Computational Methods for the Estimation of Cardiac Electrophysiological Conduction Parameters in a Patient Specific Setting

Mikael Wallman
Wolfson College



Computational Biology Research Group
Department of Computer Science
University of Oxford

Hilary Term 2013

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

Mikael Wallman
Wolfson College

Doctor of Philosophy
Hilary Term 2013

Computational Methods for the Estimation of Cardiac Electrophysiological Conduction Parameters in a Patient Specific Setting

Abstract

Cardiovascular disease is the primary cause of death globally. Although this group encompasses a heterogeneous range of conditions, many of these diseases are associated with abnormalities in the cardiac electrical propagation. In these conditions, structural abnormalities in the form of scars and fibrotic tissue are known to play an important role, leading to a high individual variability in the exact disease mechanisms. Because of this, clinical interventions such as ablation therapy and CRT that work by modifying the electrical propagation should ideally be optimized on a patient specific basis. As a tool for optimizing these interventions, computational modelling and simulation of the heart have become increasingly important. However, in order to construct these models, a crucial step is the estimation of tissue conduction properties, which have a profound impact on the cardiac activation sequence predicted by simulations. Information about the conduction properties of the cardiac tissue can be gained from electrophysiological data, obtained using electroanatomical mapping systems. However, as in other clinical modalities, electrophysiological data are often sparse and noisy, and this results in high levels of uncertainty in the estimated quantities.

In this dissertation, we develop a methodology based on Bayesian inference, together with a computationally efficient model of electrical propagation to achieve two main aims: 1) to quantify values and associated uncertainty for different tissue conduction properties inferred from electroanatomical data, and 2) to design strategies to optimise the location and number of measurements required to maximise information and reduce uncertainty. The methodology is validated in several studies performed using simulated data obtained from image-based ventricular models, including realistic fibre orientation and conduction heterogeneities. Subsequently, by using the developed methodology to investigate how the uncertainty decreases in response to added measurements, we derive an *a priori* index for placing electrophysiological measurements in order to optimise the information content of the collected data. Results show that the derived index has a clear benefit in minimising the uncertainty of inferred conduction properties compared to a random distribution of measurements, suggesting that the methodology presented in this dissertation provides an important step towards improving the quality of the spatiotemporal information obtained using electroanatomical mapping.

Contents

1 Introduction **1**
- 1.1 Motivation . . . 1
- 1.2 Goal . . . 4
- 1.3 Dissertation Outline . . . 5

2 Literature Review **8**
- 2.1 Cardiac Anatomy . . . 9
- 2.2 Cardiac Electrophysiology . . . 11
- 2.3 Normal Myocardial Conduction . . . 14
- 2.4 Myocardial Scars and Fibrosis . . . 16
- 2.5 Electroanatomical Mapping . . . 18
  - 2.5.1 Non-Contact Mapping with EnSite . . . 19
  - 2.5.2 EnSite Accuracy . . . 20
  - 2.5.3 Contact Mapping with CARTO . . . 22
  - 2.5.4 CARTO Accuracy . . . 23
  - 2.5.5 Contact Mapping with EnSite NavX . . . 24
  - 2.5.6 NavX Accuracy . . . 25
  - 2.5.7 Integration of electroanatomical data with CT/MRI using NavX Fusion and Cartomerge . . . 26
- 2.6 Computational Models of Cardiac Electrophysiology . . . 29
- 2.7 Estimation of Cardiac Conduction Parameters . . . 32
- 2.8 Conclusions . . . 36

3 Computational Methods **38**
- 3.1 Models of Cardiac Electrical Conduction . . . 39
  - 3.1.1 The Bidomain Model . . . 39
  - 3.1.2 The Monodomain Model . . . 40
  - 3.1.3 Numerical Solution Methods for the Monodomain and Bidomain Models 42
  - 3.1.4 The Eikonal Equations . . . 43
  - 3.1.5 Numerical Solution Methods for the Eikonal Equation . . . 46
- 3.2 Parameter Estimation . . . 52
  - 3.2.1 The Parameter Estimation Problem . . . 52

3.2.2 Bayesian Inference 56

4 Comparison between the eikonal, graph based and monodomain methods for activation time approximations 61
4.1 Introduction 62
4.2 Methods 63
4.2.1 Graph Based Activation Time Approximations 63
4.2.2 Convergence Properties of the Graph Based Method 66
4.2.3 Anatomical Model 67
4.2.4 Stimulation Protocols 69
4.2.5 Error Calculation and Data Fitting 69
4.2.6 Cross Validation Procedure 70
4.3 Results 70
4.3.1 Monodomain Simulations 71
4.3.2 Connectivity vs Convergence in a Cuboid Geometry 72
4.3.3 Comparison with Independently Estimated Cost Functions 76
4.3.4 Mesh Resolution and Quality 77
4.3.5 Comparison with Known Cost Functions 79
4.3.6 Fibre Orientation 85
4.3.7 Computational Performance 86
4.4 Discussion 88
4.5 Conclusions 92

5 Bayesian Inference of Conduction Properties 93
5.1 Introduction 94
5.2 Methods 96
5.2.1 Anatomical Model 97
5.2.2 Bayesian Inference 97
5.2.3 Data Generation 100
5.3 Results 102
5.3.1 Validation of Inference under Different Levels of Measurement and Model Accuracy 102
5.4 Discussion 107

6 Application of Bayesian Inference to Optimal Measurement Placement 112
6.1 Introduction 113
6.2 Methods 113
6.2.1 A Combinatorial View of Measurement Placement 114
6.2.2 An Approximate Greedy Solution for Measurement Placement 116
6.2.3 A Heuristic for Measurement Placement 116
6.2.4 Visualization of the LV Endocardium 118

6.3 Results 119
6.3.1 Effects of Measurement Placement 119
6.4 Discussion 128

7 Conclusions and Future Directions **131**
7.1 Summary 131
7.2 Key Contributions 134
7.3 Future Directions 136
7.3.1 Extensions of the Parameter Estimation Framework 136
7.3.2 Refinement of Models for Anatomy and Errors 139
7.3.3 Applications of the Developed Methods 141
7.4 Concluding Remarks 142

Acknowledgements

First and foremost I would like to thank my supervisors, Dr Blanca Rodriguez and Professor Nic Smith. Blanca, you have been a great support through both good and less good times, in academic as well as real life. You have always spoken your mind about my achievements - occasionally in less than praising terms, but always very honestly. I've learned a lot! Nic, you have given me great advise throughout the DPhil, both in terms of scientific direction and presentation. Through you I have learned a lot about how high level research is done. Further, I would like to thank the examiners for my Transfer, Confirmation and final Viva, namely Professor Kevin Burrage, Dr Jonathan Whiteley, Dr Pier Lambiase and Dr Maxime Sermesant, who have all contributed to focussing and improving the work presented in this dissertation.

Try as he might, no man is an island. Hence, I would also like to thank the members of the computational biology group, who have provided a very stimulating environment. The CHASTE team, in particular Dr Miguel Bernabeu and Dr Joe Mark Pitt-Francis, have given invaluable assistance with turning simulation ideas into synthetic data. Additionally, Dr Pitt-Francis has given great input to the problem of optimal selection of measurements treated in this thesis, for which I am much obliged. I am also especially grateful to Dr Sara Dutta and Dr Alfonso Bueno-Orovio for giving me great assistance in the practical matters involved in getting a dissertation submitted on time.

Further, I would like to acknowledge support from the European Commission as part of the euHeart project, which has provided funding for my studentship. I am also thankful for additional support from the preDiCT grant and the UK Medical Research Council, allowing me to finish my DPhil project.

Finally, none of this would have happened without the unconditional support, enthusiasm and love from my wonderful parents Åsa and Kåre, and from Anna, my far better half. You'll get your terrier, eventually...

Nomenclature

V_m	Transmembrane Potential
J_i	Intracellular Current Density
J_e	Extracellular Current Density
ϕ_i	Intracellular Potential
ϕ_e	Extracellular Potential
σ_i	Intracellular Conductivity Tensor
σ_e	Extracellular Conductivity Tensor
I_m	Total Transmembrane Current
I_{ion}	Total Ionic Current
C_m	Cell Membrane Capacitance
R_m	Cell Membrane Resistance
η	Cellular Electrophysiology State Variables
$\mathbf{f}$	Cellular Electrophysiology Dynamics Function
cv_l	Longitudinal Conduction Velocity
cv_t	Transverse Conduction Velocity
cv_n	Normal Conduction Velocity
n_i	Graph Node
$e_{i,j}$	Graph Edge
$c_{i,j}$	Graph Edge Cost
t_i	Activation Time
$\mathfrak{M}$	Model Space
$\mathcal{M}$	Model
$\mathfrak{Q}$	Parameter Space
$\mathcal{Q}$	Parameter Set
q_i	Parameter
$\mathfrak{D}$	Data Space
d_i	Activation Time Datum
ϵ_e	Activation Time Errors from Modelling
ϵ_m	Activation Time Errors from Measurements
$N(m, \sigma)$	Normal Distribution
σ_0	Standard Deviation Multiplier for Errors in Modelling
σ_m	Standard Deviation of Errors from Measurements

1

Introduction

1.1 Motivation

Cardiovascular diseases (CVDs) are the number one cause of death globally. More people die annually from CVDs than from any other cause. An estimated 17.3 million people died from CVDs in 2008, representing 30% of all global deaths, with extrapolations suggesting a staggering 23.6 million deaths in 2030. In 2003, the cost of CVDs the European Union was 169 billion Euro[1], not including those costs due to the reduction in productivity arising from mortality and morbidity. As a result, the various mechanisms underlying development of CVDs have become the focus of an intensive research effort, where each new level of understanding has the potential to improve preventive and therapeutic measures, ultimately leading to an improved quality of life on a global scale.

