ablation normal tissue, pre-ablation scar, and ablated tissue. Top: PCA 1, neighbourhood intensity. Bottom: PCA 2 (reversed), neighbourhood voxel variation.

The results are shown in Figure 5.5. The ablation points are both more enhanced and more heterogeneous than either pathological scar or pre-ablation normal tissue. For both features, there is a steady progression from normal tissue to scar to ablated tissue, suggesting the ablated tissue is ‘more scarred’ than pathological scar. For both features, the difference between pathological scar and ablated tissue is small (difference 0.25 standard deviations) but statistically significant (PCA 1: $p = 0.006$; PCA 2: $p = 0.001$ by t-test).

5.4.3 Image-Predicted Bipolar Voltage

Finally, pre-ablation normal tissue, pre-ablation fibrosis, and ablated tissue were compared by the bipolar voltage values predicted by the random forest regression model constructed in Section 4.4.3. These values are presented in Figure 5.6.




Figure 5.6: Image-predicted bipolar voltage of pre-ablation normal tissue, pre-ablation scar, and ablated tissue.

Contrary to the relationship of the direct image parameters, predicted bipolar voltage value for ablated points (0.57 ± 0.21 mV), falls between the values for normal tissue (1.08 ± 0.61 mV) and pre-ablation scar (0.32 ± 0.14 mV).

5.5 Conclusions

Three classes of comparison were made in this chapter. First, non-ablation map points and ablation points were compared between pre-procedure and post-procedure images. As expected, non-ablation points changed very little between pre- and post-procedure images. Ablation points, on the other hand, showed both more enhancement and more heterogeneity in post-procedure images.

Second, the image characteristics of post-procedure ablated tissue were compared to corresponding pre-ablation normal tissue as well as pathogenic fibrosis in the pre-ablation image. Both the mean voxel intensity and the voxel variation within the neighbourhoods of ablated points were significantly higher than the corresponding features of pathological scar in pre-procedure images. This lends credence to the explanation that ablated tissue is ‘more scarred’ than the pathological fibrosis of AF pathology, and that while the latter allows tortuous, arrhythmogenic conduction, the former blocks conduction, and thus arrhythmia. While this result is significant ($p < 0.01$), it is relatively small (0.25 standard deviations). Indeed, it is likely

that location plays at least as important a role in PCI's curative value as does character of the scar itself.

Finally, applying the random forest regressor model from Chapter 4 resulted in the unexpected result that the predicted bipolar voltage for ablated tissue is higher than that of pre-procedure scar. Understanding this result requires taking into consideration the training of that model, and its importance values (Figure 4.10, page 73). The fibrosis of AF is a global disease: while there are localised patches of scar, the character is far different from the hand-placed lines of scar produced by a radiofrequency catheter. For that reason, to expect the random forest model to produce meaningful bipolar values for post-ablation images is to commit the cardinal sin of machine learning: applying the model to data fundamentally different to the training data. If a prediction model is to be applied to images of previously-ablated hearts, it must also be trained on them.

5.5.1 Future Directions

Chapter 4 established that simple comparisons of distribution and trend in individual image features, even after reorientation with PCA, provides limited utility compared to more robust machine learning methods. While the work in this chapter has established the basic relationship between image characteristics of pre-procedure scar and ablated tissue, a prediction model might prove more clinically useful. By incorporating the post-ablation electroanatomic map into the training step, known ablation tissue bipolar values would be included, facilitating the construction of a more comprehensive model.

Given the performance of CE-MDCT imaging in Chapter 4, it would be extremely interesting to undertake a study of post-procedure ablation lesions using CT data. Unfortunately, repeated CT scans introduce undesirable radiation burden. Careful study design could minimise this risk.

Finally, the patient population evaluated in this analysis was limited by the clinical reality that each 'post-ablation' MR image was in fact the pre-procedure image for the patient's clinically indicated second ablation. Thus, only patients

with AF recurrence were included. While this does not detract from the core aim of the work – evaluating the character of scarring formed during ablation – it does limit comparative analysis between treatment responders and non-responders or predicting success based on analysis of post-ablation images. Further work on this topic should include provision for imaging treatment responders in order to expand the breadth of conclusions.

Ideally, EAM information would also be available in patients not requiring further ablation. However, the risk of an invasive mapping procedure likely precludes its use without clinical indication. Alternatively, an animal model could supply this data.

It is one object of this invention to provide a method and means for the recognition of complex patterns in a picture.

— Paul V. C. Hough^[236]

6

Annular Hough Transform for LV in 3-D Echocardiogram

Contents

6.1	Introduction	111
6.1.1	Survey of Segmentation Approaches	112
6.1.2	The Hough Transform	117
6.1.3	The Hough Transform in Cardiac Segmentation	121
6.2	Experimental Data	122
6.3	Methods	123
6.3.1	Gradient-Based Circular Hough Transform	123
6.3.2	Feature Asymmetry Edge Detector	124
6.3.3	Annular Accumulator	126
6.3.4	Radial Constraint	127
6.3.5	Comparison Criteria	128
6.4	Results	129
6.5	Conclusions	133
6.5.1	Future Directions	134

Chapter 6 Abstract

Background Segmentation of left ventricular myocardium from 3-D echocardiograms is desirable both as a first step toward computation of clinical parameters as well as for facilitating geometric registration. While ultrasound presents unique

challenges to level sets and active contours, the voting mechanism of the Hough transform is uniquely suited to missing and spurious edge data.

Methods A total of 5641 short-axis slices were extracted from 25 3-D cine (3D+t) echocardiograms. Each slice was segmented first using a simple gradient-based endocardial circular Hough transform and evaluated by reference to expert manual segmentation. The gradient-based Hough segmentation was compared to enhanced algorithms which progressively incorporated feature asymmetry as an edge detector, use of epicardial edge information in a two-step voting process, and a radial search constraint. Manual and automatic segmentations were compared based on their centre points, radii, and Dice coefficient.

Results and Conclusions The fully enhanced Hough algorithm was significantly more accurate than the basic single-surface, gradient-based method (centre point discrepancy 2.5 mm enhanced, 8.6 mm basic; Dice coefficient 0.77 enhanced, 0.58 basic). Inaccuracies were largely restricted to apical short-axis slices, where limited image data is present. This method promises utility for estimation of clinical measurements and segmentation-assisted cross-modality registration in near-real-time.

6.1 Introduction

Segmentation of the left ventricular myocardium is the critical step in automated image-based evaluation of cardiac structure and function. Blood pool segmentation alone allows computation of classic volumetric parameters such as stroke volume and ejection fraction. Segmentation of the myocardial tissue will enable determination of other clinically important measurements, such as stress-strain curves, dyssynchrony, and wall motion abnormalities.

Segmentation also facilitates cross-modality geometric registration by abstracting away modality-dependent tissue responses. The benefits of this approach are especially distinct in echocardiography, which has particularly intractable features

for registration. The ready availability of ultrasound at the interventional bedside makes it an ideal registration intermediary between pre-op diagnostic scans and operative guidance.

Below, we survey popular approaches to left ventricle segmentation in the echocardiogram, with particular focus on the Hough transform.

6.1.1 Survey of Segmentation Approaches

Echocardiogram left ventricular segmentation is a mature problem in medical computer vision, with several avenues of open investigation. The most common approaches include active contours and level set methods for boundary estimation. Alternatively, segmentation of the blood pool or myocardium may be considered a classification problem on a region or pixel level. Other approaches involve fitting active shape models and incorporating motion data.

Active Contour Models

Active contour models, also known as ‘snakes’, are iterative approaches which attempt to fit a curve to image features using an energy minimisation strategy and, optionally, manual input^[237]. The curve, classically parameterised as a series of control points, is initialised manually or automatically. From there, the parameters are iteratively adjusted to minimise the total energy function

$$E_{\text{snake}}^* = \int E_{\text{snake}} ds = \int (E_{\text{internal}} + E_{\text{image}} + E_{\text{constraint}}) ds, \quad (6.1)$$

where E_{snake} , the local energy parameterised by arc length s , is the sum of an internal energy term E_{internal} , an image-based energy term E_{image} , and optionally external constraint energy $E_{\text{constraint}}$. The internal energy term promotes regularity and smoothness of the curve, while the image term provides attraction of the curve to salient image features. Trivial examples of E_{image} are the brightness of the image itself or an image gradient function, which would facilitate convergence to bright lines or boundaries in the image, respectively. Finally, the external constraint term encodes manual guidance or external factors such as the image boundary or prior knowledge.

Convergence of active contour models depends on initialisation near the image feature as well as optimisation of the energy minimisation algorithm. Several approaches have been taken in regard to echocardiogram segmentation. Mishra *et al.*^[238] showed preliminary success in segmenting the LV blood pool using initialisation based on low-pass filtering of the 2-D image and using a genetic model for contour iteration. Mignotte and Meunier^[239] avoid local minima by constructing a multi-scale optimisation pyramid. Notably, they choose an image energy function based on region statistics using the Rayleigh distribution, rather than a more typical edge-based metric. De Alexandria and colleagues^[240] took the novel approach of constructing paired contour control points along the axial direction of an ultrasound image, using a 1-D Hilbert transform energy function. Restricting calculations to one dimension along the axial direction reduced the effect of noisy ultrasound images. Extension of active contours into 3-D imaging is relatively new. Barbosa *et al.*^[241] apply B-spline active surfaces to segment the left ventricular blood pool in 3-D echocardiograms. Their work relies heavily on automatic initialisation algorithm which detects the LV axis using a Hough transform technique and mitral valve plane using the average grey level of a normal plane along that axis. This initialisation allows the subsequent active contour approach, guided by a local statistical Yezzi energy function^[242,243], to converge on the endocardial boundary.

Level Set Methods

Level set methods are another approach to boundary fitting. The key difference compared with active contour is that the boundary is represented not as a parameterised curve, but as the zero-crossing of a higher-dimensional function^[244]. A two-dimensional curve, for example, would be represented by a 3-D function over the domain of image location (Figure 6.1). Formally, the closed curve Γ is represented by the zero level set of an auxiliary function $\varphi(x, y)$:

$$\Gamma = \{(x, y) | \varphi(x, y) = 0\} . \quad (6.2)$$




Figure 6.1: Illustration of the level set construction. The endocardial contour on the 2-D short-axis echocardiogram slice (grey) is approximated by the zero-crossing (magenta) of the 3-D function (coloured mesh).

Like active contours, level sets are an iterative technique. Given some initialisation, the curve Γ proceeds at velocity v in the normal direction when φ obeys the level set equation

$$\frac{\partial \varphi}{\partial t} = v |\nabla \varphi|. \quad (6.3)$$

Efficient numerical solutions to this equation rely on sophisticated optimisation techniques, as well as strategies to avoid storing and manipulating the value of φ for every voxel value – including the narrow band method^[245], the sparse field method^[246], the octree method^[247], and run-length encoding^[248]. The fundamental advantages of level sets over active contours is the avoidance of curve parameterisation. Level sets also allow greater topographic flexibility compared to active contours, as the curve can split into two (i.e. with multiple zero-crossings of φ).

Level sets have been used to segment the LV blood pool in echocardiography. Yan and Zhuang^[249] used an average regional image intensity gradient to drive the level set progression in order to mitigate the effect of local minima, improving stability. Lin *et al.*^[250] used a multi-scale pyramid approach to avoid local minima

in 2-D slices from 3-D ultrasound data, with success comparable to the inter-observer variation of manual segmentations. Advances in computational power and real-time 3-D ultrasound probes are also paving the way for 3-D level set implementations. Galluzzo and colleagues^[251] developed a graphics processing unit (GPU) algorithm that implements the sparse field level set method on inhomogenous 3-D echocardiogram data. Using parallel processing, they achieved segmentation times under one second with clinically meaningful volume measurements, but did require manual initialisation.

Initialisation close to the appropriate boundaries is critical for correct convergence of both active contours and level sets^[252]. Classification approaches and active shape models have been used with varying success to overcome this obstacle, though the resulting methods are constrained by the quality, quantity, and variety of training data available^[253].

Classification Approaches

Classification approaches to echocardiogram segmentation rely on local or regional image statistics to determine whether each image element is inside or outside the region of interest (e.g. the LV blood pool) rather than iteratively searching for boundary features. There are a number of algorithms to perform this classification. One popular approach, as pursued by Mignotte *et al.*^[254] and Boukerroui *et al.*^[255], is to model each pixel's classification as a Markov random field called the 'labelling field'. A global energy function is defined over all possible label maps (i.e. segmentations). This function describes the probability that each pixel – and each group of neighbouring pixels – falls into a given segmentation pattern, given their respective image features. This probabilistic construction allows a Bayesian approach, taking advantage of Bayes' theorem to incorporate knowledge about the expected image feature distribution of labelled regions into the model. This strategy is particularly useful in noisy ultrasound images, though its performance is dependent on sound *a priori* knowledge of regional statistical distributions as well as optimisation of the expectation maximisation problem.

The above approach is one of many ‘graph cuts’ methods, in which image pixels are likened to locally-linked nodes on a graph and segmentation is achieved by finding an optimal ‘cut’ through that graph. Another classification strategy is machine learning, in which a decision model is trained using regional image features to classify pixels into segmentation regions. Binder and colleagues^[256] demonstrated this technique by training an artificial neural network to classify pixels as ‘tissue’ or ‘blood pool’ given regional grey value, variance, contrast, entropy, and homogeneity. They found good performance – including estimation of clinical parameters – in 2-D echocardiograms of moderate to good image quality.

Active Shape Models

Just as Bayesian classification approaches add *a priori* knowledge of statistical distributions of segmented regions in order to improve accuracy, shape models incorporate knowledge of the shape to be segmented, presenting particular advantages when edge information is missing in the image.

One example of a shape model, the active shape model, introduces an additional constraint or energy term in an active contour or level set segmentation. Limiting the movement of a contour to a known range of shapes not only ensures the final segmentation will be compatible with prior knowledge, but also reduces the parameter space, easing the energy minimisation problem^[257]. Chen *et al.*^[258] used prior shape information to constrain the evolution of an active contour model. The group validated its method on long axis 2-D echocardiogram data, finding good agreement with manual segmentation by clinicians.

Alternative approaches incorporating shape information have been pursued. Stebbing^[259] found good results with a boundary fragment model, which considers a learned collection of edge voxel masks and selects those with good image support. Unlike active shape models, the boundary fragment model matches potential edges with image responses before any iteration. The result is a group of edge features that define the left ventricle in these images. Segmentation may be completed by fitting




Figure 6.2: Block diagram illustrating electronic plane transform circuits used in the original Hough transform, from the United States patent record^[236].

a B-spline surface to these edge features. The boundary fragment model performs quite well in segmenting the left^[259] and right^[260] ventricles in 3-D echocardiograms.

The above segmentation methods have shown varying degrees of success in echocardiography. However, especially for images produced by current-generation 3-D ultrasound probes, low spatial resolution, speckle, shadowing, and parallel-beam effects cause spurious edge responses throughout the blood pool as well as missing edges^[261], presenting difficulties for segmentation. Iterative methods require good initialisation in order to converge, and while optimisation methods are improving, there is still a role for easily parallelisable approaches. Segmentation of the epicardial boundary is further complicated by the presence of nearby tissues with similar echogenicity and proximity of the epicardial boundary to the image edge in typical ultrasound windows^[262,263].

6.1.2 The Hough Transform

The Hough transform is a generalisable feature detection algorithm^[264] that relies on a discrete parameter-space voting mechanism to identify image features. This

voting mechanism excels in detecting features despite missing edge data and spurious features. Developed in the late 1950s to detect lines in bubble chamber photographs used in particle physics^[265], the original implementation of the algorithm relied on analogue electronic plane transform circuits^[236] (Figure 6.2).

The original linear Hough transform takes advantage of the symmetry inherent in the equation of a line in $\mathbb{R}^2$. Given an image space $\mathbf{I} : (x, y) \in \mathbb{R}^2$, a line is defined by the equation $y = m \cdot x + b$, where m and b are the slope and intercept.

Therefore, any specific line in $\mathbf{I}$, $y = m_1 \cdot x + b_1$, has a one-to-one correspondence with the point (m_1, b_1) in the two-dimensional parameter space $\mathbf{H} : (m, b) \in \mathbb{R}^2$. By rearrangement, a given image point $(x_1, y_1) \in \mathbf{I}$ is analogously represented in $\mathbf{H}$ by the line

$$b = -x_1 \cdot m + y_1. \quad (6.4)$$

Geometrically, every point (m, b) that satisfies equation (6.4) represents a line in $\mathbf{I}$ that intersects (x_1, y_1) . (Figure 6.3)

To harness these reciprocal correspondences, a quantised ‘accumulator image’ is constructed over $\mathbf{H}$. The Hough algorithm scans the original image $\mathbf{I}$ and, for each edge pixel, generates a linear ‘vote’ in $\mathbf{H}$. These votes are summed, and the coordinates of the maximum accumulator value reveal the parameters (m, b) of the line in $\mathbf{I}$ with the strongest image support. (Figure 6.3)

This technique is applicable to detection of any feature defined by a finite number of parameters^[264]. Modern line detection approaches, unconstrained by analogue circuitry, represent lines in $\mathbf{I}$ as $y = -\cot \theta \cdot x + \rho \csc \theta$. This lends itself to a bounded accumulator space $\mathbf{H} : (\theta, \rho)$ wherein θ is on $[0, \pi)$ and ρ is limited by the size of the image. A point (x_1, y_1) in $\mathbf{I}$ transforms to the sinusoid $\rho = x_1 \cos \theta + y_1 \sin \theta$ in $\mathbf{H}$, rather than to a line.

Two-dimensional circle detection requires a 3-D accumulator, as the parameters include both 2-D centre position and radius (c_x, c_y, r) . The image point (x_1, y_1) is transformed in $\mathbf{H}$ onto the right circular cone

$$(c_x - x_1)^2 + (c_y - y_1)^2 = r^2. \quad (6.5)$$




Figure 6.3: Schematic of the Hough transform for line detection in a two-dimensional image. For simplicity, lines are parametrised by (m, b) , as used in the original implementation, rather than (ρ, θ) , which is more common in modern implementations.




Figure 6.4: Comparison of the accumulator vote (middle column) and sample supported shapes (right column) associated with a single image point for a circular Hough detector based on edge information only (top row) and gradient information (bottom row).

Analogously, ellipses require five parameters: centre, semimajor and semiminor axis lengths, and rotation angle. Arbitrary shapes may be detected by defining position, scale, and rotation parameters in a process that is mathematically equivalent to convolution^[264].

The original Hough transform took into account only binary edge information, using a Canny or similar edge detector as a preprocessing step on $\mathbf{I}$. The accuracy of closed shape detection is significantly increased by using gradient information as well^[264]. Figure 6.4 illustrates the mechanism of this improvement: the image gradient information allows directional pruning of each accumulator vote, reducing spurious intersections and local maxima. Because slight (even pixel-level) discrepancies in image edge position or orientation may result in accumulator votes that do not overlap – especially in higher-dimensional spaces – the accumulator image is blurred before the maximum point is identified.

6.1.3 The Hough Transform in Cardiac Segmentation

Variations on the Hough transform have been used to segment the LV blood pool from short-axis echocardiogram images. Bansod and Burkule^[266] developed a method to estimate the endocardial border of 2-D cine (2D+t) parasternal short-axis echocardiogram images, using a circular Hough transform to determine the left ventricular centre point (LVCP). After preprocessing with speckle-reducing anisotropic diffusion^[267], the researchers applied gradient-based circular Hough transform to each 2-D frame in a 2D+t image. The method gave preference to the largest circle detected, and rejected circles outside the acoustic window. The radius range was fixed according to the known acoustic window. The centre coordinates of each 2-D frame were averaged, excluding those which fell outside of a reasonable region defined based on the acoustic view.

After using the Hough transform to find the LVCP, Bansod and Burkule employed a radial gradient search method to estimate the endocardial boundary. Boundary points were identified by scanning radially from the LVCP for maximum image gradient. Significant restrictions on the search space were established by a combination of image statistics, the Hough transform radius, and temporally adjacent slices. This method achieved a mean boundary accuracy of ± 2 mm compared to manual segmentation.

Fernández-Caballero and colleagues^[268] proposed using an elliptical Hough transform to initialise an active contour approach to LV endocardial segmentation. The derivative-of-Gaussian operator (convolution with the gradient of a Gaussian kernel; as distinct from the difference-of-Gaussians, DoG) was used to generate edge maps from short-axis 2-D echocardiograms. A gradient-based, elliptical Hough transform was applied to each of these short-axis images, and an active contour was initialised from the ellipse positions.

Fernández-Caballero proposed using these short axis segmentations together with a similarly constructed long-axis segmentation to calculate clinical parameters. The long-axis segmentation used a generalised Hough transform – with *a priori*

defined long axis LV shape – to initialise an active contour. This work was a methodology proposal only, and was not validated on clinical data.

To our knowledge, a Hough segmentation approach making use of both the endocardial and epicardial surfaces for LV segmentation in echocardiogram images has not been investigated. Here, we introduce a myocardial segmentation technique which applies three enhancements to the endocardial, gradient-based circular Hough transform. First, the gradient edge detector is replaced by feature asymmetry, a tunable local phase-based method^[269]. Second, by approximating the short-axis LV myocardium with an annulus, a dual voting algorithm is developed which takes advantage of both endocardial and epicardial edge information. Finally, search constraints are applied to the accumulator space. These automated segmentation methods are compared to one another and to expert manual segmentation of 3-D cine (3D+t) echocardiogram data.

An earlier version of this work was published at the 9th IEEE International Symposium on Biomedical Imaging (ISBI)^[17] in May 2012. That version incorporated the annular modification to the circular Hough transform and included the same clinical echocardiogram data for validation. However, it did not incorporate the feature asymmetry edge detector or radial constraint described below.

6.2 Experimental Data

Twenty-five 3D+t apical view echocardiograms were obtained from the John Radcliffe Hospital in Oxford. The subjects were healthy volunteers, and each study was recorded on a Philips iE33 ultrasound system. Spatial resolution was $0.88 \times 0.88 \times 0.81$ mm³ and temporal resolution was on average 59 ms. Each 3D+t sequence consisted of between 11 and 20 3-D frames, with the first frame at end diastole. The first 11 frames of each sequence were used in the experiments, for a total of 275 3-D images. Subsequent frames were discarded in order to avoid biasing the results in favour of certain patients. Diastolic frames were chosen for discard because they are relatively similar and adynamic.

Manual segmentation was performed by three experts: for each of the 3-D images, one expert delineated the LV myocardium in every fifth short-axis slice of each volume using custom software written in Matlab. The expert recorded an epicardial and an endocardial contour, using the orthogonal long-axis view to guide each short-axis segmentation. Papillary muscles and valves were included in the blood pool. Where the myocardium was less visible due to shadowing, attenuation, or a parallel beam, expert knowledge of the shape of the heart wall, guided by orthogonal long-axis slices, was used to determine the boundaries. At the time of segmentation, the experts had no knowledge of the automated segmentation algorithm for comparison.

Given the extent of the LV in each of the 275 3-D images, a total of 5641 2-D short-axis slices were used in the experiments.

6.3 Methods

We introduce three modifications to the basic circular Hough transform. First, feature asymmetry^[269,270] is used in place of the gradient-based edge detector. This permits selection of edges in specified frequency ranges and results in more sparse edge maps than gradient methods. Second, an annular accumulator is constructed which uses both positive and negative values to incorporate endocardial and epicardial surfaces into the myocardial segmentation algorithm. Finally, a search space constraint is placed to limit the endocardial radius to enforce reasonable results in low-signal apical slices.

6.3.1 Gradient-Based Circular Hough Transform

A custom application was written in C++ to implement a standard gradient-based directional circular Hough transform to detect the LV endocardial boundary. For each short-axis slice $\mathbf{I}(x, y)$, an accumulator image $\mathbf{H}(c_x, c_y, r)$ was formed. The directional first-order central gradients $G_x(x, y)$ and $G_y(x, y)$ were computed, and the gradient magnitude image $G(x, y) = \sqrt{G_x^2 + G_y^2}$ was calculated. G was thresholded, with values below 33% of the maximum discarded.




Figure 6.5: Schematic of conical vote $V_{p_i}(c_x, c_y, r)$. The positive and negative votes each subtend an angle of $\pi/4$, centred on gradient direction θ_g . Grey indicates no vote, while white and black represent positive and negative votes respectively.

For each remaining point $p_i = (x_i, y_i)$, an accumulator vote V on $\mathbf{H}$, based on gradient direction $\theta_g(x, y) = \arctan(G_y/G_x)$, was constructed. The vote is depicted in Figure 6.5 and is defined parametrically as:

$$\begin{aligned} c_x &= x_i \mp r \sin \theta & c_y &= y_i \mp r \cos \theta \\ r &\in [0, r_{\max}] & \theta &\in \left[\theta_g - \frac{\pi}{8}, \theta_g + \frac{\pi}{8}\right] \\ V_{p_i}(c_x, c_y, r) &= \pm G(x_i, y_i). \end{aligned} \quad (6.6)$$

The votes were summed to form an accumulator image, which was smoothed with a Gaussian kernel. The parameters of the maximum value represent the centre and endocardial radius of the left ventricular blood pool,

$$c_x, c_y, r_{\text{endo}} = \underset{c_x, c_y, r}{\operatorname{argmax}} \left(\text{Gaussian} * \sum_{p_i} V_{p_i} \right). \quad (6.7)$$

This algorithm was applied to the 5641 short-axis slices, and the results were compared to manual segmentation as well as the modified algorithms detailed below.

6.3.2 Feature Asymmetry Edge Detector

The first modification of the Hough transform segmentation algorithm was implementation of feature asymmetry in place of the gradient edge detector. Feature asymmetry is a structure detection algorithm generalised from the one-dimensional

monogenic signal^[270]. Feature asymmetry demonstrates particular strengths identifying structures in cardiac ultrasound images^[269].

The Riesz transform^[270] in two dimensions is a generalisation of the one-dimensional Hilbert transform. Its operator components are defined in the frequency domain as

$$\mathcal{H}_x = \frac{i \omega_x}{\sqrt{\omega_x^2 + \omega_y^2}} \quad \text{and} \quad \mathcal{H}_y = \frac{i \omega_y}{\sqrt{\omega_x^2 + \omega_y^2}}, \quad (6.8)$$

where ω_x and ω_y are orthogonal frequency domain parameters corresponding to the x and y spatial directions in $\mathbf{I}$, and i is the imaginary unit. Each filter responds to edges along its image axis direction. Analogous to the Hilbert transform, the Riesz transform requires the selection of an isotropic low-pass filter. Here, we choose the log-Gabor filter defined over the 2-D frequency domain,

$$G(\omega_x, \omega_y) = \exp \left[\frac{-\ln \left(\frac{\sqrt{\omega_x^2 + \omega_y^2}}{\omega_0} \right)^2}{2 \ln \left(\frac{k}{\omega_0} \right)^2} \right] \quad \text{with} \quad \frac{2\pi \cdot k}{\omega_0} = 2^{\left(-\frac{b}{2\sqrt{\ln 4}} \right)}, \quad (6.9)$$

where ω_0 is the centre frequency of the filter and b is the bandwidth in octaves. The log-Gabor filter in the spatial domain is $g(x, y) = \mathcal{F}^{-1} \{G(\omega_x, \omega_y)\}$, where $\mathcal{F}^{-1}$ is the inverse Fourier transform. The odd and even responses, $o(x, y)$ and $e(x, y)$, of the 2-D image $I(x, y)$ are

$$o(x, y) = \sqrt{\mathcal{H}_1\{I * g\}^2 + \mathcal{H}_2\{I * g\}^2} \quad \text{and} \quad (6.10)$$

$$e(x, y) = |I * g|. \quad (6.11)$$

The local energy $E(x, y)$ and feature asymmetry $FA(x, y)$ are

$$E(x, y) = \sqrt{o(x, y)^2 + e(x, y)^2} \quad \text{and} \quad (6.12)$$

$$FA(x, y) = \frac{\lfloor |o(x, y)| - |e(x, y)| - T \rfloor}{E(x, y) + \epsilon}, \quad (6.13)$$

where T is a threshold value, ϵ is small, and $\lfloor x \rfloor$ is equal to x if $x > 0$ and zero otherwise. The local edge orientation ϕ is

$$\phi(x, y) = \arctan \left(\frac{\mathcal{H}_2\{I * g\}}{\mathcal{H}_1\{I * g\}} \right). \quad (6.14)$$




Figure 6.6: Comparison of gradient edge detector (left) with feature asymmetry edge detector (right). Top: short-axis slices; Bottom: long-axis slices.

Together, $FA(x, y)$ and $\phi(x, y)$ define the strength and direction of the local image edge feature. Feature asymmetry – after directional non-maximum suppression – was substituted for gradient-based measures in Equation 6.6: $FA(x, y)$ replaces $G(x, y)$ and $\phi(x, y)$ replaces $\theta_g(x, y)$ (Figure 6.6).

6.3.3 Annular Accumulator

In order to detect an annular LV myocardial cross-section in short-axis slices, both the endocardial and epicardial borders must be considered. Values of the accumulator $\mathbf{H}$ much greater than zero are strong circular responses with bright pixels on the outside (e.g. an endocardial border), and values less than zero represent edges with dark pixels on the outside (e.g. an epicardial border). The annulus of the left ventricular wall is described by these two concentric circles.

For a given potential centre point (c_x, c_y) , the endocardial and epicardial radii




Figure 6.7: Two-dimensional plot of centre likelihood scores $s(c_x, c_y)$. Brighter areas are more likely to be the centre of the left ventricle annulus.

r_{endo} and r_{epi} with the strongest image support are

$$r_{\text{endo}}(c_x, c_y) = \arg \max_r \mathbf{H}(c_x, c_y, r) \quad \text{and} \quad (6.15)$$

$$r_{\text{epi}}(c_x, c_y) = \arg \min_{r > r_{\text{endo}}} \mathbf{H}(c_x, c_y, r). \quad (6.16)$$

The accumulator values at these radii are used to assign each potential centre a likelihood score s . The score is a weighted sum of the endocardial and epicardial accumulator values (Figure 6.7),

$$s(c_x, c_y) = k \mathbf{H}(c_x, c_y, r_{\text{endo}}) - \mathbf{H}(c_x, c_y, r_{\text{epi}}), \quad (6.17)$$

where k is a weighting parameter. $k = 2$ was selected empirically for this application, as the endocardial border was more reliably present in the image than the epicardial border. The annular segmentation of the left ventricular myocardium in the given short-axis slice has the centre (c_x, c_y) that maximises s , and the corresponding endocardial and epicardial radii.