Although the CVD category encompasses a heterogeneous range of conditions, many of these diseases are associated with abnormalities in the cardiac electrical activation sequence. These abnormalities can be found in a wide spectrum of CVDs, ranging from heart failure to atrial and ventricular arrhythmias. Despite the fact that the fundamental mechanisms underlying many conduction abnormalities remain unclear, structural abnormalities in the form of scars and fibrotic tissue are known to play an important role.

We thus have a causal chain where structural abnormalities lead to electrophysiological abnormalities, which in turn cause disturbances in the cardiac function. Additionally, feedback often exists between these causal components. A number of clinical interventions have been devised that rely on directly modifying the electrical activation sequence, either by blocking or redirecting the electrical current using ablation lines, or by initiating additional activation waves using electrical pacing from implanted electrodes. However, at present these therapies are expensive and still have a comparatively low success rate. A major reason for the low success rate is an incomplete understanding of the exact disease mechanisms, a difficulty which is further complicated by the high individual variability in disease conditions.

Computational modelling has become an increasingly important tool for understanding how disease related structural changes and clinical interventions affect the cardiac electrophysiology. The power of computational modelling comes from its ability to integrate information from a wide range of sources to make predictions about various aspects of cardiac function, such as the activation sequence, under different conditions, for example in the presence of scars and fibrosis. Once constructed and validated, a predictive cardiac model can be used to optimize clinical therapies by simulating and comparing many different possible approaches to select the most beneficial one. Models also provide a valuable tool to aid the development of new forms of therapy by helping researchers understand the disease mechanisms in more detail.

A biophysically detailed cardiac electrophysiological model is typically formulated mathematically using the bidomain equations (or its simplification, the monodomain equation) [2, 3]. The model consists of two key components: an anatomical mesh defining the spatial domain of the model, and one or several specialized cell models, specifying the properties that describe the cell membrane function of the cells making up the cardiac tissue. At the interface between the anatomical mesh and the cell models are the tissue level conduction properties, representing the electrical coupling between adjacent cells in the cardiac tissue. The conduction properties describe how the cell models interact within the domain defined by the anatomical model, and

as such have a profound impact on the cardiac activation as predicted by simulations. The conduction properties may vary spatially throughout the anatomical model to represent different regions of tissue, and may also vary depending on direction, reflecting the complex spatial orientation of the cells that make up the cardiac tissue.

Thus, a necessary step in the modelling process is to infer the electrophysiological conduction properties of the cardiac tissue from available data. In an *in vitro* laboratory setting, electrical activation can be studied using optical mapping [4, 5], electrode arrays [6] or motion based imaging [7], while detailed spatial data, including cell orientation, can be obtained from histology or MRI after the experiment has been conducted [8, 9]. Contrastingly, in a clinical setting, the available data are generally comparatively sparse. Spatial data can be obtained non-invasively using imaging modalities such as MRI [10], CT [11] or ultrasound [12], but reliable electrophysiological (EP) information still needs to be obtained invasively using catheterization [13]. This limits the amount and quality of data that can be obtained, due to constraints in both time and accessibility.

Because of these inherent limitations, clinical EP contact measurements often consist of point measurements obtained exclusively from the endocardial surface. Since the data sets are typically small, and associated with various degrees of noise, extracted parameter values will be associated with a potentially large degree of uncertainty. This poses two challenges. Firstly, it is important to be able to quantify the uncertainty in extracted parameters, to prevent estimation errors from being translated via associated models to errors in therapy. To this end, the framework of Bayesian inference has gained increasing attention in recent decades [14, 15]. Secondly, particularly in the case of clinical time constraints, it becomes important to optimize the information content of the obtained data by placing the measurements in an optimal way. An additional difficulty comes from the typically high computational demands of electrophysiological simulations. A single heartbeat simulated using the bidomain model can take several hours of computation time. This makes iterative procedures, such as those typically used in

fitting parameters to measured data, impractically slow. Thus, both of the challenges mentioned above must be overcome while keeping CPU time within limits that ensure that the developed methods are practically useful.

1.2 Goal

It is the goal of this dissertation to develop a probabilistic methodology based on Bayesian inference, coupled to a computationally efficient model of the cardiac propagation, to obtain a simultaneous quantification of the cardiac conduction parameters and their associated uncertainty. The developed methodology will subsequently be used to study how the information content of measured electrophysiological data can be optimized. Rather than estimating single parameter values, the Bayesian framework generates a joint probability distribution in the parameter space of the model. This not only provides us with a way to quantify the inherent uncertainty in estimated parameters, but it also gives us a tool to compare two distinct estimates of the same parameter, thus providing us with a measure required for optimizing the information content of the data.

A probabilistic, model driven approach such as the one developed in this dissertation requires a large number of simulations to produce the probabilistic description of parameter values. This makes detailed descriptions of the cardiac conduction such as the monodomain and bidomain models unsuitable, because of their high computational demands. Our first goal is thus to develop a computationally efficient model of the cardiac electrical propagation, optimized for the purpose of estimating conduction parameters from electroanatomical data, and to compare it to the state-of-the-art in detailed and phenomenological models. Secondly, we aim to incorporate the developed model into the framework of Bayesian inference, and demonstrate the ability of the resulting methodology to produce the desired probabilistic descriptions of conduction parameter values. Thirdly, we will use the developed tools to investigate the effect of different strategies for placing electroanatomical measurements on the uncertainty associated with esti-

mated parameter values. The outcome of this will be an *a priori* computable index, predicting the optimal spatial distribution of measurements on the cardiac surface for a defined parameter estimation task.

1.3 Dissertation Outline

We now describe the structure of the remainder of this dissertation, with particular focus on how the methods mentioned above will be developed and validated, and how the aims of the dissertation will be achieved.

We begin in Chapter 2 by introducing the necessary background information on cardiac physiology. The basic structure of the heart and its electrophysiological function are described. Here, particular focus is placed on the properties important for the electrical wave propagation in diseased states, such as myocardial scar formation. Also described is the coronary system, which provides important access routes for catheterization and pacemaker lead placement. Additionally, two main electroanatomical mapping systems are reviewed, with focus on their accuracy. This information will later be used to inform the simulation studies that are used to validate the developed methodology. Also included in this chapter are accounts of computational modelling and parameter estimation, constituting a relevant background to the modelling and inference efforts reported in later chapters.

Chapter 3 describes the key computational methods that have been used in this work. The first part of the chapter focuses on describing cardiac electrophysiological modelling. The state of the art in modelling of cardiac tissue, represented by the bidomain and monodomain models, is presented, along with an overview of the main numerical methods available to solve the associated equations. This is followed by an account of the Eikonal equations, representing the main phenomenological alternative for describing cardiac electrical propagation. Also given is a description of fast marching-like (FMM) algorithms for providing solutions to the eikonal

equations. Finally, the theory and computational methods associated with parameter estimation are outlined. Here, the Bayesian inference framework is also presented and contrasted with traditional approaches for parameter estimation, highlighting the advantages of a probabilistic formulation of parameter values compared to the point estimates obtained from non-probabilistic methods.

Chapter 4 begins with the development of the novel graph based method for activation time estimation that will be used in subsequent chapters to drive the Bayesian inference of cardiac conduction properties. A comprehensive validation study of the graph based conduction model is also presented. In this study, the output from the graph based model is compared to solutions to the eikonal equation obtained with the FMM, using monodomain simulations as the gold standard. The setting for the comparison is an image-based ventricular model, incorporating realistic fibre orientation and a transmural scar. Differences in accuracy and computational performance between the three methods are presented. Parts of this chapter correspond to theory and results published in [16] and [17].

After the methodological developments in the previous chapters, Chapter 5 presents a computational study where we apply Bayesian inference, coupled to the path finding model of wave propagation developed in Chapter 4, to estimate conduction parameters from virtual electroanatomical data. The data are obtained from an image-derived human ventricular model, incorporating realistic fibre orientation and a transmural scar. Using probabilistic formulations for errors in modelling and measurements, the impact of various degrees of inaccuracy on estimated conduction parameters and their associated uncertainties are studied. Parts of this chapter correspond to theory and results presented in [18].

Next, in Chapter 6, we analyze the challenge of finding an optimal strategy for placing electroanatomical measurements in order to maximize their information content. Building on this analysis, we then proceed to define a heuristic for predicting the optimal measurement placements, thereby addressing the challenge. The performance of the heuristic is subsequently

investigated by comparing it to two alternative strategies for measurement placement within the simulated human setting from the previous chapter. Parts of this chapter correspond to theory and results presented in [18].

Finally, in Chapter 7, the work presented in previous chapters is discussed. A summary of the main contributions of this dissertation is given, including validation of the graph based method, development of the probabilistic parameter estimation framework, and the application of the probabilistic parameter representation to the problem of optimal measurement placement. Conclusions are drawn as to how measurements can be collected in order to construct personalized electrophysiology models in an efficient and reliable way. Further, the impact of uncertainty in modelling and measurements is discussed, as well as how the trade-off between accuracy and speed can be optimized. Lastly, future developments of the presented methods are treated.

2

Literature Review

Following the background and formulation of challenges given in Chapter 1, this chapter aims to provide an overview of a number of relevant background topics, laying the foundation for the methods and results presented in the rest of the dissertation. The chapter begins with an overview of the cardiac anatomy and electrophysiology, representing the object of study for this work. This is followed by an in-depth review of the main electroanatomical mapping systems, used clinically to obtain information of the in vivo *cardiac electrical propagation in human hearts. Here, the published data on spatial and temporal accuracy are summarized, which will form a context for the probabilistic methods developed in later chapters. This is followed by a review of computational modelling of the heart, giving a short historical context of the subject, as well as describing the current state-of-the-art. The chapter finishes with an account of the previous and current methods for estimating cardiac conduction parameters, further setting the scene for the parameter estimation efforts reported in this dissertation.*

2.1 Cardiac Anatomy

The heart is a hollow muscle, responsible for pumping blood through the body. An average heart weighs around 300 g and beats at a rate of 75 bpm. The heart is roughly left-right symmetrical, with two types of chambers on each side, the atria and the ventricles.

The heart receives deoxygenated blood from the body through the superior and inferior vena cava, leading it into the right atrium. From there, it passes into the right ventricle through the tricuspid valve. When the right ventricle contracts, the blood is pushed through the pulmonic valve, via the pulmonary artery and into the lungs, where it is reoxygenated and depleted of carbon dioxide. From the lungs, it then flows back into the heart through the pulmonary veins and into the left atrium. From there, it passes through the mitral valve into the left ventricle, which finally discharges it through the aortic valve, and via the aorta out in the rest of the body. For an overview of the cardiac anatomy, see Figure 2.1.

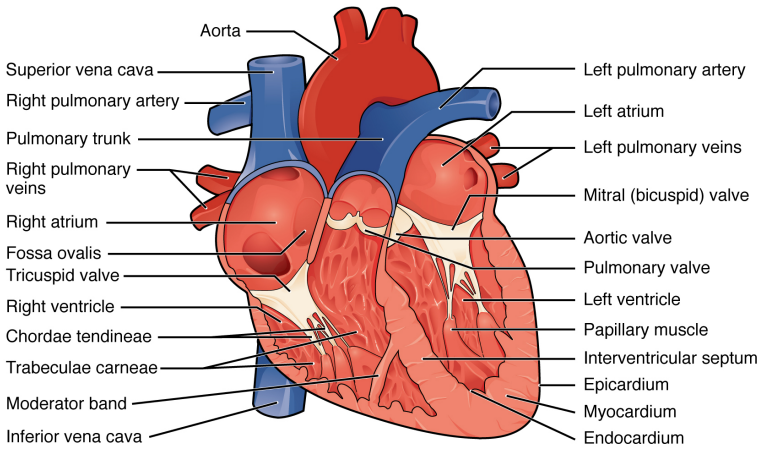


Figure 2.1: An overview of the cardiac anatomy (from www.cnx.org).

The blood supply to the heart itself comes from the right and left coronary arteries, emanating from the base of the aorta, as outlined in Figure 2.2. The right coronary artery supplies blood to the right atrium, the right ventricle, the lower parts of the left ventricle and the back of

the septum. The left coronary artery divides into the circumflex artery and the left anterior descending artery. The circumflex artery supplies blood to the left atrium and the back of the left ventricle, while the left anterior descending artery supplies blood to the front and bottom of the left ventricle, and to the front of the septum.

Deoxygenated blood arrives in the right atrium through the coronary veins. The main coronary vein connected to the right atrium is the coronary sinus (CS). This typically branches off into the small, middle and great cardiac veins. The great cardiac vein branches off into the posterior vein, left marginal (or posterolateral) vein, the marginal (or lateral) vein, and ends in the anterior interventricular (anterolateral) vein. The structure of the coronary veins is highly variable between individuals [19], which has an impact on which access routes for catheterization of the outer surface of the heart are available.

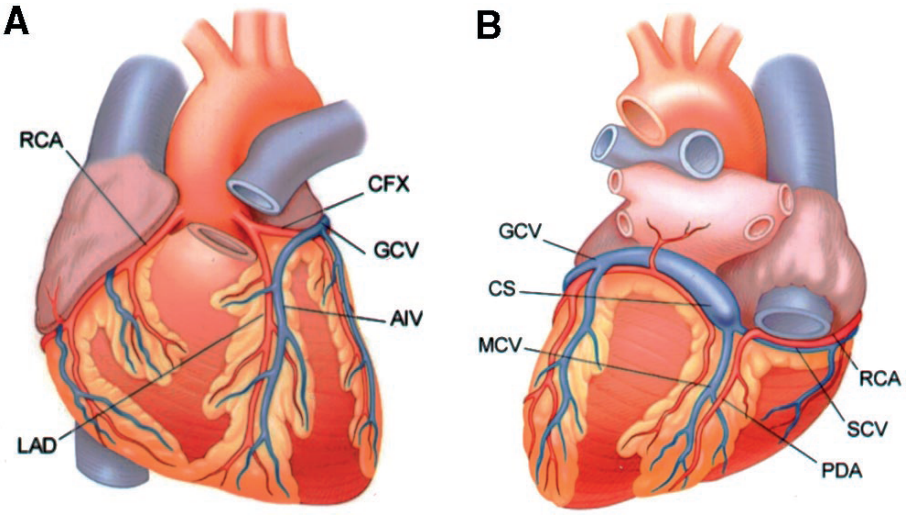


Figure 2.2: An overview of the venous and arterial systems. Anterior (A) and posterior (B) views of epicardial coronary circulation. LAD = left anterior descending coronary artery; AIV = anterior interventricular vein; CFX = circumflex coronary artery; RCA = right coronary artery; GCV = great cardiac vein; PDA = posterior descending artery; CS = coronary sinus; MCV = middle cardiac vein; SCV = small cardiac vein. From [20].

The heart is enclosed in a double-walled sac called the pericardium. The pericardial sac protects the cardiac muscle and anchors it to the structures surrounding it. The tissue making up the

muscular walls of the heart is termed the myocardium. The myocardium, in turn is composed of three layers, termed the endocardium, the mid-myocardium and the epicaridum, all with slightly different physiological properties.

The myocardium consists of a series of inter-connected sheets of tissue separated by cleavage planes. Cardiac muscle cells, called cardiomyocytes (hereafter referred to as myocytes), are locally aligned side by side, forming myocardial sheets. From a functional vantage point, the myocytes are mechanically connected to each other via intercalated disks and electrically coupled via gap junctions to form an electrically continuous syncytium. Structurally, the myocardial sheets have a complex spatial orientation, with the direction of the long axis of the cells rotating approximately 60 degrees from epicardium to endocardium. During a heartbeat, the myocytes in each sheet will contract in a roughly uniform direction and the sheets will slide over each other, maximising the volume output and mechanical efficiency of the heart.

2.2 Cardiac Electrophysiology

Myocytes contract as a result of an external electrical stimulus, usually originating from adjacent cells. When at rest, a myocyte has a negative potential across its membrane. This is known as the resting potential, and is around -85 mV for the human heart. However, myocytes are excitable cells, meaning that they have the ability to receive and propagate an electrical perturbation. The excitable properties of the myocytes stem from the fact that the cells continuously invest energy in maintaining a difference in the concentration of several ionic species between the outside and inside of their membranes. The main ionic species involved in the electrophysiological function of myocytes are calcium (Ca^{2+}), sodium (Na^{+}) and potassium (K^{+}), with chlorine (Cl^{-}) playing a minor role. Ionic pump mechanisms transport Ca^{2+} and Na^{+} out of the cells, while simultaneously transporting K^{+} into the cells. This creates an electrochemical gradient, comprising the combined effects of differences in charge and differences in concen-

tration across the cell membrane, acting as a storage of electro-chemical potential energy. The gradient, which is determined by the combination of the pumping action of the electrogenic pumps and the outside concentration of the different ionic species, allows calculation of the so called equilibrium potential (E) for each ionic species. This signifies the point where the attraction caused by the electrical potential is exactly counterbalanced by the ionic concentration. By definition, a positive equilibrium potential corresponds to a driving force into the cell, while a negative equilibrium potential corresponds to a driving force out of the cell. The equilibrium potentials for Ca^{2+}, Na^{+} and K^{+} are $E_{Ca} = +134$ mV, $E_{Na} = +52$ mV and $E_K = -96$ mV, and the force acting on the respective ionic species is proportional to the difference between these equilibrium potentials and the transmembrane potential.

When an external stimulus causes the transmembrane potential to cross a certain threshold, voltage activated ion channels open, allowing a fast flow of ions down the electrochemical gradient and into the cell. The resulting changes in transmembrane potential subsequently triggers a series of time and voltage dependent changes in the permeability of the membrane to various ionic species, causing further characteristic changes to the transmembrane potential due to ion flux down the respective gradients. The trajectory of the transmembrane potential is known as the action potential (AP), which is schematically illustrated in Fig. 2.3.

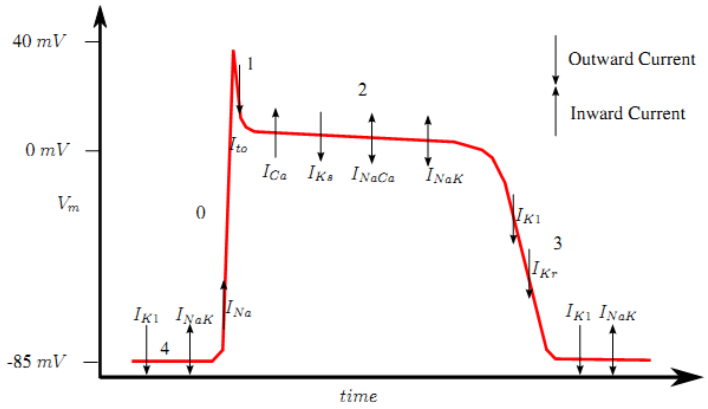


Figure 2.3: Schematic illustration of the cardiac action potential.