6.3.4 Radial Constraint

Particular issues were encountered in segmentation of the apical short axis slices, which often contain minimal edge features. In several cases, weak edge elements resulted in selection of an annulus far outside the image space. In response, the

annulus parameter search space was limited to include only radii less than or equal to the radius of facing edges found in the image slice.

For each edge pixel, appropriately aligned opposing edges – those that would be expected at the opposite side of the LV wall – were identified based on position and edge direction. The semi-distance (i.e. radius) between each such pair of matching edges was recorded, and all radial searches in accumulator space were limited to values below the 8th decile, plus standard deviation, of these edge-pair radii.

6.3.5 Comparison Criteria

Each modification of the basic Hough transform (Section 6.3.1) was applied to the algorithm sequentially. Every resulting segmentation algorithm was applied to the 5641 short-axis slices, and the results were compared to manual segmentation by the following criteria:

Discrepancy in centre coordinate and radii The centre point C' of the left ventricular cross-section was defined as the centre of the bounding box of the manual segmentation. The distance ΔC between C' and the centre as determined by the modified Hough transform, $C = (c_x, c_y)$, was calculated.

For eight angles ($\theta = \frac{n\pi}{4}$), the endocardial and epicardial radii of the manual segmentation, r'_{endo} and r'_{epi} respectively, were measured. These were compared to the corresponding radii of the Hough transform annulus, r_{endo} and r_{epi} , which for comparison purposes were also measured from the manual segmentation centre C (Figure 6.8). Comparison was by root mean square difference over the eight sets of radii measured for a given short-axis slice:

$$\Delta r_{\text{endo}} = \sqrt{\sum_{\theta} [r_{\text{endo}}(\theta) - r'_{\text{endo}}(\theta)]^2} \quad \text{and} \quad (6.18)$$

$$\Delta r_{\text{epi}} = \sqrt{\sum_{\theta} [r_{\text{epi}}(\theta) - r'_{\text{epi}}(\theta)]^2}. \quad (6.19)$$




Figure 6.8: Manual and Hough transform segmentations with evaluation criteria. (a) Short-axis slice. (b) Manual segmentation. C' is the centre point of the bounding rectangle. Radii r'_{endo} and r'_{epi} are measured from C' at eight angles θ . (c) Hough transform segmentation and its centre C . r_{endo} and r_{epi} are measured from C , the manual segmentation centre.

Dice coefficient For each slice, the manual segmentation M and modified Hough transform segmentation H of the LV myocardium were evaluated for pixel overlap, quantified with the Dice coefficient^[271,272],

$$d = \frac{2 \cdot |M \cap H|}{|M| + |H|}. \quad (6.20)$$

Segmentation failure A segmentation was determined to have ‘failed’ entirely if the manual segmentation centre did not fall within the endocardial border of the Hough transform segmentation. Because radii are measured for validation based on the manual segmentation centre, ‘failed’ segmentations could not be meaningfully compared based on radii, and were excluded from that measurement.

6.4 Results

The results of the modified Hough transformation of the left ventricular myocardium from 5641 short-axis slices are summarised in Table 6.1. The unmodified, gradient-based, circular Hough transform performed relatively poorly, exhibiting centre coordinate errors of 8.6 ± 12.4 mm (mean $\pm$ standard deviation) and endocardial radius errors of 4.5 ± 6.3 mm. A full 7.5% of these slices failed, i.e. the Hough endocardial boundary did not contain the manual segmentation centre point. The Dice coefficient was 0.58 ± 0.23 .

Algorithm Modification	Coordinate discrepancy (mm)			Dice's d	Fail (%)
	C	r_{endo}	r_{epi}		
None	8.6 ± 12.4	4.5 ± 6.3		0.58 ± 0.23	7.5
Basal 80%	7.8 ± 12.5	3.9 ± 5.3		0.62 ± 0.22	5.9
Apical 20%	11.3 ± 11.8	6.9 ± 8.9		0.44 ± 0.23	13.1
FA Edges	5.0 ± 6.3	3.7 ± 5.9		0.66 ± 0.21	1.5
Basal 80%	4.7 ± 6.1	3.4 ± 5.7		0.68 ± 0.20	1.1
Apical 20%	6.0 ± 6.7	4.6 ± 6.5		0.59 ± 0.21	4.3
Annular	3.2 ± 4.4	2.5 ± 3.6	3.9 ± 5.7	0.74 ± 0.20	1.5
Basal 80%	2.6 ± 4.0	2.0 ± 2.7	2.3 ± 3.0	0.81 ± 0.13	0.6
Apical 20%	5.4 ± 5.3	4.2 ± 5.6	9.8 ± 8.6	0.49 ± 0.23	4.8
R. Limit	2.5 ± 2.2	2.2 ± 2.6	3.1 ± 4.0	0.77 ± 0.16	1.4
Basal 80%	2.3 ± 2.0	2.0 ± 2.3	2.1 ± 2.6	0.82 ± 0.10	0.07
Apical 20%	3.5 ± 2.6	3.2 ± 3.2	6.8 ± 6.0	0.59 ± 0.20	6.3

Table 6.1: Results (mean $\pm$ standard deviation) of the modified Hough transform segmentation of short-axis slices from 3-D echocardiogram images. Modifications were added progressively, in the order described in Sections 6.3.1 to 6.3.4.

Inspecting the data by long axis coordinate (Figure 6.9), we found that the most apical selection of short-axis slices has significantly higher errors and lower d than the rest. This is to be expected: as a 2-D method operating on slices from a 3-D dataset, the Hough transform is particularly sensitive to lower image quality as often found at the most apical portion of parasternal 3-D echocardiograms. Without the context of long-axis slices, there is very little edge information available.

Incorporating feature asymmetry as an edge detector resulted in significant gains (Figure 6.10). Mean centre coordinate and radius errors dropped to 5.0 and 3.7 mm, respectively, and the Dice coefficient increased to 0.66 ± 0.21 . The failure rate also decreased dramatically, to 1.5 %. Qualitatively, there was also more consistency between sequential short-axis slices (Figure 6.10, left).

Adding the two-step voting procedure to take into account epicardial edge information had the most dramatic effects on mean centre coordinate discrepancy, which dropped from 5.0 mm to 3.2 mm. While the apical slices were still poorly segmented (Dice coefficient < 0.5), segmentation for the basal 80 % of slices was reasonable ($d = 0.81 \pm 0.13$).




Figure 6.9: Centre coordinate error (middle) and Dice coefficient (right), by long axis quintile, of traditional Hough transformation (endocardial border only, gradient feature) compared to manual segmentation. Left: Example Hough segmentation in short-axis (top) and long-axis (bottom) views.




Figure 6.10: Centre coordinate error (middle) and Dice coefficient (right), by long axis quintile, of feature asymmetry Hough transformation (endocardial border only) compared to manual segmentation. Left: Example Hough segmentation in short-axis (top) and long-axis (bottom) views.




Figure 6.11: Centre coordinate error (middle) and Dice coefficient (right), by long axis quintile, of annular Hough transformation (feature asymmetry) compared to manual segmentation. Left: Example Hough segmentation in short-axis (top) and long-axis (bottom) views.




Figure 6.12: Centre coordinate error (middle) and Dice coefficient (right), by long axis quintile, of full modified Hough transformation (annular, feature asymmetry, radial search constraint) compared to manual segmentation. Left: Example Hough segmentation in short-axis (top) and long-axis (bottom) views.




Figure 6.13: Feature asymmetry edge map of 3-D echocardiogram data, filtered for endocardial (magenta) and epicardial (cyan) edge information.

Finally, adding the matching-opposite-edges radius search limit had two major positive effects (Figure 6.12). First, the average segmentation quality in the apex – as measured by the Dice coefficient – increased from 0.49 to 0.59. Second, the failure rate in the basal 80 % of the heart dropped to nearly nil, with only three of 4416 individual segmentation tasks failing (0.07 %).

Specifically evaluating the performance of the fully modified Hough algorithm on the basal portion of the LV myocardium, results are encouraging. All mean positional errors were between 2.0 and 2.3 mm, and the Dice coefficient was 0.82.

6.5 Conclusions

The fully modified Hough transform method produces annuli that differ from the gold-standard manual segmentation by an average of 2.5 - 3.0 mm over the whole left ventricle and 2.0 - 2.3 mm in the apical 80 %. This is similar performance to other state-of-the-art techniques^[262], which are capable of achieving accuracy in the range of 1.4 - 1.8 mm using methods that consider speckle statistics and




Figure 6.14: Preliminary alignment of 3-D echocardiogram (yellow) and MR image (grey) based on left ventricular segmentation.

myocardial tissue characteristics. The Hough method carries the added benefit that, as a voting method, it may be arbitrarily parallelised in order to achieve near-real-time results. The modified Hough transform method may be useful for quickly estimating physiological parameters and basic LV shape information.

This transform may also serve to initialise iterative segmentation techniques, mitigating the pitfalls of spurious edges and data dropout which typically plague these methods in 3-D echocardiogram data. Furthermore, the approximate centre information gleaned from the Hough transform can be used to filter 3-D feature asymmetry edge information by direction (Figure 6.13), revealing clear edge maps for endocardial and epicardial borders. Proper initialisation and clear edge maps are critical in allowing iterative segmentation techniques to converge correctly.

6.5.1 Future Directions

One primary motivation for the rapid segmentation that can be achieved by the Hough transform is to allow near-real-time cross-modality image registration. When

registering pre-operative images (for example, MR or CT) with intra-operative imaging (echocardiography), speed is critical. Rapid segmentation of operative imaging allows immediate alignment by centre of mass and moments of inertia of the segmented masses. Preliminary work in this direction has yielded promising results (Figure 6.14). These ideas are explored further in Chapter 7.

The temptation to form premature theories upon insufficient data is the bane of our profession.

—Sir Arthur Conan Doyle’s Sherlock Holmes^[273]

7

Future Work

Contents

7.1	Introduction	136
7.2	MR-Compatible EP Lab Paradigm	137
7.2.1	Calibration	139
7.2.2	Projection	140
7.2.3	Conclusion	141
7.3	Pre-/Intra-Op Image Registration	141
7.3.1	3-D Segmentation-Based Registration	143
7.3.2	Temporal Alignment for 3D+t Registration	145
7.3.3	Conclusion	146
7.4	Conclusion	146

7.1 Introduction

Over three decades of development, guidance of interventional electrophysiology^[124,125] has advanced from reliance on fluoroscopy alone^[274] to incorporation of integrated 3-D imaging and advanced electroanatomic mapping techniques^[11,14,275,276] (Section 2.3). This progress has been driven by the coevolution of imaging technology and clinical approach. A typical ablation procedure today begins with a 3-D MR or CT imaging session at least several hours before the ablation. These images are analysed to ascertain the anatomic layout of the heart and to recognise functional

considerations – areas of fibrosis, dysmotility, etc. – that will inform the intervention. The images are also manually registered and incorporated into the EAM toolkit using commercial systems^[5,192].

During the ablation procedure, a commercial EAM system presents the cardiologist with a schematic of the cardiac chamber derived from a segmentation of the pre-operative MR or CT image. As the catheter is manipulated, its position – calculated by magnetic or electrical impedance localisation systems – is represented on this schematic. When local electrograms are recorded and ablation lesions are created, their properties – bipolar voltage, activation timing, etc. – appear on the myocardial schematic as dots and colour codes.

The work presented in this thesis is a series of steps toward two new EP imaging paradigms. Each paradigm involves a fundamental advance in the role of interventional imaging. Under the first scheme, the development of MR-compatible electrophysiology equipment, coupled with intra-operative registration protocols and MR thermography, will permit MR guidance and monitoring of intracardiac ablation (Section 7.2). Under the second model, rapid geometric registration techniques, paired with near-real-time segmentation of intra-operative imaging, will allow beating-heart guidance of catheter procedures (Section 7.3).

In this chapter, we present these new clinical paradigms and describe some of the remaining obstacles on the path forward.

7.2 MR-Compatible EP Lab Paradigm

As more is understood about the utility of MR in tissue characterisation (Chapters 4 and 5), the advantages of deeper integration of advanced image information in percutaneous procedures is becoming evident. By combining electrocardiographic imaging (ECGI), a developing method of back-projecting surface electrograms onto the myocardium^[8,277,278], with the bipolar information coded in MR images, it is conceivable that much of the local electrogram mapping in an electrophysiology procedure could be superseded by non-invasive imaging techniques.

Beyond procedural planning, intra-operative imaging of EP ablations holds great promise. The most feared complications of radiofrequency ablation – cardiac perforation, phrenic nerve injury, and atrioesophageal fistula – are results of thermal injury^[175,176,234]. Conversely, treatment failure is often caused by electrical reconnection of pulmonary veins to the atrial myocardium, which can result from insufficiently aggressive ablation. Though modern techniques incorporate ablation catheters with temperature probes as well as oesophageal thermometers in order to prevent injury, measurement the temperature of the tissue itself during ablation – with fine spatial discrimination – will be a significant advance in safety. MR thermography, a modality based on temperature-dependence of the proton resonance frequency, has been evaluated as a way to achieve this level of ablation monitoring^[279–281].

Traditional percutaneous EP equipment includes metals that are incompatible with the magnetic fields and radiofrequency pulses necessary for MR imaging. Enthusiastic work is underway to build MR-compatible electrophysiology laboratories, and success is being reported^[91]. Even when MR-compatible electrophysiology equipment becomes widely available, exclusive use of MR imaging to provide guidance will be limited by radiofrequency interference. The radio pulses used in MR induce noise in catheter electrical systems, rendering local electrograms temporarily unreliable. Furthermore, generating a 3-D MR image with adequate spatial density to provide direct catheter guidance cannot occur in near-real-time.

However, by establishing careful electrical isolation and taking advantage of MR-compatible EP equipment^[88,91,192], a hybrid system will be possible. Under this paradigm, a pre-operative MR image will be incorporated into the EAM system with continuous registration. During the procedure, intermediate 2-D images will be acquired to verify catheter contact with – and position on – the myocardium, and MR thermography will be used to monitor ablation.

Continuous registration implies two distinct coordinate spaces: $\mathbf{M}$, corresponding to the EAM, and $\mathbf{I}$, corresponding to the MR image. The clinician will use a hybrid catheter, which incorporates electrodes for EAM as well as tracking coils for MR. A

registration step will be included in the EP protocol to facilitate generation of a mapping function $\mathbf{R}(\mathbf{p}) : \mathbf{M} \rightarrow \mathbf{I}$ from EAM space to image space, incorporating the registration principals described in Chapter 3. For the remainder of the procedure, $\mathbf{R}$ will be employed to project the real-time catheter position as detected by EAM onto MR image space for image-based guidance.

7.2.1 Calibration

Both the MR image space $\mathbf{I}$ and the EAM space $\mathbf{M}$ are three-dimensional real spaces. While $\mathbf{I}$ has orthonormal bases and a linear mapping to physical coordinates, $\mathbf{M}$ is less well-defined, and may be only analytic over local regions, depending on the implementation of the EAM system. The goal of the clinical calibration step will be to generate a collection of N individual correspondences $r_n : \mathbf{M} \rightarrow \mathbf{I}$ with sufficient spatial density to blanket the anatomical region of interest. Each correspondence r_n is defined by the pair of corresponding vectors ($\mathbf{p}_n \in \mathbf{M}, \mathbf{p}'_n \in \mathbf{I}$). The calibration step, then, consists of finding these pairs.

Finding $\mathbf{p}_n$

$\mathbf{p}_n$ represents the location in EAM space $\mathbf{M}$ of the MR tracking coil. Because the coil cannot be directly detected via the EAM system, our ability to find $\mathbf{p}_n$ will depend on inferring its most likely position from the coordinates of EAM electrodes. In the simplest case, the calibration catheter will incorporate three tracking elements with a small, uniform spacing between them. The outer elements will be EAM electrodes, while between them is the MR coil. As $\mathbf{M}$ is locally continuous and the distance between markers is small, $\mathbf{p}_n$ may be reasonably estimated by the mean of the EAM electrode positions. Although this does not account for curvature of the catheter, this effect will likely be insignificant over the small distances involved.

In order to account for curvature, a more complex calibration catheter might comprise five elements: four regularly-spaced EAM electrodes and a centrally positioned MR tracking coil. The curve of the catheter between the four EAM

electrodes may be estimated by a cubic Bézier curve, thereby providing an estimate of $\mathbf{p}_n$ which considers catheter curvature.

Finding $\mathbf{p}'_n$

Unlike the EAM electrodes, the MR tracking coil will be easily detected by a set of three MR sequences with orthogonal gradient configurations, providing $\mathbf{p}'_n \in \mathbf{I}$ directly.

7.2.2 Projection

After recording multiple correspondences $r_n = (\mathbf{p}_n \in \mathbf{M}, \mathbf{p}'_n \in \mathbf{I})$ in the calibration step, the task will turn to providing procedural guidance. Throughout the procedure, the EAM apparatus will provide the position $\mathbf{p} \in \mathbf{M}$ of EAM electrodes. Image-based guidance will rely on determining the corresponding position $\mathbf{R}(\mathbf{p}) \in \mathbf{I}$ for overlay on the MR image.

As $\mathbf{I}$ corresponds to physical coordinates, $\mathbf{M}$ is locally analytic, and the function $\mathbf{R}(\mathbf{p}) : \mathbf{M} \rightarrow \mathbf{I}$ is locally well-behaved, it follows that $\mathbf{R}(\mathbf{p})$ is best defined by reference to local correspondences r_n . Taking a mean weighted by inverse distance, we find

$$\mathbf{R}(\mathbf{p}) = \left(\sum_{n=0}^N \frac{1}{\|\mathbf{p} - \mathbf{p}_n\|} \right)^{-1} \sum_{n=0}^N \frac{1}{\|\mathbf{p} - \mathbf{p}_n\|} \mathbf{p}'_n. \quad (7.1)$$

For computational efficiency, as the influence of correspondences r_n falls rapidly with distance $\|\mathbf{p} - \mathbf{p}_n\|$, only the few most influential terms will be included.

This method of projection will lend itself to internal validation of the calibration sequence. Each time the MR tracking sequence is executed while the calibration catheter is in the anatomical region of interest, a new correspondence $r = (\mathbf{p}, \mathbf{p}')$ will be generated. Before adding this correspondence to the mapping function $\mathbf{R}$, a local registration error ϵ can be computed as

$$\epsilon = \|\mathbf{p}' - \mathbf{R}(\mathbf{p})\|. \quad (7.2)$$

As ϵ is a distance in $\mathbf{I}$, it may be likened to the registration error in physical coordinates. During the clinical calibration step, the anatomical region of interest

will be sampled until ϵ is within an acceptable range for the intervention. During the procedure, this validation may be repeated at will to continuously monitor the quality of the registration. In the event ϵ is found to be outside a defined range, the existing set of r_n 's may be discarded and the calibration step repeated.

7.2.3 Conclusion

This section has laid out one procedural vision for the MR-compatible electrophysiology lab. Pre-operative imaging will be used to plan the procedure, with the benefit of interpretation by machine learning algorithms like those developed in Chapters 4 and 5. Intra-operatively, the image space will be registered with the electroanatomic mapping system in a continuous and self-monitoring manner. MR thermography will be used to monitor ablation, improving isolation and reducing complications.

Critically, as soon as this system is implemented, it will necessarily generate a huge amount of well-paired MR and EAM data for machine learning training. This will result in a virtuous circle whereby procedural data results in higher prediction accuracy and more effective procedures.

7.3 Pre-/Intra-Op Image Registration

Advances in 3-D cardiac ultrasound (3DUS)^[42,43] are building toward another novel approach to procedural imaging. Unlike fluoroscopy and 2-D ultrasound, 3DUS provides the physician a complete view of the working area in near-real-time. Fusion echocardiography is expanding this view further by incorporating the best features of multiple images into an expansive clinical panorama^[252,282]. Ultrasound avoids the radiation risk inherent to fluoroscopy^[96,97] and is compatible with interventional equipment. However, it suffers from low CNR, pronounced image artefacts, and low spatial resolution^[41] compared to MR and CT.

Interventional and diagnostic cardiology stand to benefit from an algorithm to quickly and accurately align pre-procedure MR and CT images with procedural 3DUS images. This algorithm will allow cardiologists to be presented MR- or CT-quality data synchronised to near-real-time 3DUS images, facilitating extremely




Figure 7.1: Pipeline approach to intraoperative guidance using preoperative advanced imaging.

accurate, beating-heart guidance. Diagnostic cardiologists will use this technology to correlate high-resolution structural and functional information with valvular and kinetic defects that are seen most clearly on echocardiograms. Chronic heart disease monitoring and long-term clinical studies will make use of multi-modality strategies to minimise the number of costly MR scans or radiation-intensive CT scans required to track pathology over time.

Currently, near-real-time multi-modality registration in three dimensions is prohibited by computational limits. We propose a pipeline approach that integrates preoperative 3DUS data in order to take the multi-modality registration step offline (Figure 7.1). The results of this offline registration will be indexed and later recalled

to be convolved with the near-real-time, single-modality registration, generating a complete transform from intraoperative 3DUS image to preoperative MR image.

Even without the near-real-time constraint, 3-D multi-modality registration is not trivial. Traditionally, image registration techniques have fallen into one of two broad categories: intensity-based and feature-based registrations. Intensity-based techniques use pixel-wise comparisons between images to measure similarity, iterating transform parameters until the similarity metric is maximised^[283–285]. By contrast, feature-based methods rely on the extraction of salient feature points, such as corners and edges, from images. Corresponding features are then brought into alignment by iterative minimisation of a distance-based energy function^[286–288].

Intensity-based techniques are most effective for single-modality registration, in which similar tissues are represented by similar relative voxel values. Gains have been made in intensity-based, multi-modality registration by using cross-correlation- and information-based metrics^[285], pre-filtering input images^[289,290], and even generating simulated ultrasound images from MR or CT data^[291]. However, the task remains difficult, especially for images with disparate spatial resolution and tissue response.

We propose a geometry-based, multi-modality image registration technique to align cardiac MR or CT and 3DUS images and image sequences. The rigid alignment of two 3-D images is defined by six degrees of freedom: three translation parameters and three rotation parameters. Accurate segmentation of the left ventricle in each of the MR and 3DUS images can be used to quickly solve five of these parameters by aligning the centre of mass and long axis of the respective segmentations. The time parameter will be solved by a more straightforward 1-D optimisation. This algorithm has been preliminarily assessed in a small number of images using the Hough results described in Chapter 6 and corresponding MR data. A representative image is shown in Figure 6.14 (page 134).

7.3.1 3-D Segmentation-Based Registration

From the ultrasound segmentation technique developed in Chapter 6, we will extract a set L_{US} of point coordinates $\mathbf{p} = (x, y, z)$ in a 3-D ultrasound image I_{US} that

represent the left ventricular myocardial segmentation. Likewise, region filling will be used to determine point sets $L_{\text{MR/CT}}$ and $R_{\text{MR/CT}}$ that represent the left and right ventricle segmentations, respectively, of the 3-D MR or CT image $I_{\text{MR/CT}}$. As this medical image data will contain embedded resolution information, the spatial scaling relationship between the image domains is known. For the purposes of this alignment, we will disregard nonlinear cross-modality effects. Therefore, there are six independent degrees of freedom in the registration problem – three translational and three rotational.

Translational alignment will be achieved by calculating the centre of mass of the LV segmentations L_{US} and $L_{\text{MR/CT}}$. The vector-valued centre-of-mass function $\text{CoM}(S)$ operates on a set of points S and is defined by the mean

$$\text{CoM}(S) = \frac{1}{|S|} \sum_{\mathbf{p} \in S} \mathbf{p}. \quad (7.3)$$

A translation transform will bring the LV centres together, resolving the three translational degrees of freedom of the alignment problem. Two of the three rotation parameters will be determined by aligning the left ventricular long axes. Consider the 3×3 moment of inertia tensor $\mathcal{I}$ of a set S of 3-D points. As our segmentations will be of uniform ‘density’, the tensor is

$$\mathcal{I}(S) = \frac{1}{|S|} \sum_{\mathbf{p} \in S} (\|\mathbf{p}\|^2 E_3 - \mathbf{p} \otimes \mathbf{p}) \quad (7.4)$$

$$= \frac{1}{|S|} \sum_{(x,y,z) \in S} \left((x^2 + y^2 + z^2) \begin{bmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{bmatrix} - \begin{bmatrix} xx & xy & xz \\ yx & yy & yz \\ zx & zy & zz \end{bmatrix} \right), \quad (7.5)$$

where E_3 is the 3×3 identity matrix, and $\otimes$ denotes the outer product. A given figure’s principal moments of inertia $\mathcal{I}_1$, $\mathcal{I}_2$, and $\mathcal{I}_3$ are the eigenvalues of its moment of inertia tensor, and the figure’s principal axes are the corresponding eigenvectors.

The long axis of the LV will correspond to the first principal axis of the segmentation (i.e. that with the smallest associated principal moment). These long axes will be calculated from the 3DUS and MR or CT segmentations, $\mathbf{l}_{\text{US}}$ and $\mathbf{l}_{\text{MR/CT}}$, respectively. In the coordinate system with the origin at the coincident

centres of mass of each image, the transform required to align these long axes will be a rotation of θ about the rotation axis $\hat{r}$ where

$$\begin{aligned} \theta &= \arccos(\hat{l}_{\text{US}} \cdot \hat{l}_{\text{MR/CT}}) \quad \text{and} \quad \begin{cases} \hat{l}_{\text{US}} = \mathbf{l}_{\text{US}} / \|\mathbf{l}_{\text{US}}\| \\ \hat{l}_{\text{MR/CT}} = \mathbf{l}_{\text{MR/CT}} / \|\mathbf{l}_{\text{MR/CT}}\| \end{cases} \\ \hat{r} &= \hat{l}_{\text{US}} \times \hat{l}_{\text{MR/CT}}. \end{aligned} \quad (7.6)$$

After translating the LV centres of mass to the origin and rotating their long axes into coincidence, one degree of freedom will remain: the angle of rotation ϕ about the long axis. In addition to rotating the LV properly, this final step will bring the right ventricle (RV) into alignment. We will rely on this RV alignment to identify ϕ , by evaluating the fit of the MRI RV segmentation at various rotations on the fusion echo image. First, an auxiliary image $A_{\text{US-RV}}$ will be computed from the fusion echo image I_{US} to emphasise the dark interior and bright edges of the RV in ultrasound:

$$A_{\text{US-RV}} = I_{\text{US}} + \|\nabla [\mathcal{G}(\sigma) * \text{Median}_s(I_{\text{US}})]\|. \quad (7.7)$$

Here, ∇ denotes the image gradient, Median_s is a median filter with radius s , and $\mathcal{G}(\sigma)$ is a Gaussian kernel. A second artificial image will be generated from the segmentation $R_{\text{MR/CT}}$ that is bright around the exterior and dark on the interior. (In our preliminary investigations, this image was generated by applying a Gaussian filter followed by a gradient magnitude filter to the mask image representing the segmentation.) These two artificial images will then be compared over various values of ϕ , to identify the last parameter for rigid alignment.

7.3.2 Temporal Alignment for 3D+t Registration

In order to use the above 3-D registration technique to align cine (3D+t) data, two final degrees of freedom will need to be resolved: temporal shift t_0 and scale t_s . Echocardiogram and cine MR images are generally provided as single-heartbeat datasets, so these values will be known approximately ($t_0 \approx 0$ and $t_s \approx 1$). However, variations in imaging technique or heart rate during imaging sessions will cause minor perturbations. A comparison of the volumes of the LV segmentations from the fusion echo and MR images over time, $v_{\text{US-LV}}(t)$ and $v_{\text{MR/CT-LV}}(t)$ respectively, will be used to refine t_0 and t_s .

As clinical MRI data will usually be more temporally dense than the fusion echo data, the MRI segmentation volume will be subsampled over time. A root mean squared difference metric $\text{RMSD}(t_0, t_s)$ will be used to evaluate the performance of the linear time transformation $T(t) = t_0 + t_s t$ over the set of time points $F_{\text{US}} = \{t : \exists v_{\text{US-LV}}(t)\}$ where ultrasound frames are defined:

$$\text{RMSD}(t_0, t_s) = \sqrt{\frac{1}{|F_{\text{US}}|} \sum_{t \in F_{\text{US}}} [v_{\text{US-LV}}(t) - v_{\text{MR/CT-LV}}^*(T(t))]^2} \quad (7.8)$$

Minimising $\text{RMSD}(t_0, t_s)$ will provide temporal alignment of the 3D+t data, completing the image registration.

7.3.3 Conclusion

This section has described one practical application for the quick geometric segmentation algorithm introduced in Chapter 6. This registration, fitted into a pipeline approach as described above, will bring beating-heart, MR- and CT-quality image data into the electrophysiology laboratory even without expensive and complex hybrid labs. Clinically, this approach will be combined with instrument segmentation techniques^[292–294] to provide a more accurate picture of catheter contact during cardiac motion. It will also facilitate operative targeting of areas of interest identified in the pre-procedure images.

7.4 Conclusion

The paradigms described above constitute two paths by which the geometric image analysis techniques developed in this thesis will make their way to the clinical electrophysiology laboratory. More broadly, as machine learning and computer vision toolsets continue to grow, they will frame a shift in our consideration of medical imaging, and especially procedural imaging. While clinical images traditionally make up a static piece of the diagnostic algorithm – to be ordered, consulted, and then stored away – our advancing data processing and machine learning abilities will change this role. Imaging studies will instead be considered a deposit in a patient’s global clinical model. The model will include the whole of what today is

considered a medical record, from tissue-level imaging data to demographics and family history, to be interrogated by sophisticated machine learning algorithms when a clinical question arises.

Tout ce que je fais, je le fais pour que cela serve.

Everything I do, I do in order that it may be of use.

—Michel Foucault^[295]

8

Conclusion

We are in a transformative era for interventional cardiology and electrophysiology. Through the first twenty years of catheter interventions, clinicians and researchers focused on replicating surgical procedures with fluoroscopy-guided tools (Chapter 2). Progress was made by introducing more manoeuvrable catheters capable of more intricate manipulation. The goals of a given procedure were well-known in advance. Where intraoperative decision-making was used, it was in a very limited context, e.g. to respond to fluoroscopic angiography during an angioplasty. While cardiac imaging continued to advance, its clinical use was partially limited by the time, expense, and expertise required to acquire useful information from an image. Even where images were informative, that information might not affect the choice among a limited range of clinical interventions.

Over the past fifteen years, however, tremendous progress has been made on three entwined paths for integrated, image-guided intervention. First, the cost and availability of advanced imaging has become less prohibitive, increasing utilisation and driving innovation. Second, the precision of interventional procedures has improved dramatically, introducing tangible clinical decisions that benefit from high-resolution tissue information. Third and finally, computational power and image processing algorithms have advanced, allowing previously unthinkable integration of different clinical images and data sets.

The results presented in this thesis are steps along this third path. A simple, efficient registration algorithm was designed to register EAM points to cardiac MR and CT images (Chapter 3). While this algorithm is not groundbreaking in its approach, its simplicity enables the robust analysis of image features in the chapters that follow.

Current discourse on catheter ablation of atrial fibrillation includes significant attention to the role of fibrosis in pulmonary vein isolation outcomes (Section 4.1.1) and the ability of advanced imaging to identify that fibrosis (Section 4.1.2). This research, however, has thus far drawn purely correlative conclusions, with the underlying presumption that predictive accuracy would be built on that foundation. The machine learning approach taken in Chapter 4 has made this leap, for the first time quantifying the predictive power of LGE-MRI and CE-MDCT in identifying atrial fibrosis.

The significance of this work was not only to demonstrate prediction of bipolar voltage values, but also to more generally underscore the utility of a machine learning approach to analyse large amounts of complex image data. While this perspective is well established in the field of computer vision, it is still gaining traction in the clinical community. With the demonstration of predictive power, it becomes reasonable to consider using image information alone to make clinical decisions – in this case, selecting likely responders to pulmonary vein isolation while guiding less appropriate candidates toward more aggressive pharmacological rate control.

In Chapter 5, we investigated the relationship between MR image characterisation of natural fibrosis and post-ablation lesion, finding that while the two were generally similar, post-ablation scar was overall a ‘more extreme’ version of natural fibrosis, with higher levels of enhancement and heterogeneity. Just as importantly, the results demonstrated that the random forest bipolar voltage regression model built using images of naïve myocardium (Section 4.4.3) was inadequate to characterise ablation scar. As machine learning becomes more fundamental to medical decision making, the limitations of a given training set must

become as deeply integrated into the academic medicine lexicon as randomisation and placebo control are today.

Finally, in Chapter 6 we validated one more piece in the near-real-time segmentation and registration repertoire: a method for rapid segmentation of left ventricular myocardium in 3-D echocardiograms. A modified Hough transform achieved near-universal success in non-apical 3-D echocardiogram slices. This approach deliberately sacrifices pixel-perfect accuracy for algorithmic simplicity and speed, with an eye towards near-real-time registration with pre-operative images in a pipeline approach (Section 7.3). Tools of this kind will gain increased importance in years to come, as quick, robust links are needed to construct the pipeline methods that bring advanced imaging into the operating room.

In the short term, advanced imaging tools, including machine learning, will greatly improve our ability to make clinical decisions using advanced imaging. Computer-aided diagnosis (CADx) is already a burgeoning technology in radiology, with demonstrated successes primarily in tumour detection^[296]. At its most basic level, prediction of bipolar value from MR and CT images is an implementation of CADx that will assist clinicians in identifying appropriate candidates for ablation of atrial fibrillation.

In the intermediate future, integrated imaging solutions will allow for better procedure planning and post-operative evaluation. High-accuracy tissue characterisation and rapid registration will allow fibrosis imaging to be fused with electrocardiogram data and integrated into a complete electrical model^[106] of the myocardium. This model, constructed before the patient is on the EP laboratory table, will identify potential arrhythmogenic foci and form a basis for predicting effective procedural induction strategies and ablation targets. This pre-operative planning will shorten procedure times and decrease the amount of ablation energy needed to achieve arrhythmia termination, improving patient safety while enhancing treatment efficacy.

The long-term implications of this work derive from widespread adoption of functional image guidance. The development of electroanatomic mapping allowed

the position of an intra-cardiac catheter – and by association all of the information encoded in local electrograms – to move from 2-D fluoroscopy planes onto 3-D virtual spaces within the memory of EAM computers. Analogously, as advanced imaging technology expands the role of intraoperative imaging, its utility will advance beyond mere instrument guidance to monitoring the progress of treatment from planning to completion to follow-up. This shift will be driven by practicality, delivered by rapid registration algorithms, utility, in the form of machine learning-driven interpretation of advanced modalities such as MR thermography, and clinical efficacy, measured in complications averted and repeat procedures avoided.

The work presented in this thesis is one step in the long journey from medical imaging as a static snapshot to imaging as an open window to structure and function, to be accessed regularly – pre-, intra-, and post-procedurally – as clinically indicated. When this open window is framed in a machine learning context, each piece of imaging information is immediately associated with a full array of clinical data – demographics, laboratory values, electrograms, clinical outcomes, and other images. This association between image features and clinical parameters will reinforce itself, creating an ever-expanding base of training data upon which to build better machine learning models. Ultimately, advanced imaging modalities interpreted through broad-based models will provide a toolset that will enhance diagnostic accuracy and treatment efficacy in tomorrow’s clinical environment.

Appendices

Cor animalium, fundamentum est vitæ, princeps omnium, Microcosmi Sol, a quo omnis vegetatio dependet, vigor omnis & robur emanat.

The heart of animals is the foundation of their life, the sovereign of everything within them, the sun of their microcosm, that upon which all growth depends, from which all power proceeds.

— William Harvey^[297]



Review of Cardiac Physiology and Electrophysiology

Contents

A.1 Anatomy	153
A.2 Mechanical Cycle	160
A.3 Electrical Cycle	161
A.4 Cellular Electromechanical Coupling	164

In this appendix, we briefly review cardiac anatomy and physiology, with a focus on details relevant to interventional electrophysiology. All information presented below is well-established in cardiology and can be found in relevant medical physiology textbooks, such as Guyton and Hall^[298], the very comprehensive Boron and Boulpaep^[299], or the accessible and beautifully illustrated Silverthorn^[300]. Unless otherwise indicated, all anatomical figures in this section are adapted from Gray's *Anatomy of the Human Body*^[50].

A.1 Anatomy

The heart is the nexus of the circulatory system, located within the membranous pericardial sac in the centre of the thoracic cavity. Functionally, the heart consists




Figure A.1: Four-chamber illustration of the human heart. Clockwise from upper-left: right atrium, left atrium, left ventricle, right ventricle.

of two distinct pumps, each composed of two chambers. The right heart receives venous blood from the head, trunk, and periphery via the superior and inferior venae cavae and pumps it through the pulmonary arteries to the lungs. The left heart receives oxygenated blood from the lungs via the pulmonary veins and pumps it at high pressure into systemic circulation through the aorta. While they function at discrete locations in the circulatory loop, the left and right pumps are physically co-located, separated by the thin cardiac septum (Figure A.1).

Each of the cardiac pumps is made up of two chambers: an atrium and a ventricle. The atria are thin-walled, compliant chambers into which incoming blood flows before passing to the ventricles. The ventricles are the primary pumping chambers, and have thicker, muscular walls. The left ventricular (LV) wall, which must pump at high pressure against the total peripheral resistance, is 0.6 - 1.0 cm thick^[235] compared with less than 0.5 cm for the right ventricle (RV), which pumps blood into the low-resistance pulmonary circulation.




Figure A.2: Left atrium in posterior cardiac illustrations. Top: intact posterior cardiac view with left atrium, centre. Bottom: left-posterior oblique view with left pulmonary veins, posterior left atrial wall, and left ventricular free wall cut away. Left atrial appendage ('auricle') visible, upper-left.

The left atrium (Figure A.2) is located superiorly (basal) and posteriorly, in direct proximity to the oesophagus and left phrenic nerve (Figure A.3). These structures are at risk of damage during catheter ablation of the left atrial myocardium or pulmonary vein ostia (Section 2.4.3). The four pulmonary veins attach at the lateral posterior aspects of the left atrium, with some individual variation (Figure A.2, top).

The left atrial appendage (LAA, classically ‘auricula’) is an extension of the left atrium located anteriorly, adjacent to the pulmonary artery (Figure A.2, bottom). There is great variation in the size and shape of the LAA, and recent research has found that this variation has a bearing on stroke risk^[301–303]. While the LAA is dynamic in a healthy heart, it is a focus of blood stasis and source for thromboembolism in AF patients^[304].

Between each atrium-ventricle pair is a valve (Figure A.4). The left atrioventricular valve is called the mitral valve, classically called the ‘bicuspid valve’ due to its two leaflets. The right valve is the tricuspid valve, and has three leaflets. Each valve opening is ringed by a fibrous annulus, and the valve flaps are anchored to papillary muscles (Figure A.4, bottom) via tendons (the chordae tendonae).

There are also valves separating the ventricles from their arterial outflow tracts. Both the aortic valve – between the left ventricle and the aorta – and the pulmonary valve – between the right ventricle and the pulmonary trunk – have three leaves. These semilunar valves are not anchored to the muscular myocardium. Instead, each derives support entirely from its annulus.

The myocardium is supplied with oxygen through the coronary circulation (Figure A.5). The left and right coronary arteries arise from the root of the aorta. The left coronary artery originates from the left aortic sinus and bifurcates after 10 - 25 mm into the left anterior descending (LAD) artery and the left circumflex (LCX) artery. The LAD descends along the anterior surface of the heart within the interventricular sulcus between the left and right ventricles, usually reaching the apex. The LCX wraps posteriorly around the left side of the heart within the coronary sulcus between the left atrium and left ventricle. The right coronary artery follows the right atrioventricular groove, giving rise to one or more right




Figure A.3: Left atrium illustrated in anatomic context. Top: cross section at the level of the RV outflow tract (just superior to the left atrium). Note proximity to the oesophagus (posterior) and left phrenic nerve (left lateral). Bottom: left-posterior view with intact left atrium.




Figure A.4: Atrioventricular valves, chordae tendinae, and papillary muscles. Top: Superior view of ventricles with atria cut away. Mitral and tricuspid valve leaflets are visible. Bottom: Right lateral view of heart. Note the chordae tendinae and papillary muscles anchoring the tricuspid valve leaflets to the right ventricular myocardium.




Figure A.5: Coronary circulation (arteries, red). Left: anterior view. Right: posterior view. Adapted from *Wikimedia Commons* under a permissive license^[305].

marginal branches. The posterior descending artery (PDA) arises from the right coronary artery in 70% of people ('right-dominant'), from the LCX in 10% ('left-dominant'), and from both in the remaining 20% ('co-dominant'). The PDA follows the posterior intraventricular sulcus to the apex of the heart.

The lateral wall of the left ventricle is supplied by diagonal arteries arising from the LAD, while its posterior wall is supplied by the PDA. The anterior two-thirds of the intraventricular septum is supplied by the LAD, while the posterior third is supplied by the PDA. The anterolateral papillary muscle is supplied jointly by the LAD and the LCX, while the posteromedial papillary muscle is supplied by the PDA alone. The right ventricular free wall is supplied by the right coronary artery and its marginal branches. The left and right atria are supplied by the LCX and the right coronary artery, respectively. The SA node is supplied by the right coronary in 60% of people and the LCX in the remaining 40%, irrespective of PDA dominance. The AV node is supplied by the PDA.




Figure A.6: Wiggers diagram of the cardiac cycle. Pressures refer to the left-side chambers. Reproduced from *Wikimedia Commons* under a permissive license^[306].

A.2 Mechanical Cycle

The cardiac mechanical cycle is summarised by the Wiggers diagram of pressure and volume curves (Figure A.6). Broadly, the cycle consists of diastole, during which the ventricles are relaxed and filling, and systole, during which the ventricles contract to expel blood from the heart. With minor exceptions, the left and right side of the heart move in unison. The language below describes the left heart.

We begin at diastasis, the middle third of diastole. During this period, the ventricle has largely finished passively filling, but is still in communication with the atrium. The atrioventricular valves are open, while the semilunar valves between the ventricles and their respective arteries are closed.

Diastasis is interrupted by the atrial systole, the final third of cardiac diastole. During atrial systole, the atrium contracts, forcing a last bit of blood through

the mitral valve to ‘top up’ the ventricular volume. This atrial ‘kick’ provides about one third of the ventricular stroke volume.

The ventricular myocardium begins to contract, marking the beginning of systole. The increased ventricular pressure immediately closes the mitral valve. Pressure continues to build as the ventricular muscle contracts against an unyielding blood volume (isovolumic contraction).

When the ventricular pressure reaches the aortic pressure, the aortic valve opens and ventricular blood is ejected through it. This ejection continues until, as the ventricle relaxes, ventricular pressure falls below aortic pressure, closing the aortic valve.

With both valves again closed, ventricular pressure drops rapidly as the muscle relaxes (isovolumic relaxation). When ventricular pressure drops below atrial pressure, the mitral valve opens, and blood rapidly flows from the atrium into the ventricle. The compliance of the ventricle soon limits flow, and the heart is again at diastasis.

The normal heart beats at 60 to 100 times per minute at rest, though trained athletes may have rates as low as the 40s. Strenuous exercise can elicit a maximal heart rate that is classically estimated as $220 - a$ beats per minute, where a is the age in years.

The stroke volume is the amount of blood expelled from each ventricle during a single heartbeat, and is approximately 70 mL in a normal heart. This value is often expressed as ejection fraction, or the proportion of the end-diastolic left ventricular volume that is ejected. An ejection fraction between 55% and 70% is normal, with lower ratios generally indicating heart failure. The cardiac output is obtained by multiplying stroke volume by heart rate, and is generally between 4 and 8 litres of blood per minute.

A.3 Electrical Cycle

The choreographed mechanical cycle of the heart is driven by carefully orchestrated electrical signals. The spread of electrical signals through the heart is regulated

almost entirely by the speed of the signal through given portions of myocardium. While gap junctions between cells allow cell-to-cell transmission of electrical action potentials between any adjacent myocytes (Section A.4), transmission is much faster between specialised electrical system cells, which have many more such junctions. The cellular refractory period (Section A.4) induces ordered ‘waves’ of depolarisation across the myocardium and prevents retrograde conduction in the healthy heart. The typical sequence of depolarisation in the course of a single heartbeat is described below. In summary, the electrical signal spreads quickly through the specialised conducting system and slowly from that system through typical contractile myocardium.

The wave of depolarisation that will trigger a normal heartbeat begins in the sinoatrial (SA) node. The SA node, often called the ‘natural pacemaker’, is a specialised bundle of conduction system cells located in the superior right atrium (Figure A.7). While several loci of cardiac cells have pacemaker potential, the SA node has the highest natural rate (60 - 100 beats per minute), and therefore dominates, as other potential pacemakers are ‘reset’ when the wave of depolarisation initiated by the SA node reaches them. The pacemaker potential is described in more detail in Section A.4.

The wave of depolarisation spreads from the SA node through the atrial myocardium, triggering atrial systole. The electrical signal travels through normal atrial myocardium at about 0.3 m/s, and through specialised internodal pathways at about 1 m/s. These pathways organise the spread through the atria, and facilitate faster conduction to the atrioventricular (AV) node (Figure A.7). One pathway, known as the anterior interatrial band or Bachman’s bundle, passes to the left atrium along the anterior atrial walls. Transmission from the SA node to the AV node takes about 30 milliseconds (ms).

The electrical impulse cannot travel directly from the atrial to the ventricular myocardium, as the fibrous valve annuli do not conduct. Instead, the AV node, located in the posterior wall of the right ventricle behind the tricuspid valve (Figure A.7) provides the conduit from the atria to the ventricles. In order to




Figure A.7: Electrical conduction system of the heart. SA: sinoatrial node. AV: atrioventricular node. 1: internodal pathways. 2: anterior interatrial band (Bachmann's bundle). 3: bundle of His. 4: right bundle branch. 5: left bundle branch. 6: anteriosuperior left bundle branch. 7: posteroinferior left bundle branch. 8: Purkinje fibres.

allow the atria to contract fully before ventricular contraction begins, the AV node creates a delay in conduction. This delay is mediated by the small number of gap junctions within the node, which causes slower transmission of the action potential from cell to cell. It takes about 90 ms for an impulse to travel through the AV node to the fibrous tissue separating the atria from the ventricles, and a further 40 ms to transit this fibrous tissue via the fascicular AV bundle. Thus, the total delay at the AV node is 130 ms, and the electrical impulse reaches the ventricular tissue a total of 160 ms after its initiation at the SA node.

Upon exiting the atrioventricular bundle, the electrical impulse spreads rapidly through the His-Purkinje system. This system is made up of large fibres with many gap junctions, facilitating conduction at 1.5 - 4.0 m/s. From the AV node, the bundle of His travels 5 - 15 mm down the intraventricular septum before dividing into the left and right bundle branches (Figure A.7). The branches progress apically, dividing into smaller branches which course circumferentially around the cardiac

chambers before turning back basally. The branches terminate in Purkinje fibres, which merge directly with normal ventricular myocytes. Conduction from the AV node through the entire His-Purkinje system takes only 30 ms. After exiting the Purkinje fibres, it takes another 30 ms for the action potential to spread throughout the ventricular muscle.

Each segment of the electrical conduction system is necessary for regular, coordinated contraction of the heart. Failure of the SA node or excessive excitability of alternate groups of myocytes leads to an ‘ectopic pacemaker’, which is often incompatible with organised flow through a conduction system carefully tuned for action potentials originating at the SA node. Excessive inhibition of transmission through the AV node leads to ‘AV block’ and ineffective coordination of atrial and ventricular contraction. In the most serious cases, AV block can inhibit ventricular contraction entirely. So-called accessory pathways between the atria and the ventricles allow atrioventricular transmission without the inbuilt delay of the AV node. Such pathways can cause a variety of arrhythmias outlined in Section B.3. Blocks in the His-Purkinje system lead to uncoordinated activation of the ventricles. Such dyssynchrony results in less effective expulsion of blood from the ventricles and can cause adverse myocardial remodelling over time. Further description of these arrhythmias, with particular attention to those with catheter ablation treatments, can be found in Section B.3.

A.4 Cellular Electromechanical Coupling

Myocytes are striated muscle cells arranged in a latticework, with individual fibres dividing and recombining frequently (Figure A.8). Intercalated discs form the barriers between individual cells. These discs feature many gap junctions, which allow easy spread of ions between cells. Thus, the myocardium is characterised as a syncytium, in which an action potential spreads readily across tissue. The atrial and ventricular syncytia are separated by the non-conducting fibrous annuli of the mitral and tricuspid valves (Section A.3).




Figure A.8: Cellular structure of the heart. (a) Schematic of branching cells connected by intercalated discs. (b) Micrograph of cardiac muscle. (c) Desmosomes provide structural support and gap junctions allow ion flow through intercalated discs. Mitochondria and actin/myosin sarcomeres also illustrated. Adapted from OpenStax College under a permissive license^[307].

The typical myocyte at rest maintains a membrane potential of about -90 mV, maintained primarily by the sodium-potassium pump and potassium leak channels. The depolarisation cycle is typically described in five phases (Figure A.9, left). When an action potential arrives at a cell, depolarising it to a threshold of about -40 mV, fast sodium channels open (Phase 0). Influx of sodium into the cell causes rapid depolarisation to $+20$ mV. As fast sodium channels close, the cell begins to repolarise via flow through open potassium channels (Phase 1). This




Figure A.9: Membrane potential of the myocardiocyte. Left: typical ventricular myocardiocyte. Note rapid Phase 0. Right: pacemaker cell in the sinoatrial node. Note higher resting potential and upsloping Phase 4.

initial repolarisation is halted by the closing of fast potassium channels and the opening of slow, voltage-gated calcium channels which allow a calcium influx from the extracellular space and the sarcoplasmic reticulum (Phase 2). The plateau phase prolongs the refractory period of cardiac muscle. As the calcium channels close and slow potassium channels open, the cell rapidly repolarises (Phase 3), and the sodium-potassium pump maintains this resting potential (Phase 4).

Myocardiocytes with pacemaker potential – notably the SA node but also the AV node and Purkinje fibres – have a higher resting membrane potential, between -70 and -50 mV. These cells uniquely express hyperpolarisation-activated, cyclic nucleotide gated (HCN) non-selective cation channel. During Phase 4, HCN channels cause slow depolarisation (Figure A.9, right). This depolarisation itself encourages closure of potassium channels, opening of calcium channels, and further opening of the HCN channels. These changing ion flows together result in an accelerating membrane depolarisation, and spontaneous arrival at the action potential threshold. Acetylcholine, delivered via the parasympathetic vagus nerve, slows the pacemaker rate by closing HCN channels, increasing potassium permeability, and reducing calcium permeability. Conversely, catecholamines, delivered as noradrenaline (norepinephrine) via sympathetic nerves and adrenaline (epinephrine) via the adrenal medulla, increase pacemaker rate by opening HCN channels and increasing the

Phase 4 calcium current. Pacemaker cells also have much lower influence from calcium channels, and their Phase 1 and Phase 2 are less well defined.

Excitation-contraction coupling in contracting myocytes is governed by calcium flow. The influx of extracellular calcium from voltage-gated calcium channels in Phase 3 triggers the opening of ryanodine receptor calcium channels (RyR) in the sarcoplasmic reticulum, resulting in further increase in cytoplasmic calcium. This calcium binds to troponin, initiating contraction of the sarcomere. Cardiac contraction lasts about 250 ms, enabling the full mechanical cycle of blood expulsion. The electrical refractory period is similar, preventing the tetanic contraction possible in skeletal muscle.

You know how people say, ‘You can’t live without love’? Well, oxygen is even more important.

— David Shore’s Dr. Gregory House^[308]

B

Heart Disease and Percutaneous Treatment

Contents

B.1	Coronary Heart Disease	169
B.2	Valvular Heart Disease	169
B.3	Cardiac Arrhythmia	170

Heart disease is the leading cause of death in the United States and the United Kingdom. Coronary artery disease alone accounts for about one in five deaths in the United States^[309]. Hypertension affects one in three U.S. adults and is associated with left ventricular hypertrophy, which is itself a major predictor of heart failure and other disease^[31,33]. Statistics in the United Kingdom are similarly troubling, with about one in three British deaths in 2010 attributed to cardiovascular disease^[310].

In this appendix, we briefly review the major categories of heart disease with established percutaneous treatment options. This overview is meant to contextualise the contributions to procedural imaging described in the body of this thesis.

B.1 Coronary Heart Disease

In 2006, 6.5% of men and 4% of women in Britain suffered from coronary heart disease, including almost a quarter of those over 75 years old^[310]. Coronary heart disease (CHD) refers to an interruption of the normal perfusion of the myocardium by the coronary artery vasculature. Coronary perfusion is necessary for myocardial function and viability. As cardiomyocytes do not have significant energy stores, effective contraction will cease within seconds of cardiac ischaemia^[311]. If ischaemia persists, myocardial infarction (MI) results.

The classic progression to MI begins with coronary arterial endothelial dysfunction. Lipids accumulate in the arterial wall followed by the formation of an atheroma, a macrophage-rich accumulation of degenerative material in the tunica intima often referred to as a ‘plaque’. The atheroma progresses to a ‘complicated fibroatheroma’ as calcification and fibrosis partially obstructs coronary flow, causing clinical angina and cardiac remodeling. Plaque rupture is the catastrophic event, exposing the thrombogenic core of the fibroatheroma. This results in rapid local thrombosis and severe-to-complete blockage of the coronary artery. Less common mechanisms of MI include embolism, coronary vasospasm, and thrombosis of a non-atherosclerotic vessel^[311].

A large and growing number of clinical tools to treat MI and acute coronary syndromes, including thrombolytics, percutaneous arterial stenting and ballooning, and coronary artery bypass grafting (CABG), have dramatically increased survival of coronary heart disease (Section 2.3). However, these improvements bring into sharp relief the cardiac sequelae of myocardial infarction, including valvular disease^[312,313] and arrhythmias^[314,315].

B.2 Valvular Heart Disease

Valvular disease affects 2.5% of Americans, including more than 13% of those over 75 years old^[316]. Valvular heart disease describes a failure in the physiologic role of any cardiac valve. Of the four valves, the mitral and aortic valves, which

are exposed to the high pressures of the left heart, are most likely to fail. The two common failure modes are regurgitation, in which a valve cannot maintain competence against backflow, and stenosis, in which a valve produces high resistance to forward flow. Though post-streptococcal rheumatic fever was once the most common cause of valvular disease^[317], degeneration driven by long-term cardiac stressors now contributes to most valvular disease in the developed world.

Mitral regurgitation is present in about 1.7% of the US adult population, including 9.3% of those over 75 years old. This is closely followed in frequency by aortic stenosis and regurgitation, which are each present in about 0.5% of the population^[316]. Their prevalence is expected to increase as the population continues to age^[318]. Significant valvular dysfunction leads to heart failure as the left ventricle succumbs to pressure or volume overload. Although surgical valve replacement or repair is often the definitive treatment^[319], a number of percutaneous options exist for patients who are poor surgical candidates^[320,321]. These low-risk approaches are becoming more common and may eventually supplant open surgery as the preferred mode of treatment.

B.3 Cardiac Arrhythmia

Cardiac arrhythmias are malignant variations of the natural progression of electrical signalling in the heart. In the normal heart, an electrical impulse travels from the sinoatrial node, through the atria, to the atrioventricular node, along the His-Purkinje system, and finally through the ventricles (Section A.3). As each cardiomyocyte can both initiate and conduct action potentials, propagation of the cardiac electrical signal is orchestrated by a hierarchy of spontaneous pacemaker frequencies, conduction rates, and refractory periods. Arrhythmias occur when one of these factors is disrupted by structural malformation, local disease, or systemic electrochemical imbalance. In adults, this is often due to myocardial infarction or ischaemia, non-ischaemic cardiomyopathy, or the side effects of medication. A huge range of arrhythmias can occur. Here we present only the most iconic, and

then only in their most common or illustrative forms (Figure B.1). For a more detailed treatment, refer to de Luna's *Clinical Arrhythmology*^[322].

The most common arrhythmia in adults is atrial fibrillation (AF)^[323], described more fully in Section 2.4. In addition to causing palpitations and syncope, AF is implicated in ischaemic strokes and progression to heart failure^[323,324]. Over the past ten years, percutaneous ablation of AF pathways has gained favour over traditional pharmacologic rate control^[325,326]. Detailed mapping procedures now allow ablation attempts even of atypical AF patterns.

Atrial tachycardia and atrial flutter are also examples of the SA node being overwhelmed by atrial activity. The former is caused by an ectopic atrial pacemaker, and the latter by an organised reentrant rhythm within the atria. The cycle timing of atrial flutter is typically at about 300 per minute, and may present as a pulse of 150 beats per minute due to 2:1 conduction through the AV node^[327]. These arrhythmias are treated similarly to AF, with both medical and ablation options^[328,329].

Ventricular tachycardia (VT) is an acute arrhythmia associated with sudden cardiac death^[330]. At high ventricular rate, there is not time for ventricular filling between contractions, and cardiac output decreases. The mechanism of VT is typically either an excitatory focus in the ventricle or an area of slow conduction which permits a reentrant loop. Each of these lesions may be the result of myocardial scar or fibrotic tissue^[314,331,332], VT can be treated by antiarrhythmic medications or catheter ablation of arrhythmogenic foci or slow-conduction regions^[333]. Many patients require hybrid therapy, with ablation increasing the efficacy of medical therapy. In addition to this substrate modification, patients often benefit from implantable cardioverter/defibrillators (ICDs), which continually monitor the patient's electrocardiogram and deliver a shock to halt VT when it is detected^[334,335].

If ventricular tachycardia becomes disorganised, the resulting rhythm is ventricular fibrillation (VF). This rhythm is incompatible with life, and must be immediately defibrillated. Like those at risk of VT, patients at risk of VF may be implanted with internal defibrillators^[334,335].




Figure B.1: Overview of some classic arrhythmias. The wide variation between and within classes of arrhythmia is simplified by necessity. Anatomic positions are approximate and representative. Heart figure adapted from Gray^[50].

Supraventricular tachycardia (SVT) refers to any tachycardia arising above the bundle of His. While this technically includes the atrial-induced tachycardias discussed above, the term SVT is more commonly used to refer to tachycardia caused by reentrant circuits.

Atrioventricular reentry tachycardia (AVRT) is an SVT which involves an accessory pathway between the atria and the ventricles, allowing the AV node to be bypassed. This pathway completes to a loop in which an electrical impulse can travel from the atria, through the AV node to the ventricles, and then back through the accessory pathway to the atria ('orthodromic', 90 - 95% of cases) or in the opposite direction ('antidromic'). The existence and location of the pathway varies. The most common example is the Bundle of Kent in Wolff-Parkinson-White syndrome. In patients with these pathways, AVRT is initiated by premature atrial or ventricular beats. In patients with AF as well as an accessory pathway, the fibrillation of the atria is not buffered by the high refractory period of the AV node, and rapid ventricular response and degeneration to ventricular fibrillation may result^[336-338]. AVRT may be corrected by catheter ablation of the accessory pathway^[339].

Atrioventricular nodal reentry tachycardia (AVNRT) is another reentrant SVT. In this arrhythmia, the AV node itself contains two pathways, typically a fast pathway and a slow pathway. Patients with this physiology will often remain in sinus rhythm (conducted via the fast pathway) until a premature atrial or ventricular impulse triggers AVNRT. Typically, impulses are carried from the atria to the ventricles by the slow pathway, and then re-enter the AV node via the fast pathway to pass back to the atria, completing the reentrant loop ('slow-fast AVNRT')^[340]. The opposite conduction pattern ('fast-slow') is seen in 10% of patients, and a 'slow-slow' is also possible. While AVNRT is often managed medically, in refractory cases ablation of the AV node – typically targeting the slow conduction pathway – is considered^[341,342].