As indicated in Fig. 2.3, the AP can be divided into 5 main phases, corresponding to changes in the permeability of specific groups of ion channels. Phase 0 is the rapid depolarization phase, or upstroke phase. The phase is initiated by an external stimulus, usually by ions leaking in from a neighbouring cell through the gap junctions. This leads to opening of the fast Na^+ channels, which in combination with the electrochemical gradient causes a rapid influx of Na^+ into the cell, known as the fast sodium current (I_{Na}). The influx of Na^+ generates a rapid increase in membrane potential, moving from the resting state at around -85 mV to the AP peak at around 50 mV during the span of a few milliseconds. The AP reaching its peak value marks the end of phase 0. Phase 1 occurs with the voltage dependent inactivation of the fast Na^+ channels. The transient net outward current causing the small downward deflection of the action potential is due to the movement of K^+ and Cl^- ions, carried by the I_{to1} and I_{to2} currents, respectively. Particularly the I_{to1} contributes to the ”notch” exhibited by ventricular action potentials in a number of species. Phase 2, commonly referred to as the plateau phase, is sustained by a balance between inward movement of Ca^{2+} (I_{Ca}) through L-type calcium channels and outward movement of K^+ through the slow delayed rectifier potassium channels, $I_{K,s}$. The sodium-calcium exchanger current, $I_{Na,Ca}$ and the sodium/potassium pump current, $I_{Na,K}$ also play minor roles during phase 2. During phase 3, denoted the repolarization phase, the L-type Ca^{2+} channels close, while the slow delayed rectifier $I_{K,s}$ K^+ channels are still open. This ensures a net outward current, corresponding to a negative change in membrane potential, thus allowing additional types of K+ channels to open. These are primarily the rapid delayed rectifier K^+ channels ($I_{K,r}$) and the inwardly rectifying K+ current, (I_{K1}). This net efflux of positive charge causes the cell to repolarize. The delayed rectifier K^+ channels close when the membrane potential is restored to about -80 to -85 mV, while I_{K1} remains conducting throughout phase 4, which signifies the fully repolarized state of the cell.

Electrical conduction is normally initiated at a small region of pacemaker cells called the sinoatrial (SA) node, situated in the right atrium. The cells in the SA node spontaneously depolarize

at regular intervals, with each depolarization initiating a new heart beat. From the SA node, the current propagates across the right atrium along four main conduction tracts, three leading down towards the atrioventricular (AV) node and the fourth (Bachmans bundle) leading to the left atrium. Once the current reaches the AV node, it passes down through the bundle of His, and out into the left and right branch of the purkinje tree, to finally emerge into the left and right ventricle. Except for along the conduction pathways described above, the atria are largely insulated from each other and from the ventricles. For an overview of the cardiac conduction system, see Figure 2.4.

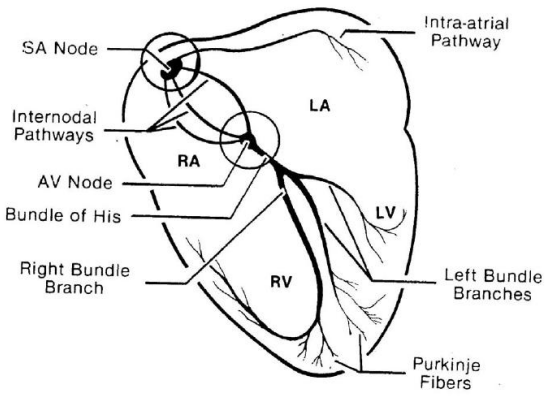


Figure 2.4: Detailed anatomy of the cardiac conduction system. RA: right atrium, LA: left atrium, LV: left ventricle, RV: right ventricle. (from www.tpub.com)

2.3 Normal Myocardial Conduction

Myocardial conduction of the electrical impulse generated from the ion channel dynamics as described in the previous section is a complex process involving the interaction of several variables. These include cell size and shape, cell-to-cell coupling, microscopical tissue structure and expression of ion channels [21]. Ventricular myocytes are roughly cylindrical in shape, and measure around $100\mu m$ in length, with a diameter of around $20\mu m$. As noted above, the my-

ocytes are connected by specialized ion channels known as gap junctions. These are arrays of densely packed protein channels that permit intercellular passage of ions and small molecules. Because current transfer between neighbouring myocytes occurs only at gap junctions, the spatial distribution and biophysical properties of gap junction channels are important determinants of the conduction properties of cardiac muscle [22].

Reconstructions of the distribution of cell–cell junctions from histochemical and electron microscopy studies have revealed tissue-specific patterns of intercellular connections, which appear to confer distinct conduction properties. Studies of intercellular connections in the canine left ventricular have revealed that individual ventricular myocytes are connected to an average of around 11 neighbors, with an approximately equal distribution between side-to-side and end-to-end connections, owing to the higher concentrations of gap junctions at the ends [23, 24].

This pattern of intercellular connections gives rise to anisotropic conduction properties in the ventricular myocardium. A wavefront traveling in the transverse direction must cross more intercellular junctions and thus encounters greater resistance and propagate more slowly than wavefronts traveling an equal distance in the longitudinal direction. Recently, studies in pig have provided added detail to the picture by highlighting the impact on conduction of the myocardial sheet structure [25], giving compelling evidence for axial conduction anisotropy.

Another factor that affects the cardiac conduction is the upstroke velocity of the action potential [26]. As explained in the previous section, when a cell receives an electrical impulse, voltage sensitive ion channels open and the voltage across the cell membrane increases. The rate of this increase in membrane potential defines a period of delay before the activated cell is able to propagate the impulse to its neighbours. On a tissue level, this phenomenon impacts the conduction velocity of the electrical activation wave.

Together, these factors interact to define heterogeneous and anisotropic conduction properties across the myocardium. At a macroscopic level, these properties can be described as a triplet of

perpendicular conduction velocities, defined in each point in the myocardial tissue. These conduction properties are also affected by disease conditions such as cardiac infarctions, providing an additional layer of complexity to the cardiac conduction.

2.4 Myocardial Scars and Fibrosis

The cardiac muscle requires a continuous supply of nutrients and oxygen. As we have seen in Section 2.1, the coronary arteries efficiently perfuse the cardiac muscle in the healthy heart. However, atherosclerosis can result in blockage of these arteries via plaque formation or thromboembolism, depriving the downstream myocardium of oxygen. If this situation is not resolved quickly, the cardiac tissue begins to die, and an infarct zone is created. The infarct zone is commonly divided into a core zone (CZ), where the most severe damage is sustained, and a surrounding border zone (BZ), where residual blood supply still exist, making the damage less severe.

The formation of a myocardial scar is a dynamic process that spans over months. Shortly after the occurrence of the infarct, various inflammatory cells migrate to the damaged region. Removal of dead and dying tissue ensues, and the release of growth factors and cytokines from inflammatory cells such as monocytes and macrophages attracts other cells that, after having reached the site of injury, develop into myofibroblasts. Subsequently, the myofibroblasts begin to synthesize large amounts of the various components that make up the extracellular matrix (ECM), including glycoproteins, proteoglycans, and fibrillar collagens including types I and III, which predominate in the ECM. The ECM synthesis is thus dramatically increased within days of injury [27].

Over several weeks following the initial injury, collagen fibers are constructed throughout the infarct zone; these are initially composed primarily of type III collagen, but are eventually replaced by the mechanically stronger type I collagen as the development of the scar is modulated

by signals from the recruited myofibroblasts. Maturation of collagen fibers, including significant cross-linking, continues to increase the tensile strength of the developing scar. Within 8 weeks of injury, the scar is fully mature [27].

Cardiomyocytes are terminally differentiated and have left the cell cycle. Thus, they fail to repopulate the scar in sufficient numbers to effect a functional repair. This results in a scar that persists for the life of the patient. Studies in human have shown that myofibroblasts in the cardiac infarct scar also can persist for many years [28]. This is likely to affect the cardiac function negatively; besides impairing cardiac contraction and relaxation, both myofibroblasts and the ECM they produce exhibit electrical properties different from those of the surrounding tissue, potentially contributing to arrhythmogenesis. The arrhythmogenic effects of myocardial scars stem from both electrophysiological and structural factors. The main risk factors arise in the BZ, where structural and electric remodeling can give rise to slow conduction, unidirectional block, and rate dependent electric instability after myocardial infarction, all of which contribute to increasing the probability of electrical reentry.

Structurally, strands of surviving myocytes have been found within the BZ of mature scars, forming a network of conduction pathways. Results from computer simulations on these networks suggest that the topology alone is sufficient to initiate reentry via unidirectional conduction block and decreased conduction velocity [29]. Electrotonic interactions between myofibroblasts and cardiomycytes have been shown to give rise to spontaneous activity in monolayer cultures of cells from rat [30], suggesting that myocardial scar formation might lead to the emergence of ectopic foci in the BZ, where fibrotic tissue and viable myocytes are both present. In the BZ of subacute canine infarcts, peak Na^+, Ca^{2+}, and K^+ currents are reduced, giving rise to decreased conduction velocity (CV) and prolongation and the action potential duration (APD) [31, 32, 33]. During the maturation of the scar in canine, reverse remodelling affect the ionic currents and work towards normalizing the APD morphology, but the decrease in CV has been shown to persist [34].

2.5 Electroanatomical Mapping

In order to study the cardiac electrical activation of the human heart in clinic, and to characterize subject specific features such as scars, measurements need to be obtained. The most reliable source of information in this case are measurements obtained from direct contact with the surface of the cardiac tissue. Endocardial contact measurements are conducted using catheters inserted into the chambers of the heart through major blood vessels. Traditionally, contrast fluid and X-ray systems have been used to visually guide the catheters to the desired position, where contact measurements of the endocardial potential can be obtained. However, in recent decades several systems have been developed which deliver simultaneous measurements of voltage and catheter position. These systems permit investigation of endocardial electrodynamics with a significantly reduced need for fluoroscopy, leading to reduced radiation doses for both patients and clinicians [35, 36].