Bradyarrhythmias, in which the heart rate is abnormally low, may be caused by failure of the SA node or a failure of the action potential to reach the myocardium and trigger effective contraction. SA node failure may be a result of infarction,

medication, or age-related deterioration. When not correctable with medication, bradyarrhythmias are treated with the implantation of artificial pacemaker devices^[343]. These devices trigger contraction of the heart by delivering exogenous sparks via leads embedded in the myocardium.

Conduction blocks, most commonly at the AV node, may also cause bradycardia. While low-level blocks, which feature prolonged conduction delay at the AV node, do not warrant treatment, high-level blocks prevent beats from reaching the ventricles, and must be treated with a pacemaker^[343]. Blocks within the His-Purkinje system, including bundle branch and fascicular blocks, cause dyssynchronous activation of the left and right ventricles, decreasing efficiency of expulsion of blood from these chambers. Resynchronisation can be achieved using advanced pacemakers which simultaneously activate the chambers as required, though this treatment is currently only considered clinically indicated in 5% - 10% of patients with dyssynchrony^[344].

*Sollten nun die neuen Strahlen nicht longitudinalen
Schwingungen im Aether zuzuschreiben sein?*

*Should not the new rays be ascribed to longitudinal
waves in the ether?*

— Wilhelm Conrad Röntgen^[345]



Machine Learning and Signal Processing by Example

Contents

C.1 Random Forests	175
C.1.1 Data	176
C.1.2 Training	177
C.1.3 Testing	179
C.1.4 Conclusion	180
C.2 Principal Component Analysis	181
C.2.1 Data	181
C.2.2 Transformation	184
C.2.3 Results and Insight	185
C.2.4 Conclusion	188

In this appendix, we illustrate two of the more complex data analysis algorithms used in the thesis – random forests and principal component analysis – using a step-by-step, example-based approach.

C.1 Random Forests

Random forests (RF) is a machine learning algorithm used to make robust predictions from a large amount of input data. At its heart, RF constructs an ensemble of ‘weak learners’ – individual decision trees – each of which has low predictive value

individually, but which in concert achieve much higher accuracy. Like all machine learning algorithms, RF does not require an understanding of how – or even if – any individual feature in the input data affects the outcome. The approach is described by Breiman *et al.*^[220] and is known for its accuracy, robustness to noise, and computational efficiency. Although training a random forest is computationally intensive, making a prediction is a very speedy process. Unless otherwise cited, all information in this section is supported by Breiman^[220]. A more formal discussion of the advantages of RF is found in Section 4.1.3.

RF can be implemented as a classifier, wherein the model uses input data to predict a categorical output – the presence or absence of a disease, a discrete stage of illness, etc. – or as a regressor, which predicts a continuous variable – blood pressure, risk of stroke, probabilistic five-year mortality, etc. In either case, building an RF model begins with a set of training data – paired input data, known as ‘feature vectors’, and their associated, known gold-standard output data, called ‘responses’.

C.1.1 Data

To step through a typical RF training and testing algorithm, consider the patient data listed in Table C.1. This data is a small subset of information collected at the Cleveland Clinic (Cleveland, Ohio, USA) by Dr. Robert Detrano^[346] and now freely available at the University of California, Irvine Machine Learning Repository^[347].

For the ten patients we consider, the feature vector is made up of age, sex, systolic blood pressure at rest, serum cholesterol, serum glucose, and presence of left ventricular hypertrophy by electrocardiogram. The output, or ‘response’, is the presence or absence of at least 50% artery blockage in any coronary artery, as evaluated by fluoroscopy.

It is important to note that although we have chosen to use ten patients, each with a feature vector of six parameters, random forests’ strength is in analysis of extremely large data sets. The advantage of using only ten patients is that we can work through the training and evaluation process step-by-step, but we should not expect impressive results on this small sample.

Patient	Age	Sex	SBP	Chol	BSG	ECG	Stenosis
A	63	M	145	233	↑	LVH	-
B	67	M	160	286	-	LVH	+
C	67	M	120	229	-	LVH	+
D	37	M	130	250	-	-	-
E	41	F	130	204	-	LVH	-
F	56	M	120	236	-	-	-
G	62	F	140	268	-	LVH	+
H	57	F	120	354	-	-	-
I	63	M	130	254	-	LVH	+
J	53	M	140	203	↑	LVH	+

Table C.1: Ten heart disease patients from the Cleveland Clinic^[346,347]. Data includes age (years), sex, systolic blood pressure at rest (SBP), serum cholesterol (Chol), serum blood sugar (BSG; ↑ indicates > 120 mg/dL), and presence of left ventricular hypertrophy (LVH) by electrocardiogram (ECG) criteria. Finally, stenosis represents presence of > 50% artery blockage of any artery by coronary angiogram.

C.1.2 Training

Ordinarily, the data set is randomly divided into training and testing sets. This prevents ‘overfitting’, a pitfall of machine learning evaluation. Overfitting describes the situation in which coincidental relationships between observed feature vectors and responses (e.g. every red-haired, normotensive woman in the dataset happens to be disease-positive) can be spuriously recognised and modelled. This error is somewhat mitigated by splitting the known values into separate training and testing data sets, so that any overfitting on the training set does not produce artificially impressive results when evaluated using the disparate testing set. Here, we will use patients A, B, C, D, E, and F for training, and reserve G, H, I, and J for testing.

The crux of training an RF is generating a large number of decision trees. These trees are known as ‘weak learners’, as they individually have low classification accuracy. The key observation of RF is that a large number of weak learners may be combined to form a strong learner. Most RF models are made up of hundreds or thousands of trees. For the sake of illustration, we will construct only three here.

As its name implies, the random forests method relies on randomisation of the training data in order to generate diverse trees. This randomisation is achieved




Figure C.1: Detailed schematic of training a random forest tree. Red circles signify patients with coronary stenosis; white circles represent patients without stenosis. Example data from Table C.1.

in two ways: bootstrap aggregation (‘bagging’) on a per-tree basis, and selection of feature subsets on a per-node basis. Bagging means randomly selecting, with replacement, a same-size subset from the training data. For example: for our first tree, from the training set A, B, C, D, E, F, we randomly select F, A, C, B, C, D (Figure C.1, box 1). This selection of patients will be used when constructing every node in the first tree.

To construct a node, we begin by randomly selecting a subset of features to consider (Figure C.1, box 2). In Breiman’s original paper, the size of the subset of features is $\log_2 f$, where f is the total number of features. Most modern implementations use $\sqrt{f}$. For values $f \leq 36$, the choice makes little difference.

After a subset of features is selected, the data points (i.e. patients) up for consideration at this node are categorised by their response (i.e. stenosis classification). The single feature and, in the case of continuous features, threshold value which best discriminates between responses is chosen (Figure C.1, box 3). A variety of algorithms exist to make this selection – Gini impurity is most commonly used in random forest, though variance minimisation and information maximisation are other options^[348]. For this illustration, the important concept is that the result is a simple binary split which best separates data points with different responses.

The process of training a node – beginning with random feature selection – is repeated until each leaf node is ‘pure’ (Figure C.1). That is, each end node contains only data points with a single response value. The tree growing process is then repeated, beginning with the bootstrap aggregation process, to generate as many trees as needed.

C.1.3 Testing

After all trees in the random forest have been trained using their individual bootstrap samples (Figure C.2), we move to evaluation of the model. This process is extremely straightforward: The feature vector for each of the data points (i.e. patients) in the testing set – G, H, I, and J – is classified using each tree in the model. When patient G is classified with the first tree, she initially goes down the right branch




Figure C.2: Three-tree random forest model for example data. Leftmost tree was derived in Figure C.1.

(LVH on ECG), and then the left (normal glucose), thereby being classified as stenosis-positive. The second tree classifies her as stenosis-negative, as her age is less than 65. The third and final tree classifies her as stenosis-positive (LVH-positive, high-cholesterol). Because two out of the three trees classify her as stenosis-positive, and she is in fact stenosis-positive, this is regarded as a successful classification.

When a similar procedure is applied to the other patients, we find that patient H is successfully classified negative (by all trees), patient I is successfully classified positive (by the first and third trees), and patient J is a false negative (as classified by all trees). Therefore, we would evaluate this model as having 75% accuracy. Of course, the extremely small data sets used in this example make drawing any reasonable conclusions from this data impossible.

Note that when many trees are used, the number of votes required to classify a data point as positive can be varied to tune the random forest model, adjusting sensitivity and specificity. This facilitates generating a receiver operating characteristic (ROC) curve (Section 4.1.4, page 54), which provides a more meaningful evaluation of the classifier than accuracy alone.

C.1.4 Conclusion

Random forest is a versatile machine learning algorithm that uses an ensemble of weak learners (individual trees) to generate a strong learner. It is ideal when dealing

with large data sets and is robust to noisy feature vectors. Though training is computationally intensive on large amounts of data, classification by binary decision tree is very quick once the model is built. Training only needs to be undertaken once, and retention of the training set is not essential once the classifier is constructed – critical for embedded devices and other limited-memory applications. While random forests have a tendency overfit noisy data, this can be ameliorated by randomly selecting disparate training and testing data sets. Our use of random forest in this thesis is illustrated in Figure 4.1, page 53, and performs strongly in predicting atrial bipolar voltage using feature vectors with 36 image parameters.

C.2 Principal Component Analysis

Principal component analysis (PCA) is a data analysis technique which isolates linearly independent components from a set of (possibly) correlated features^[225]. The components extracted are ordered by the amount of total variance explained, assuming the features have dimensions of a meaningful scale. More simply, PCA is able to take highly cross-correlated data and re-frame it in terms of independent components. This is equivalent to an introduction of orthogonal axes on the feature space, isolating the independent contributions of combination of features. A more formal description of PCA is found in Section 4.1.5. Unless otherwise cited, all information in this section is supported by Jolliffe’s definitive account^[225].

C.2.1 Data

To explore a typical use of PCA, we consider the women’s 400 metre dash at the 2012 Summer Olympics. The competitors, their heights, weights, and calculated average speed in their highest-level race (i.e. finals, semi-finals, or last heat reached), are listed in Table C.2. As a data set of convenience, this was extracted from the collaborative *Wikipedia*^[349] and excludes eight sprinters whose height or weight was not reported there. It also excludes runners who did not start, did not finish, or were disqualified.

Name	Height (m)	Weight (kg)	Speed (m/s)
Sanya Richards-Ross	1.73	62	8.07
Christine Ohuruogu	1.75	68	8.05
DeeDee Trotter	1.78	64	8.05
Novlene Williams-Mills	1.68	55	7.98
Antonia Krivoschapka	1.68	59	7.97
Francena McCorory	1.70	60	7.95
Rosemarie Whyte	1.80	64	7.88
Christine Day	1.68	51	7.81
Alina Lohvynenko	1.80	68	7.79
Omolara Omotosho	1.52	50	7.78
Nataliya Pyhyda	1.70	60	7.78
Yuliya Gushchina	1.75	62	7.74
Joanne Cuddihy	1.84	65	7.71
Lee McConnell	1.77	64	7.66
Pınar Saka	1.63	53	7.64
Sviatlana Usovich	1.65	54	7.63
Raysa Sánchez	1.78	61	7.62
Marlena Wesh	1.68	64	7.62
Tjipekapora Herunga	1.67	51	7.61
Aliann Pompey	1.68	55	7.61
Shana Cox	1.72	65	7.61
Joelma Sousa	1.74	51	7.59
Jenna Martin	1.73	65	7.57
Bianca Răzor	1.69	54	7.57
Moa Hjelmer	1.72	60	7.57
Joy Nakhumicha Sakari	1.70	60	7.55
Amy Mbacke Thiam	1.85	80	7.51
Geisa Aparecida Coutinho	1.61	55	7.49
Olesea Cojuhari	1.70	58	7.46
Marina Maslenko	1.64	57	7.45
Aauri Bokesa	1.83	71	7.45
Phumlile Ndzinisa	1.65	52	7.41
Ambwene Simukonda	1.70	70	7.38
Danielle Alakija	1.86	71	7.05
Graciela Martins	1.66	60	6.86
Maziah Mahusin	1.63	51	6.75
Hristina Risteska	1.68	60	6.57

Table C.2: Women’s 400 metre dash competitors in the 2012 Summer Olympics. Speeds calculated from highest round. Excludes competitors without a finish time, or whose height or weight is not listed in *Wikipedia*^[349].




Figure C.3: Women’s 400 metre dash competitor height, weight, and speed. Pairwise least-squares linear regression models (red) with R^2 values reported. Dashed lines indicate low R^2 .

Examining the data with pairwise least-squares linear regression models (Figure C.3), we find a high degree of correlation between athletes’ heights and weights. There is very little correlation between speed and either of the other parameters. One might reasonably conclude that although we have three data points describing each athlete (height, weight, and speed), we would be able to describe an athlete almost as well with only two (speed and either height or weight) given the strong correlation between height and weight.

PCA is a statistical method which not only formalises this intuitive result, but also provides the new parameters which capture the most variance from the data and quantifies the amount of variance contained within each. Because each of our three observed parameters has a different unit of measurement, we begin by normalising each feature. A variety of normalisation strategies are possible, but we choose the z-score here for its simplicity. The z-score represents the number of standard deviations from the mean for any value: $z_x = (x - \bar{x})/\sigma_x$. After normalisation, height values range from -2.70 to 2.05, weights range from -1.51 to 2.90, and speeds vary between -3.06 and 1.45. This linear transformation to unitless values allows easier comparison of variance. Where the original data is numerically comparable from the outset, normalisation is unnecessary, though shifting the values to a zero mean is required.

Name	Height		Weight		Speed	
	(m)	z	(kg)	z	(m/s)	z
Rosemarie Whyte	1.80	1.22	64	0.55	7.88	0.87
Christine Day	1.68	-0.46	51	-1.36	7.81	0.66
Tjipekapora Herunga	1.67	-0.60	51	-1.36	7.61	0.06
Moa Hjelder	1.72	0.10	60	-0.04	7.57	-0.06

Table C.3: Four women’s 400 metre dash competitors. Heights, weights, and speeds are given as z-scores in addition to their raw values.

C.2.2 Transformation

As described mathematically in Section 4.1.5 (page 55), PCA is a transformation which shifts the data axes such that the most global variation is represented on the fewest (orthogonal) axes. Each axis is defined by the unit loading vector $\mathbf{w}_{(k)}$, which lays out the weight of each original parameter making up the (new) principal component. In our case, there are three original parameters (height, weight, and speed), so there will be three principal components ($k = 1, 2, 3$).

While it is too cumbersome to do this matrix multiplication by hand for the entire data set, let’s briefly consider just the four sprinters and their parameters (by z-score) in Table C.3. Referring back to Equation 4.2 (page 55), our task is to maximise the magnitude of the matrix product

$$\begin{bmatrix} 1.22 & 0.55 & 0.87 \\ -0.46 & -1.36 & 0.66 \\ -0.60 & -1.36 & 0.06 \\ 0.10 & -0.04 & -0.06 \end{bmatrix} \begin{bmatrix} w_1 \\ w_2 \\ w_3 \end{bmatrix} \quad (\text{C.1})$$

under the constraint that $w_1^2 + w_2^2 + w_3^2 = 1$. Refer back to the z-values in Table C.3 to derive the composition of this data matrix.

Before solving for the maximum product directly, consider a few example choices for the loading vector $\mathbf{w}_{(1)} = (w_1, w_2, w_3)$. The vector $(1, 0, 0)$ represents a principal component equivalent to the runner height. Using this vector, the magnitude of (C.1) is $\sqrt{1.22^2 + (-0.46)^2 + (-0.60)^2 + 0.10^2} = 1.44$. When we choose values $(0, 1, 0)$, i.e. only the weight, the magnitude is 2.00. Choosing only the speed,

the magnitude is 1.10. These magnitudes represent the variance of the height, weight, and speed for this subset of data.

If we consider a loading vector which balances the height and weight equally, $(w_1, w_2, w_3) = (1/\sqrt{2}, 1/\sqrt{2}, 0)$, the magnitude is

$$\sqrt{\left(\frac{1.22}{\sqrt{2}} + \frac{0.55}{\sqrt{2}}\right)^2 + \left(\frac{-0.46}{\sqrt{2}} + \frac{-1.36}{\sqrt{2}}\right)^2 + \left(\frac{-0.60}{\sqrt{2}} + \frac{-1.36}{\sqrt{2}}\right)^2 + \left(\frac{0.10}{\sqrt{2}} + \frac{-0.04}{\sqrt{2}}\right)^2} = 2.27. \quad (\text{C.2})$$

This value represents the high covariance between height and weight.

Finally, solving for the maximum magnitude of the matrix product in (C.1), we find $(w_1, w_2, w_3) = (0.646, 0.747, 0.157)$ yields a maximum of 2.28. This $\mathbf{w}_{(1)}$ is the loading vector for the first principal component.

To find the second principal component, following Equation 4.3, we first subtract the contribution of the first principal component from the data. This leaves the remainder left to be fitted,

$$\hat{\mathbf{X}}_k = \mathbf{X} - \sum_{s=1}^{k-1} \mathbf{X} \mathbf{w}_{(s)} \mathbf{w}_{(s)}^T \quad (\text{C.3})$$

$$= \begin{bmatrix} 1.22 & 0.55 & 0.87 \\ -0.46 & -1.36 & 0.66 \\ -0.60 & -1.36 & 0.06 \\ 0.10 & -0.04 & -0.06 \end{bmatrix} - \begin{bmatrix} 0.86 & 1.00 & 0.21 \\ -0.78 & -0.90 & -0.19 \\ -0.90 & -1.04 & -0.22 \\ 0.02 & 0.02 & 0.00 \end{bmatrix} \quad (\text{C.4})$$

$$= \begin{bmatrix} 0.36 & -0.45 & 0.66 \\ 0.32 & -0.46 & 0.85 \\ 0.30 & -0.32 & 0.28 \\ 0.08 & -0.06 & -0.06 \end{bmatrix}, \quad (\text{C.5})$$

which clearly contains less residual variance in height and weight (first and second columns) compared to speed (last column). The process, beginning with maximising (C.1), is repeated to find the second and third components.

C.2.3 Results and Insight

Thankfully, modern computational technology rescues us from the need to continue matrix arithmetic by hand. Returning again to the full sprinter data set (Table C.2),




Figure C.4: PCA component definitions. Each graph represents the contributions of height (H), weight (W), and speed (S) to the given PCA component.

we use an open-source statistics package (`scikit-learn`^[229]) to compute the loading vectors,

$$\begin{aligned}
 \mathbf{w}_{(1)} &= (\quad 0.704, \quad 0.699, \quad 0.140 \quad), \\
 \mathbf{w}_{(2)} &= (\quad -0.038, \quad -0.161, \quad 0.986 \quad), \text{ and} \\
 \mathbf{w}_{(3)} &= (\quad -0.709, \quad 0.700, \quad 0.087 \quad),
 \end{aligned} \tag{C.6}$$

which are represented graphically in Figure C.4.

The first component (PCA 1), which captures 59.8% of the total variance in the original data, approximately corresponds to an athlete’s size, with comparable contributions from height and weight. This component does also contain a small contribution from speed, indicating that there was covariance between size and speed in this data set. The second component, capturing a further 33.0% of variance, represents speed. The final component, capturing the remaining 7.2% of variance, represents the ‘other direction’ on the height-weight plot: weight not explained by correlation with height.

Now that we have PCA component definitions, there are a variety of things we can do with them. One of the most common applications is feature reduction: DeeDee Trotter is 1.78 m tall, weighs 64 kg, and runs 8.05 m/s. From only her first and second PCA components (1.236, 1.243), we can estimate her height at 1.77 m, her weight at 64.8 kg, and her speed at 8.06 m/s. This very near estimate comes from only two parameters, compared to the three originally presented. From a

computational efficiency perspective, PCA in this instance would allow a researcher to eliminate the storage and processing costs associated with 33% of the data while losing only 7.2% of the total data variance.

Another important application of PCA is preventing misattributed statistical relationships. If we were studying, for example, stride length of these runners, it would be improper to come to the two conclusions ‘taller athletes have longer strides’ and ‘heavier athletes have longer strides’. The covariance between height and weight makes these statements equivalent. By applying PCA first, a researcher would conclude singly that ‘PCA component 1 correlates with stride length’, and then look at the definition of component 1 to find deeper meaning.

Finally, note that the most statistically significant PCA components are not always the most important. In a study on runners’ metabolism, for example, the last PCA component, which represents weight not attributable to height, might be very interesting.

As a further example, consider a different (hypothetical) data set: the height, weight, and four limb lengths of a large population. Intuitively, the first PCA component is likely to correlate with ‘size’, and contain contributions from each parameter. Another high-order component would probably be weight not explained by the other ‘size’ components. Still more components would represent long limbs globally and individual limb differences. Each of these PCA components might be interesting to different researchers. A bioanthropologist studying the availability of nutrition over centuries might be most interested in overall ‘size’ as well as unattributed weight. A public health researcher focusing on diabetes might be interested in unattributed weight alone. A physician studying Marfan syndrome – classically associated with arachnodactyly, i.e. long limbs – would focus on that component. And an orthotist would find value in the components indicating deviations in individual limb lengths, even though those components are probably the least significant in terms of population variance explained.

C.2.4 Conclusion

PCA is a versatile method for re-orientating observational data into linearly uncorrelated features. The method reveals internal structure in multidimensional data in the way that best explains variance, and requires no prior knowledge of those relationships. It may be used for parameter reduction or to focus on the contribution of a variable that is uncorrelated with others.

The most significant limitation of PCA is its restriction to linear relationships. Given a series of height and weight data, for example, it might be more useful to use non-linear values (such as body mass index) rather than the linear combinations PCA supplies. Although alternatives exist, PCA remains a very popular choice for data transformation due to its mathematical parsimony and the resistance to bias granted by its unsupervised nature.

A man provided with paper, pencil, and rubber, and subject to strict discipline, is in effect a universal machine.

— Alan Turing^[350]

D

Supplemental Data for Chapter 4

This appendix contains supplemental data tables and figures representing the full results of single-parameter regression models and preliminary tuning tests for random forest models in Chapter 4.




Figure D.1: For each normalisation strategy, area under median random forest MR feature classifier ROC curve for bipolar voltage fibrosis thresholds.

Feature / Diameter		IIR Method			Utah Method		
		R^2	p	AUC	R^2	p	AUC
Mean	1 mm	0.011	10^{-4}	0.607	0.004	0.012	0.586
	2 mm	0.010	10^{-4}	0.610	0.004	0.016	0.586
	3 mm	0.010	10^{-4}	0.609	0.003	0.029	0.583
	4 mm	0.010	10^{-4}	0.606	0.003	0.033	0.581
	5 mm	0.009	10^{-3}	0.606	0.002	0.045	0.577
	6 mm	0.009	10^{-3}	0.609	0.002	0.058	0.583
	7 mm	0.009	10^{-3}	0.616	0.002	0.059	0.590
	8 mm	0.008	10^{-3}	0.614	0.002	0.070	0.587
	9 mm	0.008	10^{-3}	0.613	0.002	0.102	0.584
	10 mm	0.008	10^{-3}	0.610	0.002	0.110	0.587
Minimum	1 mm	0.011	10^{-4}	0.613	0.004	0.012	0.590
	2 mm	0.013	10^{-5}	0.603	0.005	0.006	0.577
	3 mm	0.010	10^{-4}	0.589	0.003	0.036	0.566
	4 mm	0.009	10^{-3}	0.593	0.002	0.057	0.567
	5 mm	0.010	10^{-4}	0.597	0.002	0.045	0.571
	6 mm	0.009	10^{-3}	0.591	0.002	0.065	0.563
	7 mm	0.009	10^{-4}	0.593	0.002	0.063	0.565
	8 mm	0.011	10^{-4}	0.587	0.003	0.043	0.560
	9 mm	0.009	10^{-4}	0.587	0.001	0.119	0.548
	10 mm	0.010	10^{-4}	0.577	0.001	0.136	0.560
Maximum	1 mm	0.011	10^{-4}	0.614	0.004	0.012	0.594
	2 mm	0.007	10^{-3}	0.602	0.002	0.056	0.586
	3 mm	0.008	10^{-3}	0.607	0.002	0.049	0.589
	4 mm	0.008	10^{-3}	0.618	0.003	0.036	0.595
	5 mm	0.007	10^{-3}	0.606	0.002	0.064	0.586
	6 mm	0.006	0.002	0.604	0.002	0.117	0.590
	7 mm	0.005	0.004	0.613	0.001	0.154	0.590
	8 mm	0.004	0.010	0.617	0.001	0.243	0.593
	9 mm	0.003	0.034	0.605	0.000	0.432	0.580
	10 mm	0.003	0.027	0.606	0.000	0.375	0.586
Varition	3 mm	0.000	0.605	0.551	0.000	0.534	0.565
	4 mm	0.001	0.342	0.547	0.001	0.284	0.562
	5 mm	0.000	0.489	0.546	0.001	0.363	0.560
	6 mm	0.000	0.679	0.553	0.000	0.456	0.563
	7 mm	0.000	0.767	0.557	0.000	0.484	0.574
	8 mm	0.000	0.525	0.555	0.000	0.832	0.569
	9 mm	0.001	0.207	0.541	0.000	0.544	0.562
	10 mm	0.001	0.263	0.549	0.000	0.674	0.564

Table D.1: Results of single-parameter linear regression on MR image data ($BV_p = a_1 \cdot \hat{v}_p + a_0$). See Section 4.4.2.

Feature / Diameter		IIR Method			Utah Method		
		R^2	p	AUC	R^2	p	AUC
Mean	1 mm	0.031	10^{-12}	0.609	0.016	10^{-6}	0.595
	2 mm	0.031	10^{-12}	0.610	0.015	10^{-6}	0.583
	3 mm	0.031	10^{-12}	0.608	0.014	10^{-5}	0.583
	4 mm	0.031	10^{-12}	0.608	0.014	10^{-5}	0.581
	5 mm	0.031	10^{-11}	0.603	0.013	10^{-5}	0.584
	6 mm	0.030	10^{-11}	0.610	0.013	10^{-5}	0.589
	7 mm	0.031	10^{-12}	0.611	0.013	10^{-5}	0.588
	8 mm	0.033	10^{-12}	0.620	0.015	10^{-6}	0.584
	9 mm	0.031	10^{-11}	0.607	0.013	10^{-5}	0.584
	10 mm	0.030	10^{-11}	0.613	0.013	10^{-5}	0.583
Minimum	1 mm	0.031	10^{-12}	0.613	0.016	10^{-6}	0.591
	2 mm	0.034	10^{-12}	0.596	0.015	10^{-6}	0.579
	3 mm	0.027	10^{-10}	0.589	0.010	10^{-4}	0.567
	4 mm	0.025	10^{-9}	0.595	0.009	10^{-3}	0.564
	5 mm	0.028	10^{-10}	0.594	0.010	10^{-4}	0.569
	6 mm	0.026	10^{-10}	0.594	0.009	10^{-3}	0.564
	7 mm	0.026	10^{-10}	0.594	0.009	10^{-3}	0.561
	8 mm	0.029	10^{-11}	0.594	0.009	10^{-3}	0.560
	9 mm	0.025	10^{-9}	0.588	0.006	0.001	0.552
	10 mm	0.025	10^{-9}	0.583	0.006	0.002	0.552
Maximum	1 mm	0.031	10^{-12}	0.613	0.016	10^{-6}	0.590
	2 mm	0.027	10^{-10}	0.604	0.014	10^{-5}	0.587
	3 mm	0.030	10^{-11}	0.616	0.015	10^{-6}	0.589
	4 mm	0.031	10^{-11}	0.615	0.016	10^{-6}	0.588
	5 mm	0.028	10^{-10}	0.606	0.014	10^{-5}	0.585
	6 mm	0.027	10^{-10}	0.606	0.013	10^{-5}	0.581
	7 mm	0.027	10^{-10}	0.616	0.015	10^{-5}	0.592
	8 mm	0.024	10^{-9}	0.621	0.013	10^{-5}	0.585
	9 mm	0.021	10^{-8}	0.608	0.011	10^{-4}	0.588
	10 mm	0.022	10^{-8}	0.607	0.011	10^{-4}	0.588
Varition	3 mm	0.005	0.004	0.548	0.006	0.002	0.565
	4 mm	0.006	0.002	0.548	0.007	10^{-3}	0.565
	5 mm	0.004	0.009	0.540	0.005	0.003	0.557
	6 mm	0.004	0.013	0.548	0.005	0.005	0.566
	7 mm	0.004	0.011	0.557	0.006	0.002	0.572
	8 mm	0.001	0.155	0.555	0.003	0.029	0.572
	9 mm	0.000	0.501	0.539	0.002	0.101	0.561
	10 mm	0.000	0.422	0.545	0.002	0.082	0.557

Table D.2: Results of single-parameter exponential regression on MR image data ($\log(BV_p) = a_1 \cdot \hat{v}_p + a_0$). See Section 4.4.2.

Feature / Diameter		IIR Method			Utah Method		
		R^2	p	AUC	R^2	p	AUC
Mean	1 mm	0.004	10^{-4}	0.527	0.003	10^{-3}	0.527
	2 mm	0.007	10^{-6}	0.540	0.005	10^{-5}	0.545
	3 mm	0.007	10^{-7}	0.549	0.005	10^{-5}	0.548
	4 mm	0.007	10^{-7}	0.551	0.005	10^{-5}	0.542
	5 mm	0.009	10^{-8}	0.559	0.007	10^{-7}	0.560
	6 mm	0.009	10^{-9}	0.560	0.008	10^{-8}	0.563
	7 mm	0.009	10^{-9}	0.560	0.008	10^{-8}	0.563
	8 mm	0.009	10^{-9}	0.561	0.008	10^{-8}	0.568
	9 mm	0.009	10^{-9}	0.561	0.008	10^{-8}	0.567
	10 mm	0.009	10^{-9}	0.561	0.008	10^{-8}	0.568
Minimum	1 mm	0.004	10^{-4}	0.530	0.003	10^{-3}	0.531
	2 mm	0.011	10^{-11}	0.560	0.008	10^{-8}	0.552
	3 mm	0.014	10^{-14}	0.562	0.009	10^{-8}	0.547
	4 mm	0.014	10^{-14}	0.569	0.008	10^{-8}	0.548
	5 mm	0.017	10^{-16}	0.585	0.010	10^{-10}	0.561
	6 mm	0.017	10^{-17}	0.592	0.011	10^{-10}	0.571
	7 mm	0.016	10^{-15}	0.596	0.009	10^{-9}	0.576
	8 mm	0.014	10^{-14}	0.597	0.008	10^{-8}	0.577
	9 mm	0.013	10^{-13}	0.602	0.007	10^{-7}	0.583
	10 mm	0.012	10^{-12}	0.604	0.006	10^{-6}	0.584
Maximum	1 mm	0.004	10^{-4}	0.524	0.003	10^{-3}	0.525
	2 mm	0.002	0.002	0.530	0.002	0.010	0.524
	3 mm	0.000	0.455	0.519	0.000	0.837	0.519
	4 mm	0.000	0.562	0.503	0.001	0.092	0.509
	5 mm	0.001	0.122	0.500	0.002	0.003	0.522
	6 mm	0.001	0.019	0.510	0.004	10^{-4}	0.529
	7 mm	0.002	0.002	0.521	0.006	10^{-6}	0.543
	8 mm	0.003	10^{-3}	0.525	0.006	10^{-6}	0.544
	9 mm	0.003	10^{-3}	0.527	0.007	10^{-6}	0.548
	10 mm	0.004	10^{-4}	0.530	0.008	10^{-8}	0.548
Varition	3 mm	0.006	10^{-6}	0.527	0.003	10^{-3}	0.513
	4 mm	0.007	10^{-7}	0.536	0.004	10^{-4}	0.519
	5 mm	0.010	10^{-10}	0.555	0.006	10^{-6}	0.533
	6 mm	0.011	10^{-10}	0.555	0.007	10^{-7}	0.544
	7 mm	0.010	10^{-10}	0.568	0.007	10^{-6}	0.556
	8 mm	0.009	10^{-9}	0.566	0.006	10^{-6}	0.554
	9 mm	0.009	10^{-8}	0.569	0.006	10^{-6}	0.557
	10 mm	0.008	10^{-8}	0.569	0.005	10^{-5}	0.557

Table D.3: Results of single-parameter linear regression on CT image data ($BV_p = a_1 \cdot \hat{v}_p + a_0$). See Section 4.5.2.

Feature / Diameter		IIR Method			Utah Method		
		R^2	p	AUC	R^2	p	AUC
Mean	1 mm	0.009	10^{-9}	0.526	0.006	10^{-6}	0.526
	2 mm	0.016	10^{-16}	0.545	0.011	10^{-11}	0.545
	3 mm	0.020	10^{-20}	0.549	0.014	10^{-14}	0.553
	4 mm	0.021	10^{-21}	0.551	0.015	10^{-14}	0.547
	5 mm	0.026	10^{-25}	0.563	0.019	10^{-19}	0.562
	6 mm	0.027	10^{-26}	0.556	0.021	10^{-21}	0.566
	7 mm	0.028	10^{-27}	0.558	0.023	10^{-22}	0.568
	8 mm	0.030	10^{-29}	0.561	0.024	10^{-24}	0.568
	9 mm	0.031	10^{-30}	0.559	0.026	10^{-25}	0.570
	10 mm	0.031	10^{-30}	0.560	0.026	10^{-25}	0.567
Minimum	1 mm	0.011	10^{-11}	0.533	0.007	10^{-7}	0.531
	2 mm	0.026	10^{-25}	0.557	0.017	10^{-17}	0.552
	3 mm	0.032	10^{-31}	0.563	0.018	10^{-17}	0.548
	4 mm	0.033	10^{-32}	0.570	0.017	10^{-17}	0.546
	5 mm	0.041	10^{-40}	0.585	0.023	10^{-22}	0.564
	6 mm	0.043	10^{-42}	0.590	0.024	10^{-24}	0.572
	7 mm	0.043	10^{-42}	0.597	0.025	10^{-24}	0.578
	8 mm	0.041	10^{-40}	0.596	0.022	10^{-22}	0.577
	9 mm	0.041	10^{-40}	0.603	0.022	10^{-22}	0.582
	10 mm	0.039	10^{-38}	0.609	0.021	10^{-20}	0.587
Maximum	1 mm	0.007	10^{-7}	0.526	0.004	10^{-4}	0.526
	2 mm	0.006	10^{-6}	0.529	0.003	10^{-3}	0.528
	3 mm	0.002	0.001	0.517	0.001	0.070	0.514
	4 mm	0.000	0.272	0.498	0.000	0.172	0.506
	5 mm	0.000	0.649	0.503	0.003	10^{-3}	0.522
	6 mm	0.001	0.139	0.511	0.006	10^{-6}	0.528
	7 mm	0.002	0.004	0.520	0.011	10^{-11}	0.542
	8 mm	0.003	10^{-3}	0.524	0.013	10^{-13}	0.544
	9 mm	0.004	10^{-4}	0.528	0.015	10^{-15}	0.548
	10 mm	0.005	10^{-5}	0.528	0.017	10^{-17}	0.553
Varition	3 mm	0.008	10^{-8}	0.527	0.004	10^{-4}	0.519
	4 mm	0.011	10^{-10}	0.538	0.006	10^{-6}	0.520
	5 mm	0.017	10^{-17}	0.555	0.011	10^{-11}	0.536
	6 mm	0.019	10^{-18}	0.559	0.013	10^{-12}	0.541
	7 mm	0.021	10^{-20}	0.563	0.015	10^{-15}	0.552
	8 mm	0.020	10^{-19}	0.564	0.015	10^{-14}	0.552
	9 mm	0.020	10^{-20}	0.568	0.015	10^{-15}	0.560
	10 mm	0.020	10^{-20}	0.568	0.015	10^{-15}	0.558

Table D.4: Results of single-parameter exponential regression on CT image data ($\log(BV_p) = a_1 \cdot \hat{v}_p + a_0$). See Section 4.5.2.




Figure D.2: Accuracy ($\pm$ standard deviation) of random forest MR feature classifiers vs number of trees for three parameter normalisation strategies (IIR, Utah, and hybrid; see text) and two fibrosis thresholds (Penn: $BV_p < 0.27$ mV, Melbourne: $BV_p < 0.5$ mV).




Figure D.3: Top: Accuracy ($\pm$ standard deviation) of random forest MR feature regressors vs number of trees for three parameter normalisation strategies (IIR, Utah, and hybrid; see text). Bottom: For each normalisation strategy, area under median random forest regressor ROC curve for bipolar voltage thresholds.

*The first kind of intellectual and artistic personality
belongs to the hedgehogs, the second to the foxes . . .*

— Sir Isaiah Berlin^[351]

References

- [1] Hippocrates, “Aphorisms,” in *Nature of Man. Regimen in Health. Humours. Aphorisms. Regimen 1-3. Dreams. Heracleitus: On the Universe*, W. H. S. Jones, Trans., number 150 in Loeb Classical Library, pp. 97–222. Harvard University Press, Boston, MA, 1931.
- [2] Dena M. Bravata, Allison L. Gienger, Kathryn M. McDonald, Vandana Sundaram, Marco V. Perez, Robin Varghese, John R. Kapoor, Reza Ardehali, Douglas K. Owens, and Mark A. Hlatky, “Systematic review: The comparative effectiveness of percutaneous coronary interventions and coronary artery bypass graft surgery,” *Annals of Internal Medicine*, vol. 147, no. 10, pp. 703–716, 2007.
- [3] David K. Foot, Richard P. Lewis, Thomas A. Pearson, and George A. Beller, “Demographics and cardiology, 1950–2050,” *Journal of the American College of Cardiology*, vol. 35, no. 4, pp. 1067–1081, 2000.
- [4] J. Eduardo Sousa, Marco A. Costa, E. Murat Tuzcu, Jay S. Yadav, and Stephen Ellis, “New frontiers in interventional cardiology,” *Circulation*, vol. 111, no. 5, pp. 671–681, 2005.
- [5] Jun Dong, Hugh Calkins, Stephen B. Solomon, Shenghan Lai, Darshan Dalal, Al Lardo, Erez Brem, Assaf Preiss, Ronald D. Berger, Henry Halperin, and Timm Dickfeld, “Integrated electroanatomic mapping with three-dimensional computed tomographic images for real-time guided ablations,” *Circulation*, vol. 113, no. 2, pp. 186–194, 2006.
- [6] Stephan Willems, Christian Weiss, Rodolfo Ventura, Rudolf Ruppel, Tim Risius, Matthias Hoffmann, and Thomas Meinertz, “Catheter ablation of atrial flutter guided by electroanatomic mapping (CARTO): A randomized comparison to the conventional approach,” *Journal of Cardiovascular Electrophysiology*, vol. 11, no. 11, pp. 1223–1230, 2000.
- [7] Yoram Rudy, “Noninvasive electrocardiographic imaging of arrhythmogenic substrates in humans,” *Circulation Research*, vol. 112, no. 5, pp. 863–874, 2013.
- [8] Charulatha Ramanathan, Raja N. Ghanem, Ping Jia, Kyungmoo Ryu, and Yoram Rudy, “Noninvasive electrocardiographic imaging for cardiac electrophysiology and arrhythmia,” *Nature Medicine*, vol. 10, no. 4, pp. 422–428, 2004.
- [9] Maxime Sermesant, Jean-Marc Peyrat, Phani Chinchapatnam, Florence Billet, Tommaso Mansi, Kawal Rhode, Hervé Delingette, Reza Razavi, and Nicholas Ayache, “Toward patient-specific myocardial models of the heart,” *Heart Failure Clinics*, vol. 4, no. 3, pp. 289–301, 2008.
- [10] Hugo Talbot, Stéphanie Marchesseau, Christian Duriez, Maxime Sermesant, Stéphane Cotin, and Hervé Delingette, “Towards an interactive electromechanical model of the heart,” *Interface Focus*, vol. 3, no. 20120091, 2013.

- [11] Malcolm C. Finlay, Ross J. Hunter, Victoria Baker, Laura Richmond, Farai Goromonzi, Glyn Thomas, Kim Rajappan, Edward Duncan, Muzahir Tayebjee, Mehul Dhinoja, Simon Sporton, Mark J. Earley, and Richard J. Schilling, “A randomised comparison of CartoMerge vs. NavX fusion in the catheter ablation of atrial fibrillation: The CAVERN Trial,” *Journal of Interventional Cardiac Electrophysiology*, vol. 33, no. 2, pp. 161–169, 2012.
- [12] Bernhard L. Gerber, Bénédicte Belge, Gabin J. Legros, Pascal Lim, Alain Poncelet, Agnès Pasquet, Giovanna Gisellu, Emmanuel Coche, and Jean-Louis J. Vanoverschelde, “Characterization of acute and chronic myocardial infarcts by multidetector computed tomography: Comparison with contrast-enhanced magnetic resonance,” *Circulation*, vol. 113, no. 6, pp. 823–833, 2006.
- [13] Katherine C. Wu, Robert G. Weiss, David R. Thiemann, Kakuya Kitagawa, André Schmidt, Darshan Dalal, Shenghan Lai, David A. Bluemke, Gary Gerstenblith, Eduardo Marbán, Gordon F. Tomaselli, and João A. C. Lima, “Late gadolinium enhancement by cardiovascular magnetic resonance heralds an adverse prognosis in nonischemic cardiomyopathy,” *Journal of the American College of Cardiology*, vol. 51, no. 25, pp. 2414–2421, 2008.
- [14] Qian Tao, Julien Milles, Carine van Huls van Taxis, Hildo J. Lamb, Johan H. C. Reiber, Katja Zeppenfeld, and Rob J. van der Geest, “Toward magnetic resonance-guided electroanatomical voltage mapping for catheter ablation of scar-related ventricular tachycardia: A comparison of registration methods,” *Journal of Cardiovascular Electrophysiology*, vol. 23, no. 1, pp. 74–80, 2012.
- [15] Zhiyu Ling, John McManigle, Irfan Khurram, Stefan Zimmerman, Joseph Marine, David Spragg, Hiroshi Ashikaga, Hugh Calkins, and Saman Nazarian, “Association between cardiac computed tomography perfusion imaging and electronatomic mapping to evaluate left atrial fibrosis in patients with atrial fibrillation,” *Circulation*, vol. 128, no. 22 Supplement, p. A18282, 2013.
- [16] Kotaro Fukumoto, John McManigle, Irfan Khurram, Stefan Zimmerman, Vadim Zipunnikov, David Spragg, Joseph Marine, Hiroshi Ashikaga, Jack Rickard, Ronald Berger, Hugh Calkins, and Saman Nazarian, “Comparison of native versus ablation-induced scar on left atrial magnetic resonance imaging,” in *35th Heart Rhythm Society*, San Francisco, CA, May 2014.
- [17] J. E. McManigle, R. V. Stebbing, and J. A. Noble, “Modified Hough transform for left ventricle myocardium segmentation in 3-D echocardiogram images,” in *9th IEEE International Symposium on Biomedical Imaging (ISBI)*, Barcelona, May 2012, pp. 290–293.
- [18] Johann Wolfgang von Goethe, *Wilhelm Meisters Wanderjahre oder die Entsagenden*, Cotta, 1829.
- [19] Dennis D. Patton, “Röntgen and the Discovery,” in *A History of the Radiological Sciences: Diagnosis*, Raymond A. Gagliardi and Bruce L. McClelland, Eds., number 1 in A History of the Radiological Sciences, pp. 31–46. Radiology Centennial, Inc., 1996.
- [20] Herbert L. Abrams, “Cardiac radiology,” in *A History of the Radiological Sciences: Diagnosis*, Raymond A. Gagliardi and Bruce L. McClelland, Eds., number 1 in A History of the Radiological Sciences, pp. 255–269. Radiology Centennial, Inc., 1996.

- [21] Andrew Webb, *Introduction to Biomedical Imaging*, Wiley-IEEE Press, 2002.
- [22] Paul Leeson, Ed., *Cardiovascular Imaging*, Oxford Specialist Handbooks in Cardiology. Oxford University Press, Oxford, 1st edition, 2011.
- [23] Kjetil Lenes, “4-D echocardiogram” [image]. Accessed 8 Jan 2014: <http://commons.wikimedia.org/wiki/File:Apikal4D.gif>, under the Creative Commons Attribution-Share Alike 3.0 Unported license.
- [24] Wikimedia Commons, “Echocardiogram” [images]. Accessed 8 Jan 2014: http://commons.wikimedia.org/wiki/File:PLAX_Mmode.jpg and *ibid*/[File:Echocardiogram_4chambers.jpg](http://commons.wikimedia.org/wiki/File:Echocardiogram_4chambers.jpg), under public domain declarations.
- [25] Patrick J. Lynch, “Heart left lateral view of left ventricle” [image]. Accessed 8 Jan 2014: http://commons.wikimedia.org/wiki/File:Heart_left_lateral_view.jpg, under the Creative Commons Attribution 2.5 Generic license.
- [26] Wikimedia Commons, “Human body, front and side” [image]. Accessed 7 Jan 2014: http://commons.wikimedia.org/wiki/File:Human_body_front_and_side.svg, under the Creative Commons Attribution-Share Alike 3.0 Unported license.
- [27] Barry B. Goldberg, Raymond Gramiak, and Atis K. Freimanis, “Ultrasonography,” in *A History of the Radiological Sciences: Diagnosis*, Raymond A. Gagliardi and Bruce L. McClennan, Eds., number 1 in A History of the Radiological Sciences, pp. 413–442. Radiology Centennial, Inc., 1996.
- [28] Inge Edler and Arne Gustafson, “Ultrasonic cardiogram in mitral stenosis,” *Acta Medica Scandinavica*, vol. 159, no. 2, pp. 85–90, 1957.
- [29] N. Bom, C. T. Lancée, J. Honkoop, and P. G. Hugenholtz, “Ultrasonic viewer for cross-sectional analyses of moving cardiac structures,” *Biomedical Engineering*, vol. 6, no. 11, pp. 500–505, 1971.
- [30] James M. Griffith and Walter L. Henry, “A sector scanner for real time two-dimensional echocardiography,” *Circulation*, vol. 49, no. 6, pp. 1147–1152, 1974.
- [31] Robert O. Bonow, Blase A. Carabello, Kanu Chatterjee, Antonio C. de Leon, David P. Faxon, Michael D. Freed, William H. Gaasch, Bruce Whitney Lytle, Rick A. Nishimura, Patrick T. O’Gara, Robert A. O’Rourke, Catherine M. Otto, Pravin M. Shah, and Jack S. Shanewise, “ACC/AHA 2006 guidelines for the management of patients with valvular heart disease,” *Circulation*, vol. 114, no. 5, pp. e84–231, 2006.
- [32] R. B. Devereux, P. N. Casale, D. C. Wallerson, P. Kligfield, I. W. Hammond, P. R. Liebson, E. Campo, D. R. Alonso, and J. H. Laragh, “Cost-effectiveness of echocardiography and electrocardiography for detection of left ventricular hypertrophy in patients with systemic hypertension,” *Hypertension*, vol. 9, no. 2, pt. 2, pp. II69–II76, 1987.
- [33] Alan H. Gradman and Fadi Alfayoumi, “From left ventricular hypertrophy to congestive heart failure: Management of hypertensive heart disease,” *Progress in Cardiovascular Diseases*, vol. 48, no. 5, pp. 326–341, 2006.

- [34] Hongxia Niu, Wei Hua, Shu Zhang, Xin Sun, Fangzheng Wang, Keping Chen, and Xin Chen, "Prevalence of dyssynchrony derived from echocardiographic criteria in heart failure patients with normal or prolonged QRS duration," *Echocardiography*, vol. 24, no. 4, pp. 348–352, 2007.
- [35] H. H. Echeverria, M. S. Bilsker, R. J. Myerburg, and K. M. Kessler, "Congestive heart failure: Echocardiographic insights," *American Journal of Medicine*, vol. 75, no. 5, pp. 750–755, 1983.
- [36] Jo Mahenthiran, Sripal Bangalore, Siu-Sun Yao, and Farooq A. Chaudhry, "Comparison of prognostic value of stress echocardiography versus stress electrocardiography in patients with suspected coronary artery disease," *American Journal of Cardiology*, vol. 96, no. 5, pp. 628–634, 2005.
- [37] Thomas H. Marwick, "Stress echocardiography," *Heart*, vol. 89, no. 1, pp. 113–118, 2003.
- [38] Bangalore Sripal, Yao Siu-Sun, and A. Chaudhry Farooq, "Stress function index, a novel index for risk stratification and prognosis using stress echocardiography," *Journal of the American Society of Echocardiography*, vol. 18, no. 12, pp. 1335–1342, 2005.
- [39] Pamela S. Douglas, John Fiolkoski, Barbara Berko, and Nathaniel Reichek, "Echocardiographic visualization of coronary artery anatomy in the adult," *Journal of the American College of Cardiology*, vol. 11, no. 3, pp. 565–571, 1988.
- [40] Alexander Ng and Justiaan Swanevelder, "Resolution in ultrasound imaging," *Continuing Education in Anaesthesia, Critical Care & Pain*, vol. 11, no. 5, pp. 186–192, 2011.
- [41] Lissa Sugeng, Victor Mor-Avi, and Roberto M. Lang, "Three-dimensional echocardiography: Coming of age," *Heart*, vol. 94, no. 9, pp. 1123–1125, 2008.
- [42] Elif Eroglu, Jan D'hooge, Lieven Herbots, Daisy Thijs, Christophe Dubois, Peter Sinnaeve, Joseph Dens, Johan Vanhaecke, and Frank Rademakers, "Comparison of real-time tri-plane and conventional 2d dobutamine stress echocardiography for the assessment of coronary artery disease," *European Heart Journal*, vol. 27, no. 14, pp. 1719–1724, 2006.
- [43] Yoshiki Matsumura, Takeshi Hozumi, Kotaro Arai, Kenichi Sugioka, Keiji Ujino, Yasuhiko Takemoto, Hiroyuki Yamagishi, Minoru Yoshiyama, and Junichi Yoshikawa, "Non-invasive assessment of myocardial ischaemia using new real-time three-dimensional dobutamine stress echocardiography: Comparison with conventional two-dimensional methods," *European Heart Journal*, vol. 26, no. 16, pp. 1625–1632, 2005.
- [44] Allan L. Klein, R. Daniel Murray, and Richard A. Grimm, "Role of transesophageal echocardiography-guided cardioversion of patients with atrial fibrillation," *Journal of the American College of Cardiology*, vol. 37, no. 3, pp. 691–704, 2001.
- [45] T. Masuyama, K. Kodama, A. Kitabatake, H. Sato, S. Nanto, and M. Inoue, "Continuous-wave Doppler echocardiographic detection of pulmonary regurgitation and its application to noninvasive estimation of pulmonary artery pressure," *Circulation*, vol. 74, no. 3, pp. 484–492, 1986.

- [46] Shelley A. Hall, M. Elizabeth Brickner, DuWayne L. Willett, Waleed N. Irani, Imran Afridi, and Paul A. Grayburn, "Assessment of mitral regurgitation severity by Doppler color flow mapping of the vena contracta," *Circulation*, vol. 95, no. 3, pp. 636–642, 1997.
- [47] Gerhard-Michael von Reutern, Michael-Wolfgang Goertler, Natan M. Bornstein, Massimo Del Sette, David H. Evans, Michael-Wolfgang Goertler, Andreas Hetzel, Manfred Kaps, Fabienne Perren, Alexander Razumovsky, Toshiyuki Shiogai, Ekaterina Titianova, Pavel Traubner, Narayanaswamy Venketasubramanian, Lawrence K. S. Wong, and Masahiro Yasaka, "Grading carotid stenosis using ultrasonic methods," *Stroke*, vol. 43, no. 3, pp. 916–921, 2012.
- [48] J. G. Stevenson, I. Kawabori, T. Dooley, and W. G. Guntheroth, "Diagnosis of ventricular septal defect by pulsed Doppler echocardiography: Sensitivity, specificity and limitations," *Circulation*, vol. 58, no. 2, pp. 322–326, 1978.
- [49] Thomas F. Meaney, "Magnetic resonance imaging," in *A History of the Radiological Sciences: Diagnosis*, Raymond A. Gagliardi and Bruce L. McClennan, Eds., number 1 in A History of the Radiological Sciences, pp. 471–482. Radiology Centennial, Inc., 1996.
- [50] Henry Gray, *Anatomy of the Human Body*, Warren H. Lewis, Ed., Lea & Febiger, Philadelphia and New York, 20th edition, 1918.
- [51] I. R. Young, D. R. Bailes, M. Burl, A. G. Collins, D. T. Smith, M. J. McDonnell, J. S. Orr, L. M. Banks, G. M. Bydder, R. H. Greenspan, and R. E. Steiner, "Initial clinical evaluation of a whole body nuclear magnetic resonance (NMR) tomograph," *Journal of Computer Assisted Tomography*, vol. 6, no. 1, pp. 1–18, 1982.
- [52] J. Paul Finn, Kambiz Nael, Vibhas Deshpande, Osman Ratib, and Gerhard Laub, "Cardiac MR imaging: State of the technology," *Radiology*, vol. 241, no. 2, pp. 338–354, 2006.
- [53] L. Y. Hsu, A. Natanzon, P. Kellman, G. A. Hirsch, A. H. Aletras, and A. E. Arai, "Quantitative myocardial infarction on delayed enhancement MRI. Part I: Animal validation of an automated feature analysis and combined thresholding infarct sizing algorithm," *Journal of Magnetic Resonance Imaging*, vol. 23, no. 3, pp. 298–308, 2006.
- [54] L. Y. Hsu, W. P. Ingkanisorn, P. Kellman, A. H. Aletras, and A. E. Arai, "Quantitative myocardial infarction on delayed enhancement MRI. Part II: Clinical application of an automated feature analysis and combined thresholding infarct sizing algorithm," *Journal of Magnetic Resonance Imaging*, vol. 23, no. 3, pp. 309–314, 2006.
- [55] David A. Feinberg, Lawrence E. Crooks, Phillip Sheldon, John Hoenninger Iii, Jeffrey Watts, and Mitsuaki Arakawa, "Magnetic resonance imaging the velocity vector components of fluid flow," *Magnetic Resonance in Medicine*, vol. 2, no. 6, pp. 555–566, 1985.
- [56] Eike Nagel, Axel Bornstedt, Jürgen Hug, Bernhard Schnackenburg, Ernst Wellnhöfer, and Eckart Fleck, "Noninvasive determination of coronary blood flow velocity with magnetic resonance imaging: Comparison of breath-hold and navigator techniques with intravascular ultrasound," *Magnetic Resonance in Medicine*, vol. 41, no. 3, pp. 544–549, 1999.

- [57] Aravind Arepally, Parag V. Karmarkar, Clifford Weiss, E. Rene Rodriguez, Robert J. Lederman, and Ergin Atalar, "Magnetic resonance image-guided trans-septal puncture in a swine heart," *Journal of Magnetic Resonance Imaging*, vol. 21, no. 4, pp. 463–467, 2005.
- [58] M. G. Lentschig, P. Reimer, U. L. Rausch-Lentschig, T. Allkemper, M. Oelerich, and G. Laub, "Breath-hold gadolinium-enhanced MR angiography of the major vessels at 1.0 T: Dose-response findings and angiographic correlation," *Radiology*, vol. 208, no. 2, pp. 353–357, 1998.
- [59] R. C. Brasch, "Work in progress: Methods of contrast enhancement for NMR imaging and potential applications," *Radiology*, vol. 147, no. 3, pp. 781–788, 1983.
- [60] Michael P. Hartung, Thomas M. Grist, and Christopher J. François, "Magnetic resonance angiography: Current status and future directions," *Journal of Cardiovascular Magnetic Resonance*, vol. 13, p. 19, 2011.
- [61] D. Tscholakoff, C. B. Higgins, U. Sechtem, and M. T. McNamara, "Occlusive and reperfused myocardial infarcts: Effect of Gd-DTPA on ECG-gated MR imaging," *Radiology*, vol. 160, no. 2, pp. 515–519, 1986.
- [62] Orlando P. Simonetti, Raymond J. Kim, David S. Fieno, Hanns B. Hillenbrand, Edwin Wu, Jeffrey M. Bundy, J. Paul Finn, and Robert M. Judd, "An improved MR imaging technique for the visualization of myocardial infarction," *Radiology*, vol. 218, no. 1, pp. 215–223, 2001.
- [63] Raymond J. Kim, Edwin Wu, Allen Rafael, Enn-Ling Chen, Michele A. Parker, Orlando Simonetti, Francis J. Klocke, Robert O. Bonow, and Robert M. Judd, "The use of contrast-enhanced magnetic resonance imaging to identify reversible myocardial dysfunction," *New England Journal of Medicine*, vol. 343, no. 20, pp. 1445–1453, 2000.
- [64] Benjamin R. Knowles, Philip G. Batchelor, Victoria Parish, Matthew Ginks, Sven Plein, Reza Razavi, and Tobias Schaeffter, "Pharmacokinetic modeling of delayed gadolinium enhancement in the myocardium," *Magnetic Resonance in Medicine*, vol. 60, no. 6, pp. 1524–1530, 2008.
- [65] John J. Green, Jeffery S. Berger, Christopher M. Kramer, and Michael Salerno, "Prognostic value of late gadolinium enhancement in clinical outcomes for hypertrophic cardiomyopathy," *Journal of the American College of Cardiology: Cardiovascular Imaging*, vol. 5, no. 4, pp. 370–377, 2012.
- [66] Giancarlo Casolo, Sergio Minneci, Rosanna Manta, Antonio Sulla, Jacopo Del Meglio, Luigi Rega, and Gianfranco Gensini, "Identification of the ischemic etiology of heart failure by cardiovascular magnetic resonance imaging: Diagnostic accuracy of late gadolinium enhancement," *American Heart Journal*, vol. 151, no. 1, pp. 101–108, 2006.
- [67] Hui Xue, Peter Kellman, Gina LaRocca, Andrew E. Arai, and Michael S. Hansen, "High spatial and temporal resolution retrospective cine cardiovascular magnetic resonance from shortened free breathing real-time acquisitions," *Journal of Cardiovascular Magnetic Resonance*, vol. 15, no. 1, p. 102, 2013.
- [68] Yiping P. Du, Elliot R. McVeigh, David A. Bluemke, Harry A. Silber, and Thomas K.F. Foo, "A comparison of prospective and retrospective respiratory

- navigator gating in 3d MR coronary angiography,” *The international journal of cardiovascular imaging*, vol. 17, no. 4, pp. 287–296, 2001.
- [69] Krishna S. Nayak, Charles H. Cunningham, Juan M. Santos, and John M. Pauly, “Real-time cardiac MRI at 3 tesla,” *Magnetic Resonance in Medicine*, vol. 51, no. 4, pp. 655–660, 2004.
- [70] Christopher J. McGann, Eugene G. Kholmovski, Robert S. Oakes, Joshua J. E. Blauer, Marcos Daccarett, Nathan Segerson, Kelly J. Airey, Nazem Akoum, Eric Fish, Troy J. Badger, Edward V. R. DiBella, Dennis Parker, Rob S. MacLeod, and Nassir F. Marrouche, “New magnetic resonance imaging-based method for defining the extent of left atrial wall injury after the ablation of atrial fibrillation,” *Journal of the American College of Cardiology*, vol. 52, no. 15, pp. 1263–1271, 2008.
- [71] P. Kellman, C. Chefd’hotel, C. H. Lorenz, C. Mancini, A. E. Arai, and E. R. McVeigh, “High spatial and temporal resolution cardiac cine MRI from retrospective reconstruction of data acquired in real time using motion correction and resorting,” *Magnetic Resonance in Medicine*, vol. 62, no. 6, pp. 1557–64, 2009.
- [72] Mark Perazella, “Gadolinium-contrast toxicity in patients with kidney disease: Nephrotoxicity and nephrogenic systemic fibrosis,” *Current Drug Safety*, vol. 3, no. 1, pp. 67–75, 2008.
- [73] J. T. Littleton and M. L. Duricz Littleton, “Conventional tomography,” in *A History of the Radiological Sciences: Diagnosis*, Raymond A. Gagliardi and Bruce L. McClelland, Eds., number 1 in A History of the Radiological Sciences, pp. 369–401. Radiology Centennial, Inc., 1996.
- [74] Johann Radon, “On the determination of functions from their integral values along certain manifolds,” *IEEE Transactions on Medical Imaging*, vol. 5, no. 4, pp. 170–176, 1986.
- [75] Johann Radon, “Über die Bestimmung von Funktionen durch ihre Integralwerte längs gewisser Mannigfaltigkeiten,” *Akademie der Wissenschaften*, vol. 69, pp. 262–277, 1917.
- [76] Ronald G. Evens, “Computed tomography,” in *A History of the Radiological Sciences: Diagnosis*, Raymond A. Gagliardi and Bruce L. McClelland, Eds., number 1 in A History of the Radiological Sciences, pp. 403–411. Radiology Centennial, Inc., 1996.
- [77] Francesca Pugliese, Nico R. A. Mollet, Giuseppe Runza, Carlos Miegheem, Willem B. Meijboom, Patrizia Malagutti, Timo Baks, Gabriel P. Krestin, Pim J. deFeyter, and Filippo Cademartiri, “Diagnostic accuracy of non-invasive 64-slice CT coronary angiography in patients with stable angina pectoris,” *European Radiology*, vol. 16, no. 3, pp. 575–582, 2005.
- [78] T. Q. Callister, B. Cooil, S. P. Raya, N. J. Lippolis, D. J. Russo, and P. Raggi, “Coronary artery disease: Improved reproducibility of calcium scoring with an electron-beam CT volumetric method,” *Radiology*, vol. 208, no. 3, pp. 807–814, 1998.
- [79] J. F. Breen, P. F. Sheedy, R. S. Schwartz, A. W. Stanson, R. B. Kaufmann, P. P. Moll, and J. A. Rumberger, “Coronary artery calcification detected with ultrafast