Among the types of functionality provided by endocardial mapping systems are chamber reconstruction, tagging of important anatomic landmarks, interactive display of diagnostic and mapping catheter locations, activation mapping, and scar mapping. Although not yet routinely used, these endocardial mapping systems have the ability to provide clinicians with information that can greatly optimize therapeutic interventions.

The endocardial mapping systems can be further divided into two categories, contact mapping and non-contact mapping. The contact mapping category contains two major systems: CARTO and EnSite NavX, which both utilizes locatable catheters to provide measurements of the endocardial potential by touching the catheter to the cardiac surface. The non-contact mapping category comprises only one system, commonly referred to plainly as EnSite. As with the contact mapping systems, the EnSite system uses a locatable catheter to estimate the chamber geometry, but differs by utilising an inflatable balloon catheter, covered with a large number of electrodes, to estimate the potential at many endocardial points simultaneously.

2.5.1 Non-Contact Mapping with EnSite

The EnSite system (EnSite; St.Jude Medical, St. Paul, Minn) for non-contact mapping consists of a catheter and a processing unit. The catheter, from top to bottom, consists of a loop shaped tip, a ring electrode, a multi-electrode array (MEA) consisting of 64 electrodes distributed across an inflatable balloon, 2 intra-chamber ring electrodes, and a system reference-ring electrode.

The MEA measures 4.6 cm in length and has a diameter of 1.8 cm, giving it a volume of 7.5ml when inflated. For reference, the LV length and diameter for end diastole are about 8 cm and 6 cm respectively [37]. The RV length is approximately the same, whereas the shortest distance between the chamber walls is around 2.5 cm [38].

The initial step when using the EnSite system is to estimate the chamber geometry. This is done by probing the endocardium with a standard electrophysiology (EP) catheter, which is subsequently located using the EnSite system. A low-amplitude alternating current is passed from the catheter being located and alternately to the top and bottom ring electrodes on the non-contact catheter. The signal is sampled by each of the 64 electrodes of MEA. By solving a point source-sink model (accounting for ambient conductivity), using a standard nonlinear least squares algorithm (Levenberg-Marquardt), the location of the EP catheter relative to the MEA can then be determined [39].

Once sufficient data points have been collected, the hull of the resulting point cloud is computed. This rough shape is subsequently smoothed using spline interpolation to produce the final model of the chamber geometry [39].

For calculating the endocardial potential, the electrical potential in the volume enclosed by the MEA and the endocardium is described by the Laplace equation. The system is then discretized using a standard boundary element method. Since this formulation in its basic form often leads to non-physical solutions, Tikhonov regularization is applied in concert with the composite

residual and smoothing operator for selecting the regularization parameter. Using the geometry estimate, along with measurements from the 64 array electrodes, the system is then solved for the endocardial potential. As a final step, spline interpolation is used to produce 3360 potential points on the endocardial surface [39].

Because data from the entire chamber are collected simultaneously with the EnSite 3000/XP system, it can be used to map non-sustained rhythms such as premature atrial complexes, irregular rhythms such as atrial fibrillation or polymorphic VT.

2.5.2 EnSite Accuracy

Several studies have been made to validate the accuracy of the EnSite system against fluoroscopically guided contact electrograms. Gornick *et al.* validated the system in both a saline filled tank and in a canine model [39]. The tank model was used to validate catheter position determination, while the canine model was used to study the accuracy of the reconstructed electrograms. For the tank model, the authors report a mean error in catheter location of about 2.33 ± 0.44 mm for distances less than 50 mm from the MEA and 7.50 ± 1.13 mm for distances greater than 50 mm. In the canine model, both timing and morphology were compared between contact and computed electrograms. Each measured electrogram was compared to an ensemble of time shifted versions of the corresponding computed electrogram. The time shifted version with the highest correlation was chosen, and the correlation together with the optimal time shift was noted. The needed time shift was -0.64 ± 2.48 ms, rendering an average correlation of 0.966.

In a similar study in human, Schilling *et al.* compared contact and computed electrograms [37]. In this study the researchers compared non-contact electrograms with contact electrograms, using correlation waveform analysis (CWA) [40]. The signals were time shifted to obtain the maximal correlation coefficient (taking values between 1 and -1), and the time shifts (correlation

timing) and correlation coefficients were recorded. They found a correlation timing of -1.30 ± 6.64 ms for distances less than 34 mm and -3.02 ± 21.03 ms for distances greater than 34 mm. The average correlations were 0.87 ± 0.12 and 0.76 ± 0.18 for < 34 mm and > 34 mm respectively. Additionally, the error in timing of the point of maximal upstroke velocity was measured, rendering a value of -1.94 ± 7.12 ms < 34 mm and -14.16 ± 19.29 ms > 34 mm from the MEA.

Abrams and co-workers performed non-contact mapping of the human right atrium (RA) [13]. Using the same approach as Schilling and co-workers, they report a maximal correlation of 0.85 ± 0.13 for distances less than 40 mm and a correlation of 0.44 ± 0.23 for distance greater than 40 mm. The corresponding average correlation timings were 5.97 and 17.54 above and below a distance of 40 mm from the MEA. Interestingly, they also estimated that 51.1% of the RA endocardium was at a distance >40 mm from the MEA.

Shilling and co-workers report an average LV length of 8.33 ± 0.92 cm (end diastolic, from mitral valve to apex) and an average LV (maximal) diameter of 6.17 ± 0.84 cm (end diastolic) and 5.22 ± 1.06 cm (end systolic) for 13 patients [37]. Judging from this, the MEA should be able to keep a distance <34 mm to most of the wall if centrally positioned. However, if offset by a little over a centimetre, parts of the LV endocardium may end up at a distance >34 mm to the MEA, leading to a possible decrease in accuracy. Additionally, patients affected by dilated cardiomyopathy might have severely enlarged ventricles, making EnSite studies more difficult.

In conclusion, the EnSite system for non-contact mapping provides simultaneous virtual electrograms for the entire chamber investigated. When considering activation times, for distances closer than 34 mm, it appears to have a slightly negative bias with a standard deviation in the order of 5 ms. Then, somewhere between 34 and 50 mm from the MEA, both bias and standard deviation seem to grow in a non-linear way, rendering the measurement accuracy far lower, with a bias exceeding 10 ms and a standard deviation of the same order of magnitude.

2.5.3 Contact Mapping with CARTO

The CARTO system (Biosense Webster Inc., Diamond Bar, California) consists of four parts: a catheter, a reference patch, a location pad and a processing unit. The location pad, usually located under the operating table, contains three coils. Each coil emits a spherical magnetic field of low intensity (between $5*10^{-6}$ and $5*10^{-5}$ Tesla), as illustrated in Figure 2.5. The amplitude, phase and frequency of these fields are detected by sensors in the catheter and reference patch, allowing the catheter position to be determined by solving a system of algebraic equations. This determines the catheter position in six degrees of freedom, while correcting for body movements [41].

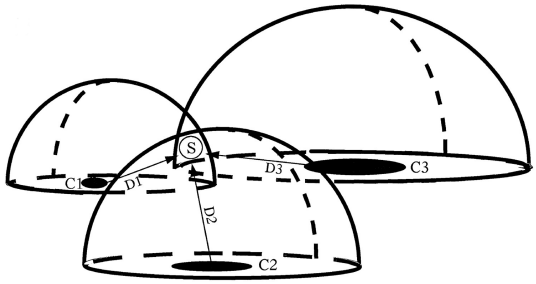


Figure 2.5: Schematic view of the magnetic fields from the location pad used by the CARTO system to locate the catheter [42].

By dragging the catheter along the endocardial surface, a large number of position-electrogram pairs can be sampled. Using a timing reference, usually an intracardiac electrogram from a fixed point, samples from the same stage of the cardiac cycle can be selected for further processing. These point samples are subsequently used by the system to estimate the chamber geometry and activation sequence.

2.5.4 CARTO Accuracy

The CARTO system has been validated in a number of studies, with respect to location accuracy and signal correlation. Gepstein *et al.* validated the system both *in vitro* and in a pig model [42]. For the *in vitro* study, a plexiglas cube, measuring $20\times20\times20$ cm was used. The cube contained 7 holes with known positions. Two protocols were used. In the first protocol, the locatable catheter was placed in each of the seven holes, and its position was determined 25 times by the CARTO system. This yielded an average standard deviation of 0.16 mm, with measurements spanning an average range of 0.55 mm. In the second protocol, the system was used to compute the distance between all pairs of holes. The results were compared to the known distances, rendering an average error of 0.42 ± 0.27 mm. For the *in vivo* study, again two protocols were used, corresponding to the ones used *in vitro*. In the first protocol, the catheter was placed in stable contact with the left ventricular endocardium at 15 different points. For each of these points, the location was determined 20 times. This yielded an average positional error of 0.54 ± 0.74 mm. In the second protocol, the catheter was marked at 10 mm intervals, and gradually introduced into the right atrium, using a guiding sheath to restrict movement to the desired direction. Location was estimated in 10mm intervals corresponding to the markings on the catheter, and distances between the measured locations were compared to the distances between the markings, giving an error of 0.73 ± 0.55 mm between known and measured distances. Finally, the correlation between the signal from the locatable catheter and that from a standard EP catheter was measured, rendering a correlation coefficient of 0.96.

Shpun *et al.* report a study of CARTO spatial accuracy in pig, using radio frequency ablation of the right atrium [43]. A point on the atrium was selected on the visual map of the CARTO system. The point was ablated, leaving a visual mark. This was repeated three times for the same point as visualised by the system, with the catheter completely withdrawn between subsequent ablations. The procedure was repeated for a total of 27 different points. The animals were subsequently killed and their hearts were removed. The distance between the centres of ablation

lesions, corresponding to the same point as visualized by the CARTO system, was found to be 2.3±2.60 mm.

In conclusion, CARTO shows good accuracy when it comes to signal correlation between the CARTO catheter and standard EP catheters. This is to be expected, since, unlike the EnSite system, it actually relies on contact measurements. For location accuracy the error is in the order of 1 mm *in vitro*, while reports on *in vivo* errors differ somewhat, with the higher estimates being around 2 mm with a standard deviation of the same magnitude.

2.5.5 Contact Mapping with EnSite NavX

The NavX system (St. Jude Medical, St Paul, Minnesota) is another system for spatial location of EP catheters. In contrast to the magnetic fields used by the CARTO system, NavX employs externally applied currents to locate a standard catheter electrode. When an electrical current is externally applied through the thorax, a voltage drop occurs across internal organs such as the heart. The resulting voltage can be recorded via standard catheter electrodes and can be used to determine electrode position in relation to the source of the current. 3 skin-electrode pairs are used to send 3 small, 1 mA currents through the thorax in 3 orthogonal directions (X,Y and Z) with slightly different frequencies of around 30 kHz used for each direction.

The X and Y fields are created using standard surface ECG electrodes. The electrodes are placed at the right and left mid-axillary lines at the fourth intercostal space (V2) level, creating the X field, and at the left shoulder and left leg, creating the Y field. For the Z field, two 10 $\times$ 15 cm skin patches are used, 1 anterior above the heart at the V2 position and the other posterior under the heart on the back. The posterior skin patch can simultaneously serve as the return electrode for RF ablation [44].

The mixture of 30 kHz signals, recorded from each catheter electrode, are then digitally separated to measure the amplitude of each of the 3 frequency components. The 3 electrical field

strengths are calculated from the difference in amplitudes. From this data, the 3D position of each electrode is then calculated by dividing each of the 3 amplitudes (V) by the corresponding electrical field strength (V/cm) [44].

The electrode positions are averaged over 1 or 2 seconds to reduce cyclic cardiac variations. Respiratory variations are too slow to be eliminated by averaging without compromising the real-time nature of the localization method. For a schematic view of the setup, see Figure 2.6

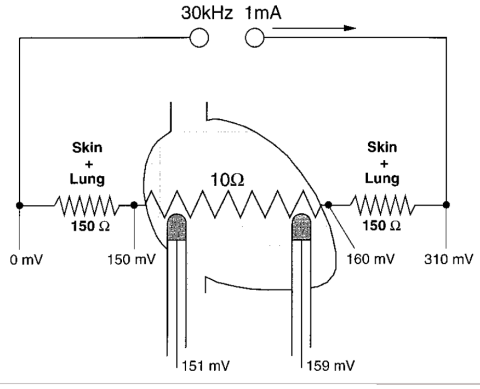


Figure 2.6: Schematic view of the NavX system [44]. An externally applied transthoracic current of 1 mA at 30 kHz results in a voltage gradient across each component of the circuit. Moving the catheter from left to right will result in a change of recorded voltage from 150 to 160 mV. Values are intended as examples and do not correspond to real measurements.

2.5.6 NavX Accuracy

Wittkampf and co-workers performed a validation study of the NavX system. In 30 patients, a deflectable decapolar electrode catheter was placed in a stable position within either the RA, right ventricle (RV), or left ventricle (LV) in 10 different patients for each group. Sequentially, each of the 10 electrodes was connected to the NavX system. Measurements were repeated at 2 other stable catheter positions within the same chamber. The 3 catheter positions were chosen such that the electrode positions covered a major portion of the cardiac chamber. The data were analysed by comparison of measured and true positions of the 10 electrodes for each catheter

position using linear regression analysis. The correlation coefficient and the residual SD around the regression line were calculated [44]. Results are summarized in Table 2.1.

Measurement	mean	SD	test cases	model
Correlation (RA)	0.996	0.004	30	human
Correlation (RV)	0.997	0.003	28	human
Correlation (LV)	0.999	0.001	28	human
Residual SD (RA)	1.7	0.8	30	human
Residual SD (RV)	1.3	0.7	28	human
Residual SD (LV)	1.0	0.4	28	human
Position SD (RA)	0.8		10	human
Position SD (RV)	1.1		10	human
Position SD (LV)	0.8		10	human

Table 2.1: Summary of validation data for the NavX system.

Earley and co-workers compared radiation exposure resulting from CARTO, NavX and conventional EP mapping [36]. They found that EnSite NavX exposes patients to the least amount of radiation, amounting to around 20% of the amount a patient is exposed to during conventional, fluoroscopically guided EP mapping. CARTO also reduces fluoroscopy time when compared with conventional mapping, but to a lesser degree than EnSite NavX. The exposure to radiation during CARTO is about 40% of that from a conventional mapping procedure. These reductions are achieved without compromising the duration, effectiveness, or safety of the procedure. This benefit however is offset by the greater expense of using these technologies with catheter and disposable costs increased by 50% over fluoroscopically guided standard catheter mapping.

2.5.7 Integration of electroanatomical data with CT/MRI using NavX Fusion and Cartomerge

Recently, efforts have been made to integrate electroanatomical mapping with geometries derived from computed tomography (CT) or magnetic resonance imaging (MRI), resulting in two new systems, NavX Fusion and Cartomerge. Early experience using these systems has demon-

strated a degree of accuracy within the range required for clinical use and with potentially important outcomes [45]. The systems work in a similar way. First, a small number of landmark points are used for a rough alignment of the two geometries. Then surface registration is performed for a large number of surface points, guided by fluoroscopic visualization [45, 46].

The most commonly reported measure of the accuracy of NavX Fusion and Cartomerge is the so-called registration error, representing an average measure of the distance between the point cloud sampled by the electroanatomical mapping system and the surface obtained from CT or MR images. Richmond and co-workers [47] report values for the registration error for Cartomerge (2.3 ± 1.8 mm) and NavX Fusion (3.2 ± 0.9 mm, $P<0.001$). Tops [48] investigated the accuracy of the Cartomerge system in the left atrium. Mean distance between the mapping points and the CT surface ranged from 1.7 ± 1.2 mm to 2.8 ± 1.8 mm. This resulted in an average of 2.1 ± 0.2 mm for the patient group as a whole. Brooks reported average NavX Fusion registration errors of 1.9 ± 0.4 mm for 55 patients. Dong *et al.* [49] studied the impact on Cartomerge of rhythm change from intrinsic activation to AF during the procedure, and found no significant change in accuracy.

Finlay and co-workers [50] report results for accuracy of the registration procedure form the CAVERN clinical trial. During AF ablation, the distance from each lesion to the nearest point on the registered image shell was calculated for around 3000 lesions for both systems, distributed across 21 patients for each system. In the study, Cartomerge lesions were found to be closer to the 3D shell than those of NavX fusion (2.0 ± 0.5 mm vs. 3.4 ± 0.9 mm, $p<0.0001$). The maximum distance from lesion to shell per patient was also lower in Cartomerge than with NavX fusion (10.2 ± 0.7 mm vs. 15.9 ± 1.1 mm, $p<0.0002$). Brooks and co-workers [45] assessed navigational accuracy for NavX Fusion by moving the catheter to 10 discrete anatomic sites and determining the distance between the registered surface, and report similar values to those of Finlay (3.3 ± 1.4 mm).

Zhong and co-workers [51] investigated the accuracy of CartoMerge on 5 patients. Here, the

electrode was guided using echocardiography into contact with an average of 16 sites along paths in the left and right pulmonary venous vestibules. At each of these sites, the location of the point was saved using the CartoMerge system. After completion of the procedure, each point was assessed on the CartoMerge image for accuracy, which was defined as the shortest Euklidean distance between the point and the lesion path in the quadrant where the point was truly located, giving an average value of 18 ± 8 mm (range 4-30 mm). In addition, the shortest linear distance from the point to any atrial surface was measured, giving a value of 3 ± 2 mm (range 0–9 mm).

It should be noted that almost all available studies of the accuracy of CartoMerge and NavX Fusion have been performed in the atria. Since the typical distance between a given point within the atria and the image-derived surface is likely to be smaller in the atria than in the ventricles, one might expect the registration errors to increase when applying the systems in the ventricles. Thus, the extrapolation of conclusions the accuracy associated with ventricular mapping and registration should be done with caution. Additionally, as suggested by Zhong *et al.* [51], the registration error might give a poor indication of the accuracy of the registered activation maps.

Finally, it should also be noted that cardiac image registration is itself a very active area of research, demonstrating the challenging nature of the problem. The heart is a non-rigid, moving object with relatively few landmarks to use for geometrical alignment, potentially making even a rough alignment non-trivial [52]. Additionally, both Carto and NavX have a limited spatial resolution, and the catheter position is sensitive to disturbances from e.g. wall motion and blood flow. Also, the optimization problem that needs to be solved to obtain the final registration between surfaces is typically high dimensional, preventing global optimization and increasing the risk of getting stuck in local minima. In summary, there are a number of important sources of error in the image registration involved in mapping electroanatomical data onto geometries derived from CT or MRI, highlighting the need to take the resulting uncertainty into account when using such data for model personalization.

2.6 Computational Models of Cardiac Electrophysiology

After having reviewed the electroanatomical mapping systems, which constitute the principal means of data acquisition for the parameter inference efforts presented in later chapters of this document, we now turn our attention to another central component of the research presented in this dissertation, namely that of computationally modelling the cardiac electrophysiology. This section will review major efforts over the last 50 years, whereas a more detailed account of the current state-of-the-art will be given in the next chapter.