- CT as an indication of coronary artery disease,” *Radiology*, vol. 185, no. 2, pp. 435–439, 1992.
- [80] Jeffrey J. Pasternak and Eric E. Williamson, “Clinical pharmacology, uses, and adverse reactions of iodinated contrast agents: A primer for the non-radiologist,” *Mayo Clinic Proceedings*, vol. 87, no. 4, pp. 390–402, 2012.
- [81] Maurice Laville and Laurent Juillard, “Contrast-induced acute kidney injury: how should at-risk patients be identified and managed?,” *Journal of Nephrology*, , no. Vol. 23 Issue 4 (July-August 2010), pp. 387–398, 2010.
- [82] Thomas C. Gerber, J. Jeffrey Carr, Andrew E. Arai, Robert L. Dixon, Victor A. Ferrari, Antoinette S. Gomes, Gary V. Heller, Cynthia H. McCollough, Michael F. McNitt-Gray, Fred A. Mettler, Jennifer H. Mieres, Richard L. Morin, and Michael V. Yester, “Ionizing radiation in cardiac imaging,” *Circulation*, vol. 119, no. 7, pp. 1056–1065, 2009.
- [83] Timo Baks, Filippo Cademartiri, Amber D. Moelker, Annick C. Weustink, Robert-Jan van Geuns, Nico R. Mollet, Gabriel P. Krestin, Dirk J. Duncker, and Pim J. de Feyter, “Multislice computed tomography and magnetic resonance imaging for the assessment of reperfused acute myocardial infarction,” *Journal of the American College of Cardiology*, vol. 48, no. 1, pp. 144–152, 2006.
- [84] Yasushi Koyama, Teruhito Mochizuki, and Jitsuo Higaki, “Computed tomography assessment of myocardial perfusion, viability, and function,” *Journal of magnetic resonance imaging: JMRI*, vol. 19, no. 6, pp. 800–815, 2004.
- [85] B. Erem, J. Coll-Font, R. Martinez Orellana, P. St’ovicek, and D.H. Brooks, “Using transmural regularization and dynamic modeling for noninvasive cardiac potential imaging of endocardial pacing with imprecise thoracic geometry,” *IEEE Transactions on Medical Imaging*, vol. 33, no. 3, pp. 726–738, 2014.
- [86] M. Sermesant, R. Chabiniok, P. Chinchapatnam, T. Mansi, F. Billet, P. Moireau, J. M. Peyrat, K. Wong, J. Relan, K. Rhode, M. Ginks, P. Lambiase, H. Delingette, M. Sorine, C. A. Rinaldi, D. Chapelle, R. Razavi, and N. Ayache, “Patient-specific electromechanical models of the heart for the prediction of pacing acute effects in CRT: A preliminary clinical validation,” *Medical Image Analysis*, vol. 16, no. 1, pp. 201–215, 2012.
- [87] Ranil de Silva, Luis F. Gutierrez, Amish N. Raval, Elliot R. McVeigh, Cengizhan Ozturk, and Robert J. Lederman, “X-ray fused with magnetic resonance imaging (XFM) to target endomyocardial injections: validation in a swine model of myocardial infarction,” *Circulation*, vol. 114, no. 22, pp. 2342–2350, 2006.
- [88] Aravindan Kolandaivelu, Albert C. Lardo, and Henry R. Halperin, “Cardiovascular magnetic resonance guided electrophysiology studies,” *Journal of Cardiovascular Magnetic Resonance*, vol. 11, p. 21, 2009.
- [89] Elliot R. McVeigh, Michael A. Guttman, Peter Kellman, Amish N. Raval, and Robert J. Lederman, “Real-time, interactive MRI for cardiovascular interventions,” *Academic Radiology*, vol. 12, no. 9, pp. 1121–1127, 2005.
- [90] Elliot R. McVeigh, Michael A. Guttman, Robert J. Lederman, Ming Li, Ozgur Kocaturk, Timothy Hunt, Shawn Kozlov, and Keith A. Horvath, “Real-time interactive MRI-guided cardiac surgery: Aortic valve replacement using a direct

- apical approach,” *Magnetic Resonance in Medicine*, vol. 56, no. 5, pp. 958–964, 2006.
- [91] S. Nazarian, A. Kollandaivelu, M. M. Zviman, G. R. Meininger, R. Kato, R. C. Susil, A. Roguin, T. L. Dickfeld, H. Ashikaga, H. Calkins, R. D. Berger, D. A. Bluemke, A. C. Lardo, and H. R. Halperin, “Feasibility of real-time magnetic resonance imaging for catheter guidance in electrophysiology studies,” *Circulation*, vol. 118, no. 3, pp. 223–229, 2008.
- [92] A. M. Jungmann, “X-rays: Samaritans of War,” in *The Book of Modern Marvels*, Waldemar Kaempffert, Ed., p. 172. Leslie-Judge Company, 1917.
- [93] Arthur G. Haus and John E. Cullinan, “Recording the Image,” in *A History of the Radiological Sciences: Diagnosis*, Raymond A. Gagliardi and Bruce L. McClennan, Eds., number 1 in A History of the Radiological Sciences, pp. 87–107. Radiology Centennial, Inc., 1996.
- [94] J S Krohmer, “Radiography and fluoroscopy, 1920 to the present.,” *RadioGraphics*, vol. 9, no. 6, pp. 1129–1153, 1989.
- [95] Mark J. Earley, “How to perform a transseptal puncture,” *Heart*, vol. 95, no. 1, pp. 85–92, 2009.
- [96] John D. Kugler, David A. Danford, Kris A. Houston, and Gary Felix, “Pediatric radiofrequency catheter ablation registry success, fluoroscopy time, and complication rate for supraventricular tachycardia: Comparison of early and recent eras,” *Journal of Cardiovascular Electrophysiology*, vol. 13, no. 4, pp. 336–341, 2002.
- [97] F. Larrazet, A. Dibie, F. Philippe, R. Palau, R. Klausz, and F. Laborde, “Factors influencing fluoroscopy time and dose-area product values during ad hoc one-vessel percutaneous coronary angioplasty,” *British Journal of Radiology*, vol. 76, no. 907, pp. 473–477, 2003.
- [98] H. I. Michelena, M. D. Abel, R. M. Suri, W. K. Freeman, R. L. Click, T. M. Sundt, H. V. Schaff, and M. Enriquez-Sarano, “Intraoperative echocardiography in valvular heart disease: An evidence-based appraisal,” *Mayo Clinic Proceedings*, vol. 85, no. 7, pp. 646–655, 2010.
- [99] F. E. Silvestry, R. E. Kerber, M. M. Brook, J. D. Carroll, K. M. Eberman, S. A. Goldstein, H. C. Herrmann, S. Homma, R. Mehran, D. L. Packer, A. F. Parisi, T. Pulerwitz, J. B. Seward, T. S. Tsang, and M. A. Wood, “Echocardiography-guided interventions,” *Journal of the American Society of Echocardiography*, vol. 22, no. 3, pp. 213–231, 2009.
- [100] A. Vegas and M. Meineri, “Core review: Three-dimensional transesophageal echocardiography is a major advance for intraoperative clinical management of patients undergoing cardiac surgery,” *Anesthesia and Analgesia*, vol. 110, no. 6, pp. 1548–73, 2010.
- [101] Joshua M. Cooper and Laurence M. Epstein, “Use of intracardiac echocardiography to guide ablation of atrial fibrillation,” *Circulation*, vol. 104, no. 25, pp. 3010–3013, 2001.
- [102] Joep L. R. M. Smeets, Shlomo A. Ben-Haim, Luz-Maria Rodriguez, Carl Timmermans, and Hein J. J. Wellens, “New method for nonfluoroscopic

- endocardial mapping in humans: Accuracy assessment and first clinical results,” *Circulation*, vol. 97, no. 24, pp. 2426–2432, 1998.
- [103] Shlomo Shpun, Lior Gepstein, Gal Hayam, and Shlomo A. Ben-Haim, “Guidance of radiofrequency endocardial ablation with real-time three-dimensional magnetic navigation system,” *Circulation*, vol. 96, no. 6, pp. 2016–2021, 1997.
- [104] Roderick Tung, Shiro Nakahara, Rafael Ramirez, Dorina Gui, Clara Magyar, Chi Lai, Michael Fishbein, and Kalyanam Shivkumar, “Accuracy of combined endocardial and epicardial electroanatomic mapping of a reperfused porcine infarct model: A comparison of electrofield and magnetic systems with histopathologic correlation,” *Heart Rhythm*, vol. 8, no. 3, pp. 439–447, 2011.
- [105] Jürgen Biermann, Christoph Bode, and Stefan Asbach, “Intracardiac echocardiography during catheter-based ablation of atrial fibrillation,” *Cardiology Research and Practice*, vol. 2012, Article ID 921746, 2012.
- [106] Kevin Wilson, Gerard Guiraudon, Douglas L. Jones, and Terry M. Peters, “Mapping of cardiac electrophysiology onto a dynamic patient-specific heart model,” *IEEE Transactions on Medical Imaging*, vol. 28, no. 12, pp. 1870–1880, 2009.
- [107] Ernest J. Ferris and Max L. Baker, “Vascular and interventional radiology,” in *A History of the Radiological Sciences: Diagnosis*, Raymond A. Gagliardi and Bruce L. McClennan, Eds., number 1 in A History of the Radiological Sciences, pp. 271–288. Radiology Centennial, Inc., 1996.
- [108] Werner Forssmann, *Experiments on Myself: Memoirs of a Surgeon in Germany*, Hilary Davies, Trans., Saint Martin’s Press, 1974.
- [109] Werner Forssmann, “Die Sondierung des Rechten Herzens,” *Klinische Wochenschrift*, vol. 8, no. 45, pp. 2085–2087, 1929.
- [110] S. I. Seldinger, “Catheter replacement of the needle in percutaneous arteriography: A new technique,” *Acta Radiologica*, vol. 39, no. 5, pp. 368–376, 1953.
- [111] H. J. Ricketts and H. L. Abrams, “Percutaneous selective coronary cine arteriography,” *Journal of the American Medical Association*, vol. 181, pp. 620–624, 1962.
- [112] Charles T. Dotter and Melvin P. Judkins, “Transluminal treatment of arteriosclerotic obstruction: Description of a new technic and a preliminary report of Its application,” *Circulation*, vol. 30, no. 5, pp. 654–670, 1964.
- [113] C. T. Dotter, “Transluminally-placed coilspring endarterial tube grafts: Long-term patency in canine popliteal artery,” *Investigative Radiology*, vol. 4, no. 5, pp. 329–332, 1969.
- [114] Andreas R. Grüntzig, Åke Senning, and Walter E. Siegenthaler, “Nonoperative dilatation of coronary-artery stenosis,” *New England Journal of Medicine*, vol. 301, no. 2, pp. 61–68, 1979.
- [115] Shaheen E. Lakhan, Anna Kaplan, Cyndi Laird, and Yaacov Leiter, “The interventionalism of medicine: Interventional radiology, cardiology, and neuroradiology,” *International Archives of Medicine*, vol. 2, p. 27, 2009.
- [116] N. Rohrhoff, J. P. Vavalle, S. Halim, T. L. Kiefer, and J. K. Harrison, “Current

- status of percutaneous PFO closure,” *Current Cardiology Reports*, vol. 16, no. 5, pp. 1–9, 2014.
- [117] Roy Beigel, Nina C. Wunderlich, Saibal Kar, and Robert J. Siegel, “The evolution of percutaneous mitral valve repair therapy: Lessons learned and implications for patient selection,” *Journal of the American College of Cardiology*, vol. 64, no. 24, pp. 2688–2700, 2014.
- [118] Jacqueline Saw and Mathieu Lempereur, “Percutaneous left atrial appendage closure: Procedural techniques and putcomes,” *Journal of the American College of Cardiology: Cardiovascular Interventions*, vol. 7, no. 11, pp. 1205–1220, 2014.
- [119] Berndt Lüderitz, “Historical perspectives on interventional electrophysiology,” *Journal of Interventional Cardiac Electrophysiology*, vol. 9, no. 2, pp. 75–83, 2003.
- [120] Albert S. Hyman, “Resuscitation of the stopped heart by intracardial therapy: II. Experimental use of an artificial pacemaker,” *Archives of Internal Medicine*, vol. 50, no. 2, pp. 283–305, 1932.
- [121] Paul M. Zoll, “Resuscitation of the heart in ventricular standstill by external electric stimulation,” *New England Journal of Medicine*, vol. 247, no. 20, pp. 768–771, 1952.
- [122] B. Lown, R. Amarasingham, and J. Neuman, “New method for terminating cardiac arrhythmias: Use of synchronized capacitor discharge,” *Journal of the American Medical Association*, vol. 182, pp. 548–555, 1962.
- [123] Cesar A. Castillo, Barouh V. Berkovits, Agustin Castellanos Jr., Louis Lemberg, George Callard, and James R. Jude, “Bifocal demand pacing,” *Chest*, vol. 59, no. 4, pp. 360–364, 1971.
- [124] Melvin M. Scheinman, Fred Morady, David S. Hess, and Rolando Gonzalez, “Catheter-induced ablation of the atrioventricular junction to control refractory supraventricular arrhythmias,” *Journal of the American Medical Association*, vol. 248, no. 7, pp. 851–855, 1982.
- [125] John J. Gallagher, Robert H. Svenson, Jack H. Kasell, Lawrence D. German, Gust H. Bardy, Archer Broughton, and Giuseppe Critelli, “Catheter technique for closed-chest ablation of the atrioventricular conduction system,” *New England Journal of Medicine*, vol. 306, no. 4, pp. 194–200, 1982.
- [126] Karen E. Thomas and Mark E. Josephson, “The role of electrophysiology study in risk stratification of sudden cardiac death,” *Progress in Cardiovascular Diseases*, vol. 51, no. 2, pp. 97–105, 2008.
- [127] H. Calkins, J. Langberg, J. Sousa, R. el Atassi, A. Leon, W. Kou, S. Kalbfleisch, and F. Morady, “Radiofrequency catheter ablation of accessory atrioventricular connections in 250 patients: Abbreviated therapeutic approach to Wolff-Parkinson-White syndrome,” *Circulation*, vol. 85, no. 4, pp. 1337–1346, 1992.
- [128] Stanley Nattel, “New ideas about atrial fibrillation 50 years on,” *Nature*, vol. 415, no. 6868, pp. 219–226, 2002.
- [129] Emelia J. Benjamin, Philip A. Wolf, Ralph B. D’Agostino, Halit Silbershatz, William B. Kannel, and Daniel Levy, “Impact of atrial fibrillation on the risk of

- death: The Framingham heart study," *Circulation*, vol. 98, no. 10, pp. 946–952, 1998.
- [130] Stephen L. Kopecky, Bernard J. Gersh, Michael D. McGoon, Jack P. Whisnant, David R. Holmes, Duane M. Ilstrup, and Robert L. Frye, "The natural history of lone atrial fibrillation," *New England Journal of Medicine*, vol. 317, no. 11, pp. 669–674, 1987.
- [131] Joseph S. Alpert, Palle Petersen, and John Godtfredsen, "Atrial fibrillation: Natural history, complications, and management," *Annual Review of Medicine*, vol. 39, no. 1, pp. 41–52, 1988.
- [132] G. Fenelon, W. Wijns, E. Andries, and P. Brugada, "Tachycardiomyopathy: Mechanisms and clinical implications," *Pacing and Clinical Electrophysiology*, vol. 19, no. 1, pp. 95–106, 1996.
- [133] Maurits C. E. F. Wijffels, Charles J. H. J. Kirchhof, Rick Dorland, and Maurits A. Allesie, "Atrial fibrillation begets atrial fibrillation: A study in awake chronically instrumented goats," *Circulation*, vol. 92, no. 7, pp. 1954–1968, 1995.
- [134] Maurits Allesie, Jannie Ausma, and Ulrich Schotten, "Electrical, contractile and structural remodeling during atrial fibrillation," *Cardiovascular Research*, vol. 54, no. 2, pp. 230–246, 2002.
- [135] Carlos A. Morillo, George J. Klein, Douglas L. Jones, and Colette M. Guiraudon, "Chronic rapid atrial pacing: Structural, functional, and electrophysiological characteristics of a new model of sustained atrial fibrillation," *Circulation*, vol. 91, no. 5, pp. 1588–1595, 1995.
- [136] Wen-Chung Yu, Shih-Huang Lee, Ching-Tai Tai, Chin-Feng Tsai, Ming-Hsiung Hsieh, Chien-Cheng Chen, Yu-An Ding, Mau-Song Chang, and Shih-Ann Chen, "Reversal of atrial electrical remodeling following cardioversion of long-standing atrial fibrillation in man," *Cardiovascular Research*, vol. 42, no. 2, pp. 470–476, 1999.
- [137] Warren J. Manning, David I. Silverman, Sarah E. Katz, Marilyn F. Riley, Patricia C. Come, Rosalie M. Doherty, Jiyi T. Munson, and Pamela S. Douglas, "Impaired left atrial mechanical function after cardioversion: Relation to the duration of atrial fibrillation," *Journal of the American College of Cardiology*, vol. 23, no. 7, pp. 1535–1540, 1994.
- [138] Eric K. Louie, Dong Liu, Sandra I. Reynertson, Henry S. Loeb, Thomas L. McKiernan, Patrick J. Scanlon, and Robert J. Hariman, "'Stunning' of the left atrium after spontaneous conversion of atrial fibrillation to sinus rhythm: Demonstration by transesophageal Doppler techniques in a canine model," *Journal of the American College of Cardiology*, vol. 32, no. 7, pp. 2081–2086, 1998.
- [139] Jannie Ausma, Maurits Wijffels, Fred Thoné, Luc Wouters, Maurits Allesie, and Marcel Borgers, "Structural changes of atrial myocardium due to sustained atrial fibrillation in the goat," *Circulation*, vol. 96, no. 9, pp. 3157–3163, 1997.
- [140] Gerrit D. Dispersyn, Jannie Ausma, Fred Thoné, Willem Flameng, Jean-Louis J. Vanoverschelde, Maurits A. Allesie, Frans C. S. Ramaekers, and Marcel Borgers, "Cardiomyocyte remodelling during myocardial hibernation and atrial fibrillation: Prelude to apoptosis," *Cardiovascular Research*, vol. 43, no. 4, pp. 947–957, 1999.

- [141] K. U. Thiedemann and V. J. Ferrans, "Left atrial ultrastructure in mitral valvular disease.," *The American Journal of Pathology*, vol. 89, no. 3, pp. 575–604, 1977.
- [142] L. Wouters, G. S. Liu, W. Flameng, V. L. Thijssen, F. Thone, and M. Borgers, "Structural remodeling of atrial myocardium in patients with cardiac valve disease and atrial fibrillation," *Experimental & Clinical Cardiology*, vol. 5, pp. 158–163, 2001.
- [143] George W. H. Bailey, Blaine A. Braniff, E. William Hancock, and Keith E. Cohn, "Relation of left atrial pathology to atrial fibrillation in mitral valvular disease," *Annals of Internal Medicine*, vol. 69, no. 1, pp. 13–20, 1968.
- [144] Dalane W. Kitzman and William D. Edwards, "Age-related changes in the anatomy of the normal human heart," *Journal of Gerontology*, vol. 45, no. 2, pp. M33–M39, 1990.
- [145] P. A. Wolf, R. D. Abbott, and W. B. Kannel, "Atrial fibrillation as an independent risk factor for stroke: The Framingham Study," *Stroke*, vol. 22, no. 8, pp. 983–988, 1991.
- [146] Timothy Watson, Eduard Shantsila, and Gregory Y. H. Lip, "Mechanisms of thrombogenesis in atrial fibrillation: Virchow's triad revisited," *The Lancet*, vol. 373, no. 9658, pp. 155–166, 2009.
- [147] R. M. Heppell, K. E. Berkin, J. M. McLenachan, and J. A. Davies, "Haemostatic and haemodynamic abnormalities associated with left atrial thrombosis in non-rheumatic atrial fibrillation," *Heart*, vol. 77, no. 5, pp. 407–411, 1997.
- [148] Gregory Y. H. Lip, Robby Nieuwlaat, Ron Pisters, Deirdre A. Lane, and Harry J. G. M. Crijns, "Refining clinical risk stratification for predicting stroke and thromboembolism in atrial fibrillation using a novel risk factor-based approach: The euro heart survey on atrial fibrillation," *Chest*, vol. 137, no. 2, pp. 263–272, 2010.
- [149] B. A. Steinberg and J. P. Piccini, "Anticoagulation in atrial fibrillation," *British Medical Journal*, vol. 348, p. g2116, 2014.
- [150] David E. Haines, Michelle Mead-Salley, Miguel Salazar, Francis E. Marchlinski, Erica Zado, Hugh Calkins, Hiran Yarmohammadi, Koonlawee Nademanee, Montawatt Amnueypol, Allan C. Skanes, and Pradyot Saklani, "Dabigatran versus warfarin anticoagulation before and after catheter ablation for the treatment of atrial fibrillation," *Journal of Interventional Cardiac Electrophysiology*, vol. 37, no. 3, pp. 233–239, 2013.
- [151] David R. Holmes, Vivek Y. Reddy, Zoltan G. Turi, Shephal K. Doshi, Horst Sievert, Maurice Buchbinder, Christopher M. Mullin, and Peter Sick, "Percutaneous closure of the left atrial appendage versus warfarin therapy for prevention of stroke in patients with atrial fibrillation: A randomised non-inferiority trial," *The Lancet*, vol. 374, no. 9689, pp. 534–542, 2009.
- [152] John W. Eikelboom, Lars Wallentin, Stuart J. Connolly, Mike Ezekowitz, Jeff S. Healey, Jonas Oldgren, Sean Yang, Marco Alings, Scott Kaatz, Stefan H. Hohnloser, Hans-Christoph Diener, Maria Grazia Franzosi, Kurt Huber, Paul Reilly, Jeanne Varrone, and Salim Yusuf, "Risk of bleeding with 2 doses of dabigatran compared with warfarin in older and younger patients with atrial

- fibrillation: An analysis of the randomized evaluation of long-term anticoagulant therapy (RE-LY) trial,” *Circulation*, vol. 123, no. 21, pp. 2363–2372, 2011.
- [153] A. J. Sanfilippo, V. M. Abascal, M. Sheehan, L. B. Oertel, P. Harrigan, R. A. Hughes, and A. E. Weyman, “Atrial enlargement as a consequence of atrial fibrillation: A prospective echocardiographic study,” *Circulation*, vol. 82, no. 3, pp. 792–797, 1990.
- [154] Ramdas G. Pai, Padmini Varadarajan, and Masato Tanimoto, “Effect of atrial fibrillation on the dynamics of mitral annular area,” *The Journal of Heart Valve Disease*, vol. 12, no. 1, pp. 31–37, 2003.
- [155] Christian T. Ruff, Robert P. Giugliano, Eugene Braunwald, Elaine B. Hoffman, Naveen Deenadayalu, Michael D. Ezekowitz, A. John Camm, Jeffrey I. Weitz, Basil S. Lewis, Alexander Parkhomenko, Takeshi Yamashita, and Elliott M. Antman, “Comparison of the efficacy and safety of new oral anticoagulants with warfarin in patients with atrial fibrillation: A meta-analysis of randomised trials,” *The Lancet*, vol. 383, no. 9921, pp. 955–962, 2014.
- [156] The AFFIRM Investigators, “A comparison of rate control and rhythm control in patients with atrial fibrillation,” *New England Journal of Medicine*, vol. 347, no. 23, pp. 1825–1833, 2002.
- [157] Peter Zimetbaum and Mark E. Josephson, “Is there a role for maintaining sinus rhythm in patients with atrial fibrillation?,” *Annals of Internal Medicine*, vol. 141, no. 9, pp. 720–726, 2004.
- [158] Jerold S. Shinbane, Mark A. Wood, D. Nick Jensen, Kenneth A. Ellenbogen, Adam P. Fitzpatrick, and Melvin M. Scheinman, “Tachycardia-induced cardiomyopathy: A review of animal models and clinical studies,” *Journal of the American College of Cardiology*, vol. 29, no. 4, pp. 709–715, 1997.
- [159] Isabelle C. Van Gelder, Hessel F. Groenveld, Harry J.G.M. Crijns, Ype S. Tuininga, Jan G. P. Tijssen, A. Marco Alings, Hans L. Hillege, Johanna A. Bergsma-Kadijk, Jan H. Cornel, Otto Kamp, Raymond Tukkie, Hans A. Bosker, Dirk J. Van Veldhuisen, and Maarten P. Van den Berg, “Lenient versus strict rate control in patients with atrial fibrillation,” *New England Journal of Medicine*, vol. 362, no. 15, pp. 1363–1373, 2010.
- [160] Vincenza Snow, Kevin B. Weiss, Michael LeFevre, Robert McNamara, Eric Bass, Lee A. Green, Keith Michl, Douglas K. Owens, Jeffrey Susman, Deborah I. Allen, and Christel Mottur-Pilson, “Management of newly detected atrial fibrillation: A clinical practice guideline from the American Academy of Family Physicians and the American College of Physicians,” *Annals of Internal Medicine*, vol. 139, no. 12, pp. 1009–1017, 2003.
- [161] Paul Dorian, “Rate control in atrial fibrillation,” *New England Journal of Medicine*, vol. 362, no. 15, pp. 1439–1441, 2010.
- [162] Brian H. Sarter and Francis E. Marchlinski, “Redefining the role of digoxin in the treatment of atrial fibrillation,” *The American Journal of Cardiology*, vol. 69, no. 18, pp. 71–81, 1992.
- [163] Abhay Bajpai, Irina Savelieva, and A. John Camm, “Treatment of atrial fibrillation,” *British Medical Bulletin*, vol. 88, no. 1, pp. 75–94, 2008.

- [164] Denis Roy, Mario Talajic, Paul Dorian, Stuart Connolly, Mark J. Eisenberg, Martin Green, Teresa Kus, Jean Lambert, Marc Dubuc, Pierre Gagné, Stanley Nattel, and Bernard Thibault, "Amiodarone to prevent recurrence of atrial fibrillation," *New England Journal of Medicine*, vol. 342, no. 13, pp. 913–920, 2000.
- [165] J. L. Cox, R. B. Schuessler, H. J. D'Agostino, C. M. Stone, B. C. Chang, M. E. Cain, P. B. Corr, and J. P. Boineau, "The surgical treatment of atrial fibrillation. III. Development of a definitive surgical procedure," *The Journal of Thoracic and Cardiovascular Surgery*, vol. 101, no. 4, pp. 569–583, 1991.
- [166] Sydney L. Gaynor, Richard B. Schuessler, Marci S. Bailey, Yosuke Ishii, John P. Boineau, Marye J. Gleva, James L. Cox, and Ralph J. Damiano Jr, "Surgical treatment of atrial fibrillation: Predictors of late recurrence," *The Journal of Thoracic and Cardiovascular Surgery*, vol. 129, no. 1, pp. 104–111, 2005.
- [167] J. J. Defauw, G. M. Guiraudon, N. M. van Hemel, F. E. Vermeulen, J. H. Kingma, and J. M. de Bakker, "Surgical therapy of paroxysmal atrial fibrillation with the "corridor" operation," *The Annals of Thoracic Surgery*, vol. 53, no. 4, pp. 564–570; discussion 571, 1992.
- [168] Gustavo Glotz de Lima, Renato A. K. Kalil, Tiago L. L. Leiria, Domingos M. Hatem, Claudio L. Kruse, Rogério Abrahão, João R. M. Sant'anna, Paulo R. Prates, and Ivo A. Nesralla, "Randomized study of surgery for patients with permanent atrial fibrillation as a result of mitral valve disease," *The Annals of Thoracic Surgery*, vol. 77, no. 6, pp. 2089–2094, 2004.
- [169] Roberto Mantovan, Antonio Raviele, Gianfranco Buja, Emanuele Bertaglia, Francesco Cesari, Alessandra Pedrocco, Claudio Zussa, Gino Gerosa, Carlo Valfrè, Paolo Stritoni, and North-eastern Italian Study on Radiofrequency Surgical Treatment of Atrial Fibrillation Investigators, "Left atrial radiofrequency ablation during cardiac surgery in patients with atrial fibrillation," *Journal of Cardiovascular Electrophysiology*, vol. 14, no. 12, pp. 1289–1295, 2003.
- [170] Hans Kottkamp, Gerhard Hindricks, Rüdiger Autschbach, Beate Krauss, Bernhard Strasser, Petra Schirdewahn, Alexander Fabricius, Gerhard Schuler, and Friedrich-Wilhelm Mohr, "Specific linear left atrial lesions in atrial fibrillation: Intraoperative radiofrequency ablation using minimally invasive surgical techniques," *Journal of the American College of Cardiology*, vol. 40, no. 3, pp. 475–480, 2002.
- [171] Bruno Chiappini, Sofia Martín-Suàrez, Antonino LoForte, Giorgio Arpesella, Roberto Di Bartolomeo, and Giuseppe Marinelli, "Cox/Maze III operation versus radiofrequency ablation for the surgical treatment of atrial fibrillation: A comparative study," *The Annals of Thoracic Surgery*, vol. 77, no. 1, pp. 87–92, 2004.
- [172] M. Haïssaguerre, P. Jaïs, D. C. Shah, L. Gencel, V. Pradeau, S. Garrigues, S. Chouairi, M. Hocini, P. Le Métayer, R. Roudaut, and J. Clémenty, "Right and left atrial radiofrequency catheter therapy of paroxysmal atrial fibrillation," *Journal of Cardiovascular Electrophysiology*, vol. 7, no. 12, pp. 1132–1144, 1996.
- [173] Michel Haïssaguerre, Pierre Jaïs, Dipen C. Shah, Atsushi Takahashi, Méléze Hocini, Gilles Quiniou, Stéphane Garrigue, Alain Le Mouroux, Philippe Le Métayer, and Jacques Clémenty, "Spontaneous initiation of atrial fibrillation by ectopic beats

- originating in the pulmonary veins,” *New England Journal of Medicine*, vol. 339, no. 10, pp. 659–666, 1998.
- [174] Carlo Pappone, Salvatore Rosanio, Giuseppe Oreto, Monica Tocchi, Filippo Gugliotta, Gabriele Vicedomini, Adriano Salvati, Cosimo Dicandia, Patrizio Mazzone, Vincenzo Santinelli, Simone Gulletta, and Sergio Chierchia, “Circumferential radiofrequency ablation of pulmonary vein ostia: A new anatomic approach for curing atrial fibrillation,” *Circulation*, vol. 102, no. 21, pp. 2619–2628, 2000.
- [175] Aakriti Gupta, Tharani Perera, Anand Ganesan, Thomas Sullivan, Dennis H. Lau, Kurt C. Roberts-Thomson, Anthony G. Brooks, and Prashanthan Sanders, “Complications of catheter ablation of atrial fibrillation: A systematic review,” *Circulation: Arrhythmia and Electrophysiology*, vol. 6, no. 6, pp. 1082–1088, 2013.
- [176] Abhishek Deshmukh, Nileshkumar J. Patel, Sadip Pant, Neeraj Shah, Ankit Chothani, Kathan Mehta, Peeyush Grover, Vikas Singh, Srikanth Vallurupalli, Ghanshyambhai T. Savani, Apurva Badheka, Tushar Tuliani, Kaustubh Dabhadkar, George Dibbu, Y. Madhu Reddy, Asif Sewani, Marcin Kowalski, Raul Mitrani, Hakan Paydak, and Juan F. Viles-Gonzalez, “In-hospital complications associated with catheter ablation of atrial fibrillation in the United States between 2000 and 2010: Analysis of 93,801 procedures,” *Circulation*, vol. 128, no. 19, pp. 2104–2112, 2013.
- [177] Fiorenzo Gaita, Riccardo Riccardi, Domenico Caponi, Dipen Shah, Lucia Garberoglio, Laura Vivalda, Alessandro Dulio, Andrea Chiecchio, Eric Manasse, and Roberto Gallotti, “Linear cryoablation of the left atrium versus pulmonary vein cryoisolation in patients with permanent atrial fibrillation and valvular heart disease correlation of electroanatomic mapping and long-term clinical results,” *Circulation*, vol. 111, no. 2, pp. 136–142, 2005.
- [178] Fiorenzo Gaita, Roberto Gallotti, Leonardo Calò, Eric Manasse, Riccardo Riccardi, Lucia Garberoglio, Francesco Nicolini, Marco Scaglione, Paolo Di Donna, Domenico Caponi, and Giorgio Franciosi, “Limited posterior left atrial cryoablation in patients with chronic atrial fibrillation undergoing valvular heart surgery,” *Journal of the American College of Cardiology*, vol. 36, no. 1, pp. 159–166, 2000.
- [179] Hung-Fat Tse, Sven Reek, Carl Timmermans, Kathy Lai-Fun Lee, J. Christoph Geller, Luz-Maria Rodriguez, Benoit Ghaye, Gregory M. Ayers, Harry J. G. M. Crijns, Helmut U. Klein, and Chu-Pak Lau, “Pulmonary vein isolation using transvenous catheter cryoablation for treatment of atrial fibrillation without risk of pulmonary vein stenosis,” *Journal of the American College of Cardiology*, vol. 42, no. 4, pp. 752–758, 2003.
- [180] Jason G. Andrade, Paul Khairy, Peter G. Guerra, Marc W. Deyell, Lena Rivard, Laurent Macle, Bernard Thibault, Mario Talajic, Denis Roy, and Marc Dubuc, “Efficacy and safety of cryoballoon ablation for atrial fibrillation: A systematic review of published studies,” *Heart Rhythm*, vol. 8, no. 9, pp. 1444–1451, 2011.
- [181] Alvaro V. Sarabanda, T. Jared Bunch, Susan B. Johnson, Srijoy Mahapatra, Mark A. Milton, Luiz R. Leite, G. Keith Bruce, and Douglas L. Packer, “Efficacy and safety of circumferential pulmonary vein isolation using a novel cryothermal

- balloon ablation system,” *Journal of the American College of Cardiology*, vol. 46, no. 10, pp. 1902–1912, 2005.
- [182] Vivek Y. Reddy, Petr Neuzil, Sakis Themistoclakis, Stephan B. Danik, Aldo Bonso, Antonio Rossillo, Antonio Raviele, Robert Schweikert, Sabine Ernst, Karl-Heinz Kuck, and Andrea Natale, “Visually-guided balloon catheter ablation of atrial fibrillation: Experimental feasibility and first-in-human multicenter clinical outcome,” *Circulation*, vol. 120, no. 1, pp. 12–20, 2009.
- [183] Andrea Natale, Ennio Pisano, Jeannie Shewchik, Dianna Bash, Raffaele Fanelli, Domenico Potenza, Pietro Santarelli, Robert Schweikert, Richard White, Walid Saliba, Logan Kanagaratnam, Patrick Tchou, and Michael Lesh, “First human experience with pulmonary vein isolation using a through-the-balloon circumferential ultrasound ablation system for recurrent atrial fibrillation,” *Circulation*, vol. 102, no. 16, pp. 1879–1882, 2000.
- [184] Vivek Y. Reddy, Petr Neuzil, Andre d’Avila, Margaret Laragy, Zachary J. Malchano, Stepan Kralovec, Steven J. Kim, and Jeremy N. Ruskin, “Balloon catheter ablation to treat paroxysmal atrial fibrillation: What is the level of pulmonary venous isolation?,” *Heart Rhythm*, vol. 5, no. 3, pp. 353–360, 2008.
- [185] Sanjiv M. Narayan, David E. Krummen, Kalyanam Shivkumar, Paul Clopton, Wouter-Jan Rappel, and John M. Miller, “Treatment of atrial fibrillation by the ablation of localized sources: CONFIRM (conventional ablation for atrial fibrillation with or without focal impulse and rotor modulation) trial,” *Journal of the American College of Cardiology*, vol. 60, no. 7, pp. 628–636, 2012.
- [186] David Lin, David S. Frankel, Erica S. Zado, Edward Gerstenfeld, Sanjay Dixit, David J. Callans, Michael Riley, Mathew Hutchinson, Fermin Garcia, Rupa Bala, Ralph Verdino, Joshua Cooper, and Francis E. Marchlinski, “Pulmonary vein antral isolation and nonpulmonary vein trigger ablation without additional substrate modification for treating longstanding persistent atrial fibrillation,” *Journal of Cardiovascular Electrophysiology*, vol. 23, no. 8, pp. 806–813, 2012.
- [187] Rukshen Weerasooriya, Paul Khairy, Jean Litalien, Laurent Macle, Meleze Hocini, Frederic Sacher, Nicolas Lellouche, Sebastien Knecht, Matthew Wright, Isabelle Nault, Shinsuke Miyazaki, Christophe Scavee, Jacques Clementy, Michel Haissaguerre, and Pierre Jais, “Catheter ablation for atrial fibrillation: Are results maintained at 5 years of follow-up?,” *Journal of the American College of Cardiology*, vol. 57, no. 2, pp. 160–166, 2011.
- [188] Feifan Ouyang, Roland Tilz, Julian Chun, Boris Schmidt, Erik Wissner, Thomas Zerm, Kars Neven, Bulent Köktürk, Melanie Konstantinidou, Andreas Metzner, Alexander Fuernkranz, and Karl-Heinz Kuck, “Long-term results of catheter ablation in paroxysmal atrial fibrillation: Lessons from a 5-year follow-up,” *Circulation*, vol. 122, no. 23, pp. 2368–2377, 2010.
- [189] Hippocrates, “Precepts,” in *Ancient Medicine. Airs, Waters, Places. Epidemics 1 and 3. The Oath. Precepts. Nutriment*, W. H. S. Jones, Trans., number 147 in Loeb Classical Library, pp. 303–334. Harvard University Press, Boston, MA, 1923.
- [190] S. Nazarian, D. A. Bluemke, and H. R. Halperin, “Applications of cardiac magnetic resonance in electrophysiology,” *Circulation: Arrhythmia and Electrophysiology*, vol. 2, no. 1, pp. 63–71, 2009.

- [191] Christian Mahnkopf, Troy J. Badger, Nathan S. Burgon, Marcos Daccarett, Thomas S. Haslam, Christopher T. Badger, Christopher J. McGann, Nazem Akoum, Eugene Kholmovski, Rob S. Macleod, and Nassir F. Marrouche, "Evaluation of the left atrial substrate in patients with lone atrial fibrillation using delayed-enhanced MRI: Implications for disease progression and response to catheter ablation," *Heart Rhythm*, vol. 7, no. 10, pp. 1475–1481, 2010.
- [192] Jun Dong, Timm Dickfeld, Darshan Dalal, Aamir Cheema, Chandrasekhar R. Vasamreddy, Charles A. Henrikson, Joseph E. Marine, Henry R. Halperin, Ronald D. Berger, Joao A. C. Lima, David A. Bluemke, and Hugh Calkins, "Initial experience in the use of integrated electroanatomic mapping with three-dimensional MR/CT images to guide catheter ablation of atrial fibrillation," *Journal of Cardiovascular Electrophysiology*, vol. 17, no. 5, pp. 459–466, 2006.
- [193] Anthony G. Brooks, Lauren Wilson, Pawel Kuklik, Martin K. Stiles, Bobby John, Shashidhar, Hany Dimitri, Dennis H. Lau, Ross L. Roberts-Thomson, Christopher X. Wong, Glenn D. Young, and Prashanthan Sanders, "Image integration using NavX fusion: Initial experience and validation," *Heart Rhythm*, vol. 5, no. 4, pp. 526–535, 2008.
- [194] Hua Zhong, Joan M. Lacomis, and David Schwartzman, "On the accuracy of CartoMerge for guiding posterior left atrial ablation in man," *Heart Rhythm*, vol. 4, no. 5, pp. 595–602, 2007.
- [195] S. Roujol, T. A. Basha, A. Tan, V. Khanna, R. H. Chan, M. H. Moghari, H. Rayatzadeh, J. L. Shaw, M. E. Josephson, and R. Nezafat, "Improved multimodality data fusion of late gadolinium enhancement MRI to left ventricular voltage maps in ventricular tachycardia ablation," *IEEE Transactions on Biomedical Engineering*, vol. 60, no. 5, pp. 1308–1317, 2013.
- [196] S. Umeyama, "Least-squares estimation of transformation parameters between two point patterns," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 13, no. 4, pp. 376–380, 1991.
- [197] Per-Erik Danielsson, "Euclidean distance mapping," *Computer Graphics and Image Processing*, vol. 14, no. 3, pp. 227–248, 1980.
- [198] Terry S. Yoo, Michael J. Ackerman, William E. Lorensen, Will Schroeder, Vikram Chalana, Stephen Aylward, Dimitris Metaxas, and Ross Whitaker, "Engineering and algorithm design for an image processing API: A technical report on ITK – the Insight Toolkit," *Studies in Health Technology and Informatics*, vol. 85, pp. 586–592, 2002.
- [199] A. Boldt, U. Wetzel, J. Lauschke, J. Weigl, J. Gummert, G. Hindricks, H. Kottkamp, and S. Dhein, "Fibrosis in left atrial tissue of patients with atrial fibrillation with and without underlying mitral valve disease," *Heart*, vol. 90, no. 4, pp. 400–405, 2004.
- [200] L. F. Tops and M. J. Schalij, "Left atrial fibrosis and recurrent arrhythmia," *Journal of the American Medical Association*, vol. 311, no. 22, p. 2335, 2014.
- [201] Atul Verma, Oussama M. Wazni, Nassir F. Marrouche, David O. Martin, Fethi Kilicaslan, Stephen Minor, Robert A. Schweikert, Walid Saliba, Jennifer Cummings, J. David Burkhardt, Mandeep Bhargava, William A. Belden, Ahmad Abdul-Karim,

- and Andrea Natale, “Pre-existent left atrial scarring in patients undergoing pulmonary vein antrum isolation: An independent predictor of procedural failure,” *Journal of the American College of Cardiology*, vol. 45, no. 2, pp. 285–292, 2005.
- [202] Prashanthan Sanders, Joseph B. Morton, Neil C. Davidson, Steven J. Spence, Jitendra K. Vohra, Paul B. Sparks, and Jonathan M. Kalman, “Electrical remodeling of the atria in congestive heart failure electrophysiological and electroanatomic mapping in humans,” *Circulation*, vol. 108, no. 12, pp. 1461–1468, 2003.
- [203] Prashanthan Sanders, Joseph B. Morton, Peter M. Kistler, Steven J. Spence, Neil C. Davidson, Azlan Hussin, Jitendra K. Vohra, Paul B. Sparks, and Jonathan M. Kalman, “Electrophysiological and electroanatomic characterization of the atria in sinus node disease: Evidence of diffuse atrial remodeling,” *Circulation*, vol. 109, no. 12, pp. 1514–1522, 2004.
- [204] Suraj Kapa, Benoit Desjardins, David J. Callans, Francis E. Marchlinski, and Sanjay Dixit, “Contact electroanatomic mapping derived voltage criteria for characterizing left atrial scar in patients undergoing ablation for atrial fibrillation,” *Journal of Cardiovascular Electrophysiology*, vol. 25, no. 10, pp. 1044–1052, 2014.
- [205] P. Bhagirath, A. W. M. van der Graaf, R. Karim, V. J. H. M. van Driel, H. Ramanna, K. S. Rhode, N. M. S. de Groot, and M. J. W. Götte, “Multimodality imaging for patient evaluation and guidance of catheter ablation for atrial fibrillation — Current status and future perspective,” *International Journal of Cardiology*, vol. 175, no. 3, pp. 400–408, 2014.
- [206] Carlo Pappone, Salvatore Rosanio, Giuseppe Augello, Giuseppe Gallus, Gabriele Vicedomini, Patrizio Mazzone, Simone Gulletta, Filippo Gugliotta, Alessia Pappone, Vincenzo Santinelli, Valter Tortoriello, Simone Sala, Alberto Zangrillo, Giuseppe Crescenzi, Stefano Benussi, and Ottavio Alfieri, “Mortality, morbidity, and quality of life after circumferential pulmonary vein ablation for atrial fibrillation: Outcomes from a controlled nonrandomized long-term study,” *Journal of the American College of Cardiology*, vol. 42, no. 2, pp. 185–197, 2003.
- [207] Hakan Oral, Aman Chugh, Mehmet Özyayın, Eric Good, Jackie Fortino, Sundar Sankaran, Scott Reich, Petar Igic, Darryl Elmouchi, David Tschopp, Alan Wimmer, Sujoya Dey, Thomas Crawford, Frank Pelosi, Krit Jongnarangsin, Frank Bogun, and Fred Morady, “Risk of thromboembolic events after percutaneous left atrial radiofrequency ablation of atrial fibrillation,” *Circulation*, vol. 114, no. 8, pp. 759–765, 2006.
- [208] Sarinya Puwanant, Brandon C. Varr, Kevin Shrestha, Sarah K. Hussain, W.H. Wilson Tang, Ruvin S. Gabriel, Oussama M. Wazni, Mandeep Bhargava, Walid I. Saliba, James D. Thomas, Bruce D. Lindsay, and Allan L. Klein, “Role of the CHADS2 score in the evaluation of thromboembolic risk in patients with atrial fibrillation undergoing transesophageal echocardiography before pulmonary vein isolation,” *Journal of the American College of Cardiology*, vol. 54, no. 22, pp. 2032–2039, 2009.
- [209] Dana C. Peters, John V. Wylie, Thomas H. Hauser, Kraig V. Kissinger, René M. Botnar, Vidal Essebag, Mark E. Josephson, and Warren J. Manning, “Detection of pulmonary vein and left atrial scar after catheter ablation with three-dimensional

- navigator-gated delayed enhancement MR imaging: Initial experience,” *Radiology*, vol. 243, no. 3, pp. 690–695, 2007.
- [210] Christian Sohns, Rashed Karim, James Harrison, Aruna Arujuna, Nick Linton, Richard Sennett, Hendrik Lambert, Giovanni Leo, Steven Williams, Reza Razavi, Matt Wright, Tobias Schaeffter, Mark O’neill, and Kawal Rhode, “Quantitative magnetic resonance imaging analysis of the relationship between contact force and left atrial scar formation after catheter ablation of atrial fibrillation,” *Journal of Cardiovascular Electrophysiology*, vol. 25, no. 2, pp. 138–145, 2014.
- [211] Robert S. Oakes, Troy J. Badger, Eugene G. Kholmovski, Nazem Akoum, Nathan S. Burgon, Eric N. Fish, Joshua J. E. Blauer, Swati N. Rao, Edward V. R. DiBella, Nathan M. Segerson, Marcos Daccarett, Jessiciah Windfelder, Christopher J. McGann, Dennis Parker, Rob S. MacLeod, and Nassir F. Marrouche, “Detection and quantification of left atrial structural remodeling with delayed-enhancement magnetic resonance imaging in patients with atrial fibrillation,” *Circulation*, vol. 119, no. 13, pp. 1758–1767, 2009.
- [212] N. Marrouche, “Left atrial fibrosis and recurrent arrhythmia—reply,” *Journal of the American Medical Association*, vol. 311, no. 22, pp. 2335–2336, 2014.
- [213] Nassir F. Marrouche, David Wilber, Gerhard Hindricks, Pierre Jais, Nazem Akoum, Francis Marchlinski, Eugene Kholmovski, Nathan Burgon, Nan Hu, Lluís Mont, Thomas Deneke, Mattias Duytschaever, Thomas Neumann, Moussa Mansour, Christian Mahnkopf, Bengt Herweg, Emile Daoud, Erik Wissner, Paul Bansmann, and Johannes Brachmann, “Association of atrial tissue fibrosis identified by delayed enhancement MRI and atrial fibrillation catheter ablation: The DECAAF study,” *Journal of the American Medical Association*, vol. 311, no. 5, pp. 498–506, 2014.
- [214] I. M. Khurram, R. Beinart, H. Yarmohammadi, T. Sasaki, D. D. Spragg, R. D. Berger, H. R. Halperin, H. Calkins, S. L. Zimmerman, and S. Nazarian, “Image intensity ratio, a novel magnetic resonance-based measure for quantification of left atrial fibrosis, correlates with the distribution of atrial bipolar voltage,” *Heart Rhythm*, vol. 9, no. 11, pp. 1917–1918, 2012.
- [215] Irfan M. Khurram, Roy Beinart, Vadim Zipunnikov, Jane Dewire, Hira Yarmohammadi, Takeshi Sasaki, David D. Spragg, Joseph E. Marine, Ronald D. Berger, Henry R. Halperin, Hugh Calkins, Stefan L. Zimmerman, and Saman Nazarian, “Magnetic resonance image intensity ratio, a normalized measure to enable interpatient comparability of left atrial fibrosis,” *Heart Rhythm*, vol. 11, no. 1, pp. 85–92, 2014.
- [216] Jeroen J. Bax, Nina Ajmone Marsan, and Victoria Delgado, “Non-invasive imaging in atrial fibrillation: Focus on prognosis and catheter ablation,” *Heart*, vol. 101, no. 2, pp. 94–100, 2015.
- [217] Jean-Benoît le Polain de Waroux, Anne-Catherine Pouleur, Céline Goffinet, Agnès Pasquet, Jean-Louis Vanoverschelde, and Bernhard L. Gerber, “Combined coronary and late-enhanced multidetector-computed tomography for delineation of the etiology of left ventricular dysfunction: Comparison with coronary angiography and contrast-enhanced cardiac magnetic resonance imaging,” *European Heart Journal*, vol. 29, no. 20, pp. 2544–2551, 2008.
- [218] Jennifer I. Berliner, Aya Kino, James C. Carr, Robert O. Bonow, and Lubna

- Choudhury, “Cardiac computed tomographic imaging to evaluate myocardial scarring/fibrosis in patients with hypertrophic cardiomyopathy: A comparison with cardiac magnetic resonance imaging,” *The International Journal of Cardiovascular Imaging*, vol. 29, no. 1, pp. 191–197, 2013.
- [219] Jing Tian, Jean Jeudy, Mark F. Smith, Alejandro Jimenez, Xianghua Yin, Patricia A. Bruce, Peng Lei, Aharon Turgeman, Aharon Abbo, Raj Shekhar, Magdi Saba, Stephen Shorofsky, and Timm Dickfeld, “Three-dimensional contrast-enhanced multidetector CT for anatomic, dynamic, and perfusion characterization of abnormal myocardium to guide ventricular tachycardia ablations,” *Circulation: Arrhythmia and Electrophysiology*, vol. 3, no. 5, pp. 496–504, 2010.
- [220] Leo Breiman, “Random forests,” *Machine Learning*, vol. 45, no. 1, pp. 5–32, 2001.
- [221] Kelly H. Zou, A. James O’Malley, and Laura Mauri, “Receiver-operating characteristic analysis for evaluating diagnostic tests and predictive models,” *Circulation*, vol. 115, no. 5, pp. 654–657, 2007.
- [222] M. H. Zweig and G. Campbell, “Receiver-operating characteristic (ROC) plots: A fundamental evaluation tool in clinical medicine,” *Clinical Chemistry*, vol. 39, no. 4, pp. 561–577, 1993.
- [223] Sofus Macaskassy and Foster Provost, “Confidence bands for ROC curves: Methods and an empirical study,” in *Proceedings of the First Workshop on ROC Analysis in AI*, Valencia, Spain, Aug. 2004.
- [224] Gregory Campbell, “Advances in statistical methodology for the evaluation of diagnostic and laboratory tests,” *Statistics in Medicine*, vol. 13, pp. 499–508, 1994.
- [225] I. T. Jolliffe, *Principal Component Analysis*, Springer, New York, 2nd edition, 2013.
- [226] Harpreet K. Pannu, William Alvarez, and Elliot K. Fishman, “ β -blockers for cardiac CT: A primer for the radiologist,” *American Journal of Roentgenology*, vol. 186, pp. S341–S345, 2006.
- [227] Hiroshi Hayashi, Meiso Hayashi, Yasushi Miyauchi, Kenta Takahashi, Shunsuke Uetake, Ippei Tsuboi, Kenji Yodogawa, Yu-Ki Iwasaki, and Wataru Shimizu, “Left atrial wall thickness and outcomes of catheter ablation for atrial fibrillation in patients with hypertrophic cardiomyopathy,” *Journal of Interventional Cardiac Electrophysiology: An International Journal of Arrhythmias and Pacing*, vol. 40, no. 2, pp. 153–160, 2014.
- [228] Christopher McGann, Nazem Akoum, Amit Patel, Eugene Kholmovski, Patricia Revelo, Kavitha Damal, Brent Wilson, Josh Cates, Alexis Harrison, Ravi Ranjan, Nathan S. Burgon, Tom Greene, Dan Kim, Edward V. R. DiBella, Dennis Parker, Rob S. MacLeod, and Nassir F. Marrouche, “Atrial fibrillation ablation outcome is predicted by left atrial remodeling on MRI,” *Circulation: Arrhythmia and Electrophysiology*, vol. 7, no. 1, pp. 23–30, 2014.
- [229] Fabian Pedregosa, Gaël Varoquaux, Alexandre Gramfort, Vincent Michel, Bertrand Thirion, Olivier Grisel, Mathieu Blondel, Peter Prettenhofer, Ron Weiss, Vincent Dubourg, Jake Vanderplas, Alexandre Passos, David Cournapeau, Matthieu Brucher, Matthieu Perrot, and Édouard Duchesnay, “Scikit-learn: Machine learning in Python,” *Journal of Machine Learning Research*, vol. 12, pp. 2825–2830, 2011.

- [230] Samuel Shem, *The House of God*, Dell Publishing, New York, 1979.
- [231] Christopher McGann, Eugene Kholmovski, Joshua Blauer, Sathya Vijayakumar, Thomas Haslam, Joshua Cates, Edward DiBella, Nathan Burgon, Brent Wilson, Alton Alexander, Marcel Prastawa, Marcos Daccarett, Gaston Vergara, Nazem Akoum, Dennis Parker, Rob MacLeod, and Nassir Marrouche, “Dark regions of no-reflow on late gadolinium enhancement magnetic resonance imaging result in scar formation after atrial fibrillation ablation,” *Journal of the American College of Cardiology*, vol. 58, no. 2, pp. 177–185, 2011.
- [232] Dana C. Peters, John V. Wylie, Thomas H. Hauser, Reza Nezafat, Yuchi Han, Jeong Joo Woo, Jason Taclas, Kraig V. Kissinger, Beth Goddu, Mark E. Josephson, and Warren J. Manning, “Recurrence of atrial fibrillation correlates with the extent of post-procedural late gadolinium enhancement: A pilot study,” *Journal of the American College of Cardiology: Cardiovascular Imaging*, vol. 2, no. 3, pp. 308–316, 2009.
- [233] James L. Harrison, Henrik K. Jensen, Sarah A. Peel, Amedeo Chiribiri, Anne K. Grøndal, Lars Ø Bloch, Steen F. Pedersen, Jacob F. Bentzon, Christoph Kolbitsch, Rashed Karim, Steven E. Williams, Nick W. Linton, Kawal S. Rhode, Jaswinder Gill, Michael Cooklin, C. A. Rinaldi, Matthew Wright, Won Y. Kim, Tobias Schaeffter, Reza S. Razavi, and Mark D. O’Neill, “Cardiac magnetic resonance and electroanatomical mapping of acute and chronic atrial ablation injury: A histological validation study,” *European Heart Journal*, vol. 35, no. 22, pp. 1486–1495, 2014.
- [234] Hana Hoyt, Aditya Bhonsale, Karuna Chilukuri, Fawaz Alhumaid, Matthew Needleman, David Edwards, Ashul Govil, Saman Nazarian, Alan Cheng, Charles A. Henrikson, Sunil Sinha, Joseph E. Marine, Ronald Berger, Hugh Calkins, and David D. Spragg, “Complications arising from catheter ablation of atrial fibrillation: Temporal trends and predictors,” *Heart Rhythm*, vol. 8, no. 12, pp. 1869–1874, 2011.
- [235] Roberto M. Lang, Michelle Bierig, Richard B. Devereux, Frank A. Flachskampf, Elyse Foster, Patricia A. Pellikka, Michael H. Picard, Mary J. Roman, James Seward, Jack Shanewise, Scott Solomon, Kirk T. Spencer, Martin St John Sutton, and William Stewart, “Recommendations for chamber quantification,” *European Journal of Echocardiography*, vol. 7, no. 2, pp. 79–108, 2006.
- [236] Paul V. C. Hough, “Method and means for recognizing complex patterns,” U.S. Patent 3 069 654, 18 Dec. 1962.
- [237] Michael Kass, Andrew Witkin, and Demetri Terzopoulos, “Snakes: Active contour models,” *International Journal of Computer Vision*, vol. 1, no. 4, pp. 321–331, 1988.
- [238] A. Mishra, P. K. Dutta, and M. K. Ghosh, “A GA based approach for boundary detection of left ventricle with echocardiographic image sequences,” *Image and Vision Computing*, vol. 21, no. 11, pp. 967–976, 2003.
- [239] M. Mignotte and J. Meunier, “A multiscale optimization approach for the dynamic contour-based boundary detection issue,” *Computerized Medical Imaging and Graphics*, vol. 25, no. 3, pp. 265–275, 2001.

- [240] Auzuir Ripardo de Alexandria, Paulo César Cortez, Jessyca Almeida Bessa, John Hebert da Silva Félix, José Sebastião de Abreu, and Victor Hugo C. de Albuquerque, “pSnakes: A new radial active contour model and its application in the segmentation of the left ventricle from echocardiographic images,” *Computer Methods and Programs in Biomedicine*, vol. 116, no. 3, pp. 260–273, 2014.
- [241] Daniel Barbosa, Thomas Dietenbeck, Brecht Heyde, Helene Houle, Denis Friboulet, Jan D’hooge, and Olivier Bernard, “Fast and fully automatic 3-D echocardiographic segmentation using B-spline explicit active surfaces: Feasibility study and validation in a clinical setting,” *Ultrasound in Medicine & Biology*, vol. 39, no. 1, pp. 89–101, 2013.
- [242] S. Lankton and A. Tannenbaum, “Localizing region-based active contours,” *IEEE Transactions on Image Processing*, vol. 17, no. 11, pp. 2029–2039, 2008.
- [243] Jr. Yezzi, A., A. Tsai, and A. Willsky, “A statistical approach to snakes for bimodal and trimodal imagery,” in *The Proceedings of the Seventh IEEE International Conference on Computer Vision, 1999*, Kerkyra, Greece, 1999, vol. 2, pp. 898–903 vol.2.
- [244] Stanley Osher and James A Sethian, “Fronts propagating with curvature-dependent speed: Algorithms based on Hamilton-Jacobi formulations,” *Journal of Computational Physics*, vol. 79, no. 1, pp. 12–49, 1988.
- [245] David Adalsteinsson and James A. Sethian, “A fast level set method for propagating interfaces,” *Journal of Computational Physics*, vol. 118, no. 2, pp. 269–277, 1995.
- [246] Ross T. Whitaker, “A level-set approach to 3d reconstruction from range data,” *International Journal of Computer Vision*, vol. 29, no. 3, pp. 203–231, 1998.
- [247] Chohong Min and Frédéric Gibou, “A second order accurate level set method on non-graded adaptive cartesian grids,” *Journal of Computational Physics*, vol. 225, no. 1, pp. 300–321, 2007.
- [248] Ben Houston, Michael B. Nielsen, Christopher Batty, Ola Nilsson, and Ken Museth, “Hierarchical RLE level set: A compact and versatile deformable surface representation,” *ACM Trans. Graph.*, vol. 25, no. 1, pp. 151–175, 2006.
- [249] JiaYong Yan and TianGe Zhuang, “Applying improved fast marching method to endocardial boundary detection in echocardiographic images,” *Pattern Recognition Letters*, vol. 24, no. 15, pp. 2777–2784, 2003.
- [250] Ning Lin, Weichuan Yu, and James S. Duncan, “Combinative multi-scale level set framework for echocardiographic image segmentation,” *Medical Image Analysis*, vol. 7, no. 4, pp. 529–537, 2003.
- [251] F. Galluzzo, D. Barbosa, H. Houle, N. Speciale, D. Friboulet, J. D’hooge, and O. Bernard, “A GPU level-set segmentation framework for 3d echocardiography,” in *2012 IEEE International Ultrasonics Symposium*, Dresden, Germany, Oct. 2012, pp. 2639–2642.
- [252] Kashif Rajpoot, Vicente Grau, J. Alison Noble, Harald Becher, and Cezary Szmigielski, “The evaluation of single-view and multi-view fusion 3d echocardiography using image-driven segmentation and tracking,” *Medical Image Analysis*, vol. 15, no. 4, pp. 514–528, 2011.