Early computational models of cardiac electrophysiology were based on the mathematical model developed by Hodgkin and Huxley in 1952 to model the ionic mechanisms of the squid giant axon [53]. This body of research took place well before the introduction of the microelectrode, making the squid axon one of the few suitable targets due to its unusual size. The work on the squid axon subsequently earned Hodgkin and Huxley a Nobel prize for Physiology or Medicine in 1963. In 1960 Denis Noble initiated the field of cardiac modelling by publishing work building on the Hodgkin Huxley model to model Purkinje fibre action potentials [54]. Over the decades since the publication of Noble's initial paper, the field of mathematical modelling of cardiac myocytes has seen a rapid development, beginning with a refined model published by the same author [55]. For the next decades, the rate of model development was fairly slow due to the lack of viable measuring techniques for the wide variety of ionic currents present in cardiac myocytes. However, in the mid 1970s, the introduction of the patch clamp technique gave new fuel to the budding field. Among the early adopters of the new technique were McAllister and coworkers, whose 1975 model [56] made several improvements to Noble's Purkinje fibre model, most notably the introduction of calcium currents. Data from the new patch clamp technique also gave rise to the first ventricular myocyte action potential model [57]. The final iteration of Noble's model was described in [58], which was notable as the first model to incorporate the sodium-potassium exchanger. From this point onwards a number of groups began to create mathematical models of ventricular myocytes. Here, a notable group of related models

are the ventricular models developed by Rudy and co-workers, originally designed for use in studies of arrhythmia [59, 60, 61, 62, 63, 64, 65, 66]. The development of human ventricular myocyte models began with [67], though this model was largely based on guinea pig data. The work presented in [68, 69] provided a more comprehensive human ventricular model. More recent and highly detailed models of ventricular myocytes have been developed, both for human [65, 70] and rabbit [71]. For an excellent review of recent developments in cardiac cell modelling, see [72].

It should also be mentioned that although the general trend in ionic cell models has been the incorporation of more and more biophysical detail, there also exists a class of models that take a phenomenological approach to describing the cellular electrophysiology [73, 74, 75, 76, 77]. These models aim to reduce the number of parameters in order to increase the computational efficiency, and to reduce uncertainty in parameter estimation, two challenges that have become increasingly relevant with the advent of subject-specific whole ventricular simulations. The more recent of these reduced cell models are capable of replicating the human tissue level dynamics of more detailed models [78, 79].

In the late 1980s, the commonly available computational power had increased to the point where tissue level simulations became viable. Early tissue level models were mainly concerned with simple domains and mathematical understanding of the bidomain equations [80, 81, 82, 83, 84]. This work moved to the study of the mechanisms underlying re-entry and the effects of defibrillation, but still mainly using regular geometries [85, 86]. This fundamental work on regular geometries eventually paved the way for application of bidomain simulations to more realistic geometries, this time relying on anatomically detailed ventricular meshes built from imaging data. The first such mesh was created for a dog heart [87], though this mesh was not initially designed for electrophysiology simulation. The rabbit ventricular geometric model introduced by Vetter and McCulloch has been perhaps the most successful ventricular model [88], extensively used for bidomain simulations of reentry [89, 90, 91, 92]. A human model developed

from computed tomography data was used to assess the difference between monodomain and bidomain models [93]. Another human model was developed by digitizing sectioned slices of a human heart and used to investigate the organization of ventricular fibrillation in human [94, 95].

More recently, the advent of high-resolution magnetic resonance imaging has allowed personalized models to be constructed [96, 97]. In combination with suitable contrast agents, MRI has enabled the imaging of cardiac scars. Models of cardiac scars have been incorporated into atrial and ventricular image-based models in a number of studies [98, 99, 100, 101]. Recently, highly detailed scar models have been derived from confocal microscopy imaging [29], but simulations on this level of detail have so far been confined to blocks of isolated tissue.

Just as there exist phenomenological descriptions of the cellular electrophysiology, phenomenological descriptions of the electrophysiology at the tissue level have also been developed in order to reduce complexity and computational costs. These are typically models of electrical activation, incorporating little or no information about the cellular mechanisms. In fact, phenomenological tissue models have existed far longer than their more detailed cousins. However, their deeper theoretical anchoring has required comparison to biophysically detailed models of cardiac tissue, such as the bidomain model mentioned above. A very early example of a phenomenological cardiac tissue model was put forth by Wiener and Rosenblueth in 1947 [102]. The model, which was far ahead of its time, was essentially a collection of cellular automata (CA), and was employed to predict both vulnerability to reentry and spiral waves. Later, in 1964, Moe et al. [103] put forth a 2D CA model. The goal of their study was to elucidate the arrhythmogenic nature of cardiac tissue, and consequently, rules for determining refractoriness were also included. The model permitted investigation of the effects of inhomogeneities leading to re-entry and fibrillation. A subsequent three-dimensional model by Okajima et al. [104] was concerned with the simulation of activation of normal tissue. The effect of varying the initiation of excitation was examined as an approach to the study of the behavior of the con-

duction system, and the authors also explored the effect of conduction abnormalities with the model. In the mid 1970s, Eifler and Plonsey [105] described a model based on an anisotropic structure to investigate anisotropic velocities and isochronal shapes in the presence of ischemic and infarcted tissue. Building on this approach, Eifler later used a similar phenomenological model to compute body surface electrograms [106]. In the early 1990s, work by Colli-Franzone et al. [107] and Keener [26] elucidated the relationship between the eikonal equation and the bidomain model, which will be described in the next chapter. This provided a stable theoretical foundation for many phenomenological wave propagation-based descriptions of the cardiac activation.

2.7 Estimation of Cardiac Conduction Parameters

Personalization of cardiac model conduction properties from clinical data is a fairly new field, largely owing to the relatively recent advent of *in vivo* measurement techniques such as electroanatomical mapping systems, reviewed in Section 2.5. However, more generally, attempts to characterize the conduction properties in cardiac tissue *in vitro* have a long history. The earliest efforts to study cardiac conduction anisotropy, as reviewed in Section 2.3, include those of Sano and coworkers [108], who directionally paced strips of endocardial and epicardial ventricular tissue from dog, and those of Draper and Mya-Tu [109], who extended the scope to include both atrial and ventricular tissue from dog, as well as cat, goat and rabbit. Clerc later performed a similar experiment on strips of calf tribecular muscle [110], but now relating the measured values to the parameters of a cable model of the cardiac tissue under study. Later, these essentially one dimensional measurements were contrasted with measurements performed by Roberts and co-workers [111, 112], who used a 7x12 matrix of electrodes to measure epicardial wave front shapes *in situ* in dog, using an elliptical approximation to the wavefront shape to estimate directional conduction velocity and resistivity. The family of electrode matrix techniques has

recently been extended to 3D by Hooks and Trew *(*et al.) [25, 113], who constructed and used a 3D plunge electrode matrix for measuring the shape of the activation wave within the porcine myocardial wall *in vitro*. Directional conduction parameters were then estimated directly using isochrone distances. A recent related approach presented in [114] relies on a ring shaped multi-electrode to measure the conduction velocity of a passing wavefront by fitting the electrode activation times to a cosine function. In [115], Shilling and coworkers measured isochrone distances perpendicular to the wavefronts to calculate directional conduction velocities.

Another line of investigation aiming to estimate cardiac conduction properties relies on the so-called four electrode technique [116]. The method utilizes a linear array of four equally spaced electrodes, the outer two of which carry an applied current, while the inner two measure the resulting voltage. Subsequently, a suitable linear transformation is applied to the resulting data, under which the transformed medium becomes isotropic while the electrodes retain their spatial configuration. One can then obtain resistivity values for the isotropic case, from which an inverse transformation yields the desired resistivity values in the original, anisotropic medium. Plonsey and Barr provided a comprehensive mathematical analysis of the technique as applied to estimating resistivity parameters in cardiac tissue under assumptions of equal anisotropy rations between the intracellular and extracellular domains [117]. The basic method has later seen several adaptations, notably to avoid the assumption about equal anisotropy ratios [118, 119, 120].

Due to the increasing computational power routinely available to researchers, the estimation of cardiac conduction parameters has gradually shifted towards model based methods, relying on raw computational power to minimize the difference between model output and simulated values using optimization techniques. In this research, both surface and volumetric models have been used together with fast phenomenological descriptions of current propagation, such as the eikonal-diffusion (ED) equation [107], or simple dynamical models such as the Mitchell-Schaeffer (MS) [77] and Aliev-Panfilov (AP) [75] models to estimate anisotropic conduction

parameters through repeated simulations [121].

Recently, Sermesant and collegues [122] estimated local conduction velocities in an MRI derived model of the human ventricles. The estimation was done by comparing simulated activation times from the ED equation to endocardial non-contact mapping data. The estimation was done sequentially: firstly, the endocardial onset of activation was determined from the activation data. Secondly, the endocardial conduction velocity was determined by iterative minimization of the difference between measured and simulated activation times, using an adaptive domain decomposition. Finally, the remaining myocardium was divided into a small number of regions, the conduction properties of which were estimated using iterative minimization of a new error function, now including the difference in total activation time as measured using the QRS duration of the ECG.

A similar approach was taken by Relan and co-workers [123], but now combining eikonal and dynamical modelling to estimate parameters based on a combination of CT/MRI and non-contact mapping data from human. Again, the ED equation was used to estimate regional conduction velocities, which were then used to set the conductivity tensors for a monodomain formulation incorporating the MS model. Subsequently, remaining parameters of the MS model were estimated from action potential duration restitution curves using APD data from different pacing protocols.