- [253] Tobias Heimann and Hans-Peter Meinzer, “Statistical shape models for 3d medical image segmentation: A review,” *Medical Image Analysis*, vol. 13, no. 4, pp. 543–563, 2009.
- [254] Max Mignotte, Jean Meunier, and Jean-Claude Tardif, “Endocardial boundary estimation and tracking in echocardiographic images using deformable template and Markov random fields,” *Pattern Analysis & Applications*, vol. 4, no. 4, pp. 256–271, 2001.
- [255] Djamel Boukerroui, Atilla Baskurt, J. Alison Noble, and Olivier Basset, “Segmentation of ultrasound images—multiresolution 2d and 3d algorithm based on global and local statistics,” *Pattern Recognition Letters*, vol. 24, no. 4–5, pp. 779–790, 2003.
- [256] T. Binder, M. Süssner, D. Moertl, T. Strohmer, H. Baumgartner, G. Maurer, and G. Porenta, “Artificial neural networks and spatial temporal contour linking for automated endocardial contour detection on echocardiograms: a novel approach to determine left ventricular contractile function,” *Ultrasound in Medicine & Biology*, vol. 25, no. 7, pp. 1069–1076, 1999.
- [257] T. F. Cootes, C. J. Taylor, D. H. Cooper, and J. Graham, “Active shape models—Their training and application,” *Computer Vision and Image Understanding*, vol. 61, no. 1, pp. 38–59, 1995.
- [258] Yunmei Chen, Feng Huang, Hemant D. Tagare, and Murali Rao, “A coupled minimization problem for medical image segmentation with priors,” *International Journal of Computer Vision*, vol. 71, no. 3, pp. 259–272, 2006.
- [259] Richard V. Stebbing and J. Alison Noble, “Delineating anatomical boundaries using the boundary fragment model,” *Medical Image Analysis*, vol. 17, no. 8, pp. 1123–1136, 2013.
- [260] Richard V. Stebbing, Ana I. L. Namburete, Ross Upton, Paul Leeson, and J. Alison Noble, “Data-driven shape parameterization for segmentation of the right ventricle from 3d+t echocardiography,” *Medical Image Analysis*, vol. 21, no. 1, pp. 29–39, 2015.
- [261] J. A. Noble and D. Boukerroui, “Ultrasound image segmentation: A survey,” *IEEE Transactions on Medical Imaging*, vol. 25, no. 8, pp. 987–1010, 2006.
- [262] Yun Zhu, Xenophon Papademetris, Albert J. Sinusas, and James S. Duncan, “A coupled deformable model for tracking myocardial borders from real-time echocardiography using an incompressibility constraint,” *Medical Image Analysis*, vol. 14, no. 3, pp. 429–448, 2010.
- [263] Andriy Myronenko, Xubo Song, and David J. Sahn, “LV motion tracking from 3d echocardiography using textural and structural information,” in *Proceedings of the 10th International Conference on Medical Image Computing and Computer-Assisted Intervention*, Berlin, Heidelberg, 2007, MICCAI’07, pp. 428–435, Springer-Verlag.
- [264] D. H. Ballard, “Generalizing the Hough transform to detect arbitrary shapes,” *Pattern Recognition*, vol. 13, no. 2, pp. 111–122, 1981.
- [265] P. V. C. Hough, “Machine analysis of bubble chamber pictures,” in *International Conference on High Energy Accelerators and Instrumentation*, L. Kowarski, Ed.,

- Geneva, Sept. 1959, pp. 554–556, Organisation Européenne pour la Recherche Nucléaire.
- [266] P. Bansod, U. B. Desai, and N. Burkule, “Automatic detection of left ventricle in echocardiographic sequences using radial search and temporal smoothing,” in *15th International Conference on Digital Signal Processing*, Cardiff, UK, July 2007, pp. 139–142, IEEE.
 - [267] Yongjian Yu and Scott T. Acton, “Speckle reducing anisotropic diffusion,” *IEEE Transactions on Image Processing*, vol. 11, no. 11, pp. 1260–1270, 2002.
 - [268] Antonio Fernández-Caballero and José M. Vega-Riesco, “Determining heart parameters through left ventricular automatic segmentation for heart disease diagnosis,” *Expert Systems with Applications*, vol. 36, no. 2, pp. 2234–2249, 2009.
 - [269] Kashif Rajpoot, Alison Noble, Vicente Grau, and Nasir Mahmood Rajpoot, “Feature detection from echocardiography images using local phase information,” in *12th Medical Image Understanding and Analysis*, Dundee, Scotland, UK, 2008.
 - [270] Keith Langley and Stephen J. Anderson, “The Riesz transform and simultaneous representations of phase, energy and orientation in spatial vision,” *Vision Research*, vol. 50, no. 17, pp. 1748–1765, 2010.
 - [271] Lee R. Dice, “Measures of the amount of ecologic association between species,” *Ecology*, vol. 26, no. 3, pp. 297–302, 1945.
 - [272] Kelly H. Zou, Simon K. Warfield, Aditya Bharatha, Clare M. C. Tempany, Michael R. Kaus, Steven J. Haker, William M. Wells, Ferenc A. Jolesz, and Ron Kikinis, “Statistical validation of image segmentation quality based on a spatial overlap index,” *Academic Radiology*, vol. 11, no. 2, pp. 178–189, 2004.
 - [273] Arthur Conan Doyle, *The Valley of Fear*, George H. Doran Company, New York, 1915.
 - [274] R. Gonzalez, M. Scheinman, W. Margaretten, and M. Rubinstein, “Closed-chest electrode-catheter technique for His bundle ablation in dogs,” *American Journal of Physiology - Heart and Circulatory Physiology*, vol. 241, no. 2, pp. H283–H287, 1981.
 - [275] Bruno P. Valdigem, Fabio B. F. C. G. Pereira, Nilton J. Carneiro da Silva, Cristiano O. Dietrich, Ricardo Sobral, Fernando Lopes Nogueira, Roberto C. Berber, Fabricio Mallman, Ibraim M. Pinto, Gilberto Szarf, Claudio Cirenza, and Angelo A. V. de Paola, “Ablation of ventricular tachycardia in chronic chagasic cardiomyopathy with giant basal aneurysm Carto Sound, CT, and MRI merge,” *Circulation: Arrhythmia and Electrophysiology*, vol. 4, no. 1, pp. 112–114, 2011.
 - [276] Emanuele Bertaglia, Paolo Della Bella, Claudio Tondo, Alessandro Proclemer, Nicola Bottoni, Roberto De Ponti, Maurizio Landolina, Maria Grazia Bongiorni, Leonardo Corò, Giuseppe Stabile, Antonio Dello Russo, Roberto Verlato, Massimo Mantica, and Franco Zoppo, “Image integration increases efficacy of paroxysmal atrial fibrillation catheter ablation: Results from the CartoMerge Italian Registry,” *Europace*, vol. 11, no. 8, pp. 1004–1010, 2009.
 - [277] Anselma Intini, Robert N. Goldstein, Ping Jia, Charulatha Ramanathan, Kyungmoo Ryu, Bartolomeo Giannattasio, Robert Gilkeson, Bruce S. Stambler, Pedro Brugada, William G. Stevenson, Yoram Rudy, and Albert L. Waldo,

- “Electrocardiographic imaging (ECGI), a novel diagnostic modality used for mapping of focal left ventricular tachycardia in a young athlete,” *Heart Rhythm*, vol. 2, no. 11, pp. 1250–1252, 2005.
- [278] Martyn P. Nash and Andrew J. Pullan, “Challenges facing validation of noninvasive electrical imaging of the heart,” *Annals of Noninvasive Electrocardiology*, vol. 10, no. 1, pp. 73–82, 2005.
- [279] K. Will, J. Krug, K. Jungnickel, F. Fischbach, J. Ricke, G. Rose, and A. Omar, “MR-compatible RF ablation system for online treatment monitoring using MR thermometry,” in *2010 Annual International Conference of the IEEE Engineering in Medicine and Biology Society*, Aug. 2010, pp. 1601–1604.
- [280] Baudouin Denis de Senneville, Sébastien Roujol, Pierre Jaïs, Chrit T. W. Moonen, Gwenaél Herigault, and Bruno Quesson, “Feasibility of fast MR-thermometry during cardiac radiofrequency ablation,” *NMR in Biomedicine*, vol. 25, no. 4, pp. 556–562, 2012.
- [281] Aravindan Kolandaivelu, Menekhem M. Zviman, Valeria Castro, Albert C. Lardo, Ronald D. Berger, and Henry R. Halperin, “Noninvasive assessment of tissue heating during cardiac radiofrequency ablation using MRI thermography,” *Circulation: Arrhythmia and Electrophysiology*, vol. 3, no. 5, pp. 521–529, 2010.
- [282] D. Augustine, M. Yaqub, C. Szmigielski, E. Lima, S. Neubauer, S. E. Petersen, H. Becher, A. Noble, and P. Leeson, “3d fusion echocardiography improves 3d left ventricular assessment: comparison with 2d contrast echocardiography,” *European Heart Journal*, vol. 34, no. suppl 1, p. 5864, 2013.
- [283] William M. Wells III, Paul Viola, Hideki Atsumi, Shin Nakajima, and Ron Kikinis, “Multi-modal volume registration by maximization of mutual information,” *Medical Image Analysis*, vol. 1, no. 1, pp. 35–51, 1996.
- [284] F. Maes, A. Collignon, D. Vandermeulen, G. Marchal, and P. Suetens, “Multimodality image registration by maximization of mutual information,” *IEEE Transactions on Medical Imaging*, vol. 16, no. 2, pp. 187–198, 1997.
- [285] G. P. Penney, J. Weese, J. A. Little, P. Desmedt, D. L. G. Hill, and D. J. Hawkes, “A comparison of similarity measures for use in 2-D-3-D medical image registration,” *IEEE Transactions on Medical Imaging*, vol. 17, no. 4, pp. 586–595, 1998.
- [286] André Collignon, Dirk Vandermeulen, Paul Suetens, and Guy Marchal, “Registration of 3d multi-modality medical images using surfaces and point landmarks,” *Pattern Recognition Letters*, vol. 15, no. 5, pp. 461–467, 1994.
- [287] L.-Y. Hsu and M. H. Loew, “Fully automatic 3d feature-based registration of multi-modality medical images,” *Image and Vision Computing*, vol. 19, pp. 75–85, 2001.
- [288] W. Cheung and G. Hamarneh, “N-SIFT: N-dimensional scale invariant feature transform for matching medical images,” in *4th IEEE International Symposium on Biomedical Imaging: From Nano to Macro, ISBI 2007*, Apr. 2007, pp. 720–723.
- [289] Josien Pluim, J. Maintz, and Max Viergever, “Image registration by maximization of combined mutual information and gradient information,” in *Medical Image Computing and Computer-Assisted Intervention – MICCAI 2000*, Scott Delp,

- Anthony DiGoia, and Branislav Jaramaz, Eds., vol. 1935 of *Lecture Notes in Computer Science*, pp. 103–129. Springer Berlin / Heidelberg, 2000.
- [290] Dinggang Shen, “Image registration by local histogram matching,” *Pattern Recognition*, vol. 40, no. 4, pp. 1161–1172, 2007.
- [291] Wolfgang Wein, Shelby Brunke, Ali Khamene, Matthew R. Callstrom, and Nassir Navab, “Automatic CT-ultrasound registration for diagnostic imaging and image-guided intervention,” *Medical Image Analysis*, vol. 12, no. 5, pp. 577–585, 2008.
- [292] Marius George Linguraru, Nikolay V. Vasilyev, Pedro J. Del Nido, and Robert D. Howe, “Statistical segmentation of surgical instruments in 3-D ultrasound images,” *Ultrasound in Medicine & Biology*, vol. 33, no. 9, pp. 1428–1437, 2007.
- [293] Paul M. Novotny, Jeff A. Stoll, Nikolay V. Vasilyev, Pedro J. del Nido, Pierre E. Dupont, Todd E. Zickler, and Robert D. Howe, “GPU based real-time instrument tracking with three-dimensional ultrasound,” *Medical Image Analysis*, vol. 11, no. 5, pp. 458–464, 2007.
- [294] Cristian A. Linte, John Moore, Chris Wedlake, and Terry M. Peters, “Evaluation of model-enhanced ultrasound-assisted interventional guidance in a cardiac phantom,” *IEEE Transactions on Bio-Medical Engineering*, vol. 57, no. 9, pp. 2209–2218, 2010.
- [295] Michel Foucault, *Dits et Écrits: 1976-1988*, Daniel Defert and François Ewald, Eds., vol. 2, Gillimard, Paris, 2nd edition, 2001.
- [296] K Doi, “Current status and future potential of computer-aided diagnosis in medical imaging,” *British Journal of Radiology*, vol. 78, pp. s3–s19, 2005.
- [297] William Harvey, *Exercitatio Anatomica de Motu Cordis et Sanguinis in Animalibus*, Sumptibus Guilielmi Fitzeri, Frankfort, 1st edition, 1628.
- [298] John E. Hall, *Guyton and Hall Textbook of Medical Physiology*, Saunders, Philadelphia, PA, 12th edition, 2010.
- [299] Walter F. Boron and Emile L. Boulpaep, Eds., *Medical Physiology*, Saunders, Philadelphia, PA, 2nd edition, 2011.
- [300] Dee Unglaub Silverthorn, *Human Physiology: An Integrated Approach*, Pearson, Boston, 6th edition, 2012.
- [301] Luigi Di Biase, Pasquale Santangeli, Matteo Anselmino, Prasant Mohanty, Ilaria Salvetti, Sebastiano Gili, Rodney Horton, Javier E. Sanchez, Rong Bai, Sanghamitra Mohanty, Agnes Pump, Mauricio Cereceda Brantes, G. Joseph Gallinhouse, J. David Burkhardt, Federico Cesarani, Marco Scaglione, Andrea Natale, and Fiorenzo Gaita, “Does the left atrial appendage morphology correlate with the risk of stroke in patients with atrial fibrillation? Results from a multicenter study,” *Journal of the American College of Cardiology*, vol. 60, no. 6, pp. 531–538, 2012.
- [302] Takehiro Kimura, Seiji Takatsuki, Kohei Inagawa, Yoshinori Katsumata, Takahiko Nishiyama, Nobuhiro Nishiyama, Kotaro Fukumoto, Yoshiyasu Aizawa, Yoko Tanimoto, Kojiro Tanimoto, Masahiro Jinzaki, and Keiichi Fukuda, “Anatomical characteristics of the left atrial appendage in cardiogenic stroke with low CHADS2 scores,” *Heart Rhythm*, vol. 10, no. 6, pp. 921–925, 2013.

- [303] Sumit Verma, “Left atrial appendage and thromboembolism: Shape matters,” *Heart Rhythm*, vol. 10, no. 6, pp. 926–927, 2013.
- [304] N. M. Al-Saady, O. A. Obel, and A. J. Camm, “Left atrial appendage: Structure, function, and role in thromboembolism,” *Heart*, vol. 82, no. 5, pp. 547–554, 1999.
- [305] Blausen Medical, “Coronary circulation” [images]. Accessed 3 Jan 2014: http://commons.wikimedia.org/wiki/File:Blausen_0260_CoronaryVessels_Anterior.png and *ibid*/File:Blausen_0261_CoronaryVessels_Posterior.png, under the Creative Commons Attribution 3.0 Unported license.
- [306] Wikimedia Commons, “Wiggers diagram” [image]. Accessed 2 Jan 2014: http://commons.wikimedia.org/wiki/File:Wiggers_Diagram.svg, under the Creative Commons Attribution-ShareAlike 2.5 Generic license.
- [307] OpenStax College, “Cardiac Muscle and Electrical Activity,” in *Anatomy & Physiology*, 2013, accessed 4 Jan 2014: <http://cnx.org/contents/302812e9-2d2d-4e44-8075-4bb75db53f36@3>, under the Creative Commons Attribution 4.0 license.
- [308] Russel Friend and Garrett Lerner (Writers), Dan Attias (Director), “Acceptance,” Television Episode In: Paul Attanasio (Executive Producer), *House MD*, series 2, ep. 1, Fox Network, Los Angeles, CA, 20 Sept 2005.
- [309] “Heart disease and stroke statistics – 2009 update,” Tech. Rep., American Heart Association, Dallas, Texas, 2009.
- [310] Nick Townsend, Kremlin Wickramasinghe, Prachi Bhatnagar, Kate Smolina, Mel Nichols, Jose Leal, Ramon Luengo-Fernandez, and Mike Rayner, *Coronary Heart Disease Statistics 2012 Edition*, Peter Weissberg, Ed., British Heart Foundation, London, 2012.
- [311] Allen P. Burke and Renu Virmani, “Pathophysiology of acute myocardial infarction,” *Medical Clinics of North America*, vol. 91, no. 4, pp. 553–572, 2007.
- [312] J. Magne, M. Senechal, J. G. Dumesnil, and P. Pibarot, “Ischemic mitral regurgitation: A complex multifaceted disease,” *Cardiology*, vol. 112, no. 4, pp. 244–259, 2009.
- [313] K. M. Chan, E. Amirak, M. Zakkar, M. Flather, J. R. Pepper, and P. P. Punjabi, “Ischemic mitral regurgitation: In search of the best treatment for a common condition,” *Progress in Cardiovascular Diseases*, vol. 51, no. 6, pp. 460–471, 2009.
- [314] H. M. Haqqani and F. E. Marchlinski, “Electrophysiologic substrate underlying postinfarction ventricular tachycardia: Characterization and role in catheter ablation,” *Heart Rhythm*, vol. 6, no. 8 Suppl, pp. S70–S76, 2009.
- [315] R. J. Myerburg, “Implantable cardioverter-defibrillators after myocardial infarction,” *New England Journal of Medicine*, vol. 359, no. 21, pp. 2245–2253, 2008.
- [316] Vuyisile T. Nkomo, Julius M. Gardin, Thomas N. Skelton, John S. Gottdiener, Christopher G. Scott, and Maurice Enriquez-Sarano, “Burden of valvular heart diseases: A population-based study,” *The Lancet*, vol. 368, no. 9540, pp. 1005–1011, 2006.

- [317] “Rheumatic fever and rheumatic heart disease,” *World Health Organization Technical Report Service*, vol. 923, pp. 1–122, 2004.
- [318] S. Stewart, K. MacIntyre, S. Capewell, and J. J. V. McMurray, “Heart failure and the aging population: An increasing burden in the 21st century?,” *Heart*, vol. 89, no. 1, pp. 49–53, 2003.
- [319] Amir M. Sheikh and Steven A. Livesey, “Surgical management of valve disease in the early 21st century,” *Clinical Medicine*, vol. 10, no. 2, pp. 177–181, 2010.
- [320] R. J. Cubeddu and I. F. Palacios, “Percutaneous techniques for mitral valve disease,” *Cardiology Clinics*, vol. 28, no. 1, pp. 139–153, 2010.
- [321] W. T. Brinkman and M. J. Mack, “Transcatheter cardiac valve interventions,” *Surgical Clinics of North America*, vol. 89, no. 4, pp. 951–966, 2009.
- [322] Antoni Bayés de Luna, *Clinical Arrhythmology*, Wiley-Blackwell, Chichester, West Sussex, UK and Hoboken, NJ, USA, 1st edition, 2011.
- [323] Alan S. Go, Elaine M. Hylek, Kathleen A. Phillips, YuChiao Chang, Lori E. Henault, Joe V. Selby, and Daniel E. Singer, “Prevalence of diagnosed atrial fibrillation in adults: National implications for rhythm management and stroke prevention: the AnTicoagulation and Risk factors In Atrial fibrillation (ATRIA) study,” *Journal of the American Medical Association*, vol. 285, no. 18, pp. 2370–2375, 2001.
- [324] Elad Anter, Mariell Jessup, and David J. Callans, “Atrial fibrillation and heart failure: Treatment considerations for a dual epidemic,” *Circulation*, vol. 119, no. 18, pp. 2516–2525, 2009.
- [325] Pierre Jais, Michel Haissaguerre, Dipen C. Shah, Salah Chouairi, Laurent Gencel, Meleze Hocini, and Jacques Clementy, “A focal source of atrial fibrillation treated by discrete radiofrequency ablation,” *Circulation*, vol. 95, no. 3, pp. 572–576, 1997.
- [326] Riccardo Cappato, Hugh Calkins, Shih-Ann Chen, Wyn Davies, Yoshito Iesaka, Jonathan Kalman, You-Ho Kim, George Klein, Andrea Natale, Douglas Packer, Allan Skanes, Federico Ambrogi, and Elia Biganzoli, “Updated worldwide survey on the methods, efficacy, and safety of catheter ablation for human atrial fibrillation: Clinical perspective,” *Circulation: Arrhythmia and Electrophysiology*, vol. 3, no. 1, pp. 32–38, 2010.
- [327] N. Saoudi, F. Cosío, A. Waldo, S. A. Chen, Y. Iesaka, M. Lesh, S. Saksena, J. Salerno, and W. Schoels, “A classification of atrial flutter and regular atrial tachycardia according to electrophysiological mechanisms and anatomical bases,” *European Heart Journal*, vol. 22, no. 14, pp. 1162–1182, 2001.
- [328] Michael D. Lesh and George F. Van Hare, “Status of ablation in patients with atrial tachycardia and flutter,” *Pacing and Clinical Electrophysiology*, vol. 17, no. 5, pp. 1026–1033, 1994.
- [329] Aman Chugh, Hakan Oral, Eric Good, Jihn Han, Kamala Tamirisa, Kristina Lemola, Darryl Elmouchi, David Tschopp, Scott Reich, Petar Igic, Frank Bogun, JR Pelosi, Frank, and Fred Morady, “Catheter ablation of atypical atrial flutter and atrial tachycardia within the coronary sinus after left atrial ablation for atrial fibrillation,” *Journal of the American College of Cardiology*, vol. 46, no. 1, pp. 83–91, 2005.

- [330] Wilbert S. Aronow, Stanley Epstein, Mordecai Koenigsberg, and Kenneth S. Schwartz, "Usefulness of echocardiographic left ventricular hypertrophy, ventricular tachycardia and complex ventricular arrhythmias in predicting ventricular fibrillation or sudden cardiac death in elderly patients," *The American Journal of Cardiology*, vol. 62, no. 16, pp. 1124–1125, 1988.
- [331] J. O. O'Neill, R. C. Starling, Y. Khaykin, P. M. McCarthy, J. B. Young, M. Hail, N. M. Albert, N. Smedira, and M. K. Chung, "Residual high incidence of ventricular arrhythmias after left ventricular reconstructive surgery," *Journal of Thoracic and Cardiovascular Surgery*, vol. 130, no. 5, pp. 1250–1256, 2005.
- [332] William G. Stevenson, "Ventricular scars and ventricular tachycardia," *Transactions of the American Clinical and Climatological Association*, vol. 120, pp. 403–412, 2009.
- [333] Fred Kusumoto, "A comprehensive approach to management of ventricular arrhythmias," *Cardiology Clinics*, vol. 26, no. 3, pp. 481–496, 2008.
- [334] Erik Wissner, William G. Stevenson, and Karl-Heinz Kuck, "Catheter ablation of ventricular tachycardia in ischaemic and non-ischaemic cardiomyopathy: Where are we today?," *European Heart Journal*, vol. 33, no. 12, pp. 1440–1450, 2012.
- [335] Stuart J. Connolly, Michael Gent, Robin S. Roberts, Paul Dorian, Denis Roy, Robert S. Sheldon, L. Brent Mitchell, Martin S. Green, George J. Klein, and Bernard O'Brien, "Canadian implantable defibrillator study (CIDS): A randomized trial of the implantable cardioverter defibrillator against amiodarone," *Circulation*, vol. 101, no. 11, pp. 1297–1302, 2000.
- [336] George J. Klein, Thomas M. Bashore, T. D. Sellers, Edward L. C. Pritchett, William M. Smith, and John J. Gallagher, "Ventricular fibrillation in the Wolff-Parkinson-White syndrome," *New England Journal of Medicine*, vol. 301, no. 20, pp. 1080–1085, 1979.
- [337] R. W. Campbell, R. A. Smith, J. J. Gallagher, E. L. Pritchett, and A. G. Wallace, "Atrial fibrillation in the preexcitation syndrome," *The American Journal of Cardiology*, vol. 40, no. 4, pp. 514–520, 1977.
- [338] Takashi Hamada, Tatsuro Hiraki, Hisao Ikeda, Ichiro Kubara, Teruhisa Yoshida, Masanobu Ohga, and Tsutomu Imaizumi, "Mechanisms for atrial fibrillation in patients with Wolff-Parkinson-White syndrome," *Journal of Cardiovascular Electrophysiology*, vol. 13, no. 3, pp. 223–229, 2002.
- [339] R. K. Thakur, G. J. Klein, and R. Yee, "Radiofrequency catheter ablation in patients with Wolff-Parkinson-White syndrome.," *Canadian Medical Association Journal*, vol. 151, no. 6, pp. 771–776, 1994.
- [340] Kevin F. Kwaku and Mark E. Josephson, "Typical AVNRT: An update on mechanisms and therapy," *Cardiac Electrophysiology Review*, vol. 6, no. 4, pp. 414–421, 2002.
- [341] J. J. Langberg, M. C. Chin, M. Rosenqvist, J. Cockrell, N. Dullet, G. Van Hare, J. C. Griffin, and M. M. Scheinman, "Catheter ablation of the atrioventricular junction with radiofrequency energy," *Circulation*, vol. 80, no. 6, pp. 1527–1535, 1989.

- [342] James H. Baker II, Vance J. Plumb, Andrew E. Epstein, and G. Neal Kay, “Predictors of recurrent atrioventricular nodal reentry after selective slow pathway ablation,” *The American Journal of Cardiology*, vol. 73, no. 11, pp. 765–769, 1994.
- [343] Gabriel Gregoratos, Melvin D. Cheitlin, Alicia Conill, Andrew E. Epstein, Christopher Fellows, T. Bruce Ferguson, Roger A. Freedman, Mark A. Hlatky, Gerald V. Naccarelli, Sanjeev Saksena, Robert C. Schlant, and Michael J. Silka, “ACC/AHA guidelines for implantation of cardiac pacemakers and antiarrhythmia devices: Executive summary,” *Circulation*, vol. 97, no. 13, pp. 1325–1335, 1998.
- [344] Michele Brignole, Angelo Auricchio, Gonzalo Baron-Esquivias, Pierre Bordachar, Giuseppe Boriani, Ole-A. Breithardt, John Cleland, Jean-Claude Deharo, Victoria Delgado, Perry M. Elliott, Bulent Gorenek, Carsten W. Israel, Christophe Leclercq, Cecilia Linde, Lluís Mont, Luigi Padeletti, Richard Sutton, and Panos E. Vardas, “2013 ESC Guidelines on cardiac pacing and cardiac resynchronization therapy,” *Europace*, vol. 15, no. 8, pp. 1070–1118, 2013.
- [345] Wilhelm Conrad Röntgen, “Über eine neue Art von Strahlen,” *Annalen der Physik*, vol. 300, no. 1, pp. 1–11, 1898.
- [346] Robert Detrano, Andras Janosi, Walter Steinbrunn, Matthias Pfisterer, Johann-Jakob Schmid, Sarbjit Sandhu, Kern H. Guppy, Stella Lee, and Victor Froelicher, “International application of a new probability algorithm for the diagnosis of coronary artery disease,” *The American Journal of Cardiology*, vol. 64, no. 5, pp. 304–310, 1989.
- [347] Andras Janosi, William Steinbrunn, Matthias Pfisterer, and Robert Detrano, “Heart disease data set,” *UCI Machine Learning Repository*, K. Bache and M. Lichman, Eds., University of California, Irvine, CA, USA, 1988, Accessed 12 Jan 2014: <http://archive.ics.uci.edu/ml/datasets/Heart+Disease>
- [348] Leo Breiman, Jerome Friedman, Charles J. Stone, and R. A. Olshen, *Classification and Regression Trees*, Chapman and Hall/CRC, New York, NY, USA, 1st edition, 1984.
- [349] Wikipedia contributors, “Athletics at the 2012 Summer Olympics,” *Wikipedia, the Free Encyclopedia*, Accessed 8 Jan 2014: http://en.wikipedia.org/wiki/Athletics_at_the_2012_Summer_Olympics.
- [350] Alan Turing, “Intelligent machinery,” in *Cybernetics: Key Papers*, A. D. J. Robertson and C. R. Evans, Eds. University Park Press, Baltimore, MD, USA and Manchester, UK, 1968.
- [351] Isaiah Berlin, *The Hedgehog and the Fox: An Essay on Tolstoy’s View of History*, Henry Hardy, Ed., Princeton University Press, 2nd edition, 2013.