In [124], Camara and co-workers used a genetic algorithm to estimate the distribution of purkinje terminals in a porcine model, constructed from MRI data. Activation times resulting from each generation of terminal distributions were computed using an anisotropic eikonal formulation and compared to corresponding activation times measured using optical mapping. Subsequently the average absolute error between simulated and measured values was used to calculate the individual fitness functions for selection on the next generation in the genetic algorithm.

A full Bayesian treatment of the ventricular personalization problem was recently given by

Konukoglu and co-workers in [125]. Here, a novel inference approach based on a polynomial chaos formulation ED equation was used to estimate a conductivity function for the LV endocardium from electroanatomical data. Since this work is especially similar to some of the results presented in this dissertation, it will be given further attention in Chapter 5 below.

Initial efforts have also been made to directly adapt monodomain conductances (often together with other parameters) using various error minimization methods, ranging from manual tuning to genetic algorithms. In [126], Aguado-Sierra and co-workers iteratively adjusted the conductivities the monodomain model to minimize an error measure based on a combination of the difference between electroanatomical and simulated data, and on the difference in total activation time as obtained from the QRS duration.

In [127], Pop and *et al.* used combined data from MRI and optical mapping to estimate electrophysiological parameters in a porcine monodomain model, with cellular dynamics described by the AP model. Initially, the average tissue conductivity ratio the a parameter of the AP model were estimated from isochrones around a stimulating electrode. Subsequently, the ventricular volume was divided into 26 disjoint regions, and the conductivity of each regions was iteratively modified to improve the model fitting by employing a trust region method, until a pre-specified threshold on the absolute difference between simulated and measured depolarization times was reached.

Szilagyi *et al.* [128] used a genetic algorithm approach to simultaneously adapt parameters for a coupled torso and electromechanical cardiac model. The method relied on a human MRI and CT derived geometry, ECG and ultrasound data, combined with a parameter reduction scheme to achieve the computational efficiency required to treat the optimization problem.

Wang and co-workers [129] used a probabilistic approach based on the unscented Kalman filter to estimate personalized cardiac electrophysiological properties based on optical mapping and MRI data. A low dimensional mesh free method boundary element method was used to derive a

state space formulation of the monodomain equation using the AP model for cellular dynamics. Subsequently, probabilistic simulations were performed while accounting for observed data to constrain the solution space. By iterating these simulations over a range of model parameter values, probabilistic estimates of the parameter values were obtained.

Finally, it should be mentioned that the imaging of cardiac activation directly from high resolution body surface potential maps [130, 131] is currently reaching a mature state and is beginning to see clinical use [132]. Wang and co-workers have recently presented a statistical framework to reconstruct volumetric activation inside the 3-D myocardium from noninvasive BSP recordings [133], relying on a stochastic state-space representation to take into account model and data uncertainties. This has been used for e.g. non-invasive scar imaging [134]. Although strictly more concerned with the estimation of dynamical states rather than parameters, these techniques have the potential to provide valuable data that can be used for model personalization.

2.8 Conclusions

In this chapter, we have described how the electrical activation wave propagates throughout the cardiac tissue, ultimately leading to its coordinated contraction. Further, we reviewed normal and diseased conduction properties of the myocardial tissue, which will constitute the object of study for the parameter estimation efforts in Chapters 5 and 6. Additionally, we reviewed the major electroanatomical mapping systems, representing the principal means of data collection for parameter estimation of conduction properties in the aforementioned chapters. In the review of the electroanatomical mapping systems we also saw that the collected data are often associated with a significant amount of uncertainty, motivating the use of probabilistic methods in Chapters 5 and 6. We also gave an overview of previous efforts in cellular and tissue modelling, providing a background to the graph based conduction model presented and validated in Chapter 4. Finally, we gave an overview of previous work relating to the estimation of conduction

parameters, giving further context for the work presented in Chapters 5 and 6.

3

Computational Methods

In the previous chapter, we gave the general context and background for the work presented in this dissertation. In this chapter we expand on the essential details of the computational and mathematical techniques that underpin the development of ventricular models of cardiac electrophysiology, as well as faster phenomenological alternatives, leading up to the graph based conduction model presented in the following chapter. We begin this chapter by introducing the bidomain model in Section 3.1.1, which represents the current state of the art in describing the electrical propagation in syncytial tissue. In Section 3.1.2 we account for the reduction of the bidomain model to the monodomain model. In Section 3.1.4 we present the eikonal equations, giving the context of the work presented in Chapter 4. Finally, in Section 3.2 we treat the problem of parameter estimation, giving relevant background for the results presented in Chapters 5 and 6.

3.1 Models of Cardiac Electrical Conduction

3.1.1 The Bidomain Model

For simulations of the cardiac electrophysiology, the major part of the work done during the last decade is based on the bidomain equations [2, 3, 135]. These equations are a result of a homogenisation approach, describing the cardiac tissue as two continuous and interpenetrating domains. The domains represent the volume-averaged properties of the intracellular and extracellular spaces respectively, with the behaviour of current flow between them governed by a model of the average electrophysiological properties of the cell. Within this formulation, two electrical potentials exist at each point in space, one intra-cellular (ϕ_i), and one extra-cellular (ϕ_e). This also implies a transmembrane potential, V_m, at each point in space, defined as

$$V_m = \phi_i - \phi_e. \tag{3.1}$$

The movement of charged particles within the two domains gives rise to currents. To describe this, we define two current densities, one intra-cellular (J_i) and one extra-cellular (J_e). By application of Ohm's law, these quantities can be related to the two potentials as

$$\begin{aligned} J_i &= -\sigma_i \nabla \phi_i \\ J_e &= -\sigma_e \nabla \phi_e, \end{aligned} \tag{3.2}$$

where σ_i and σ_e are the intracellular and extracellular conductivity tensors. Requiring conservation of current and considering a small patch of membrane, the current densities can be directly related to the transmembrane current (I_m), according to

$$\begin{aligned} \beta I_m &= \nabla \cdot (\sigma_i \nabla \phi_i) \\ \beta I_m &= -\nabla \cdot (\sigma_e \nabla \phi_e), \end{aligned} \tag{3.3}$$

where β is the cellular membrane surface to volume ratio. Finally, the total transmembrane current is defined according to

$$I_m = C_m \frac{\partial V_m}{\partial t} + I_{ion}(\eta, V_m). \tag{3.4}$$

Here, C_m is the membrane capacitance and I_{ion} is the ionic current, a function dependent on the model of the ion channel dynamics. η is a vector containing the state variables for the cell model. This leads us to the bidomain equations:

$$\begin{aligned} \nabla \cdot (\sigma_i \nabla \phi_i) &= \beta(C_m \frac{\partial V_m}{\partial t} + I_{ion}(\eta, V_m)) \\ \nabla \cdot (\sigma_e \nabla \phi_e) &= -\beta(C_m \frac{\partial V_m}{\partial t} + I_{ion}(\eta, V_m)). \end{aligned} \tag{3.5}$$

The ion channel dynamics are governed by a set of non-linear differential equations

$$\frac{\partial \eta}{\partial t} = \mathbf{f}(\eta, V_m) \tag{3.6}$$

where f is a vector valued function. The function $\mathbf{f}$ encapsulates the cellular level electrophysiology as reviewed in Section 2.2, and can be represented by a wide variety of models [136].

3.1.2 The Monodomain Model

The bidomain equations can be very computationally expensive to solve, which limits their applicability in whole organ simulation scenarios. This is especially true in scenarios where multiple simulations need to be performed, such as parameter estimation. However, under certain conditions, the bidomain model can be reduced to a simplified form, known as the monodomain model, without significant loss of accuracy [93]. The computational cost of simulating the monodomain model is significantly lower than corresponding simulations using the bidomain model.

As the name suggests, while the bidomain model consists of two spatially interpenetrating domains, representing the two sides of the cell membrane, the monodomain model combines these domains into a single one. This prevents the model from being used in studies where the special properties of the extracellular domain plays an important role, such as in simulations of defibrillation [93]. Additionally, it has been shown that conduction near the surface of the tissue behaves differently between the monodomain and bidomain formulations [137]. Thus, the monodomain approximation is better suited for situations where the myocardium is relatively thick, since a smaller fraction of the tissue is close to the boundary in this situation. This implies that the monodomain model can be expected to give more accurate results in the case of a human geometry than in the case of smaller hearts, such as a rabbit geometry.

The basic assumption behind the monodomain model is that the conductivity tensors in the extracellular and intracellular domains are related by a constant factor. We start the derivation of the monodomain equation by defining the total current density

$$J_t = J_i + J_e. \tag{3.7}$$

By substituting this relation into Eq. (3.2) and rearranging, we can express the intracellular current density as

$$\nabla\phi_i = (\sigma_i + \sigma_e)^{-1}(\sigma_e \nabla V_m - J_t). \tag{3.8}$$

Together with the relation between intracellular current density and the transmembrane current in Eq. (3.3), this allows us to write

$$\beta I_m = \nabla \cdot (\sigma_i(\sigma_i + \sigma_e)^{-1}\sigma_e \nabla V_m) - \nabla \cdot (\sigma_i(\sigma_i + \sigma_e)^{-1})J_t. \tag{3.9}$$

The second term of the right hand side of Eq. (3.9) describes the contribution to the transmembrane current of the divergence of the total current. However, we know that $\nabla J_t = 0$, which implies that if the term $(\sigma_i(\sigma_i + \sigma_e)^{-1})$ is proportional to the identity matrix, the divergence

term vanishes. This happens if and only if σ_i and σ_e are related by a constant factor, such that $\sigma_e = \alpha\sigma_i$. If this is the case, the bidomain equations can be reduced to the monodomain equation:

$$\beta(C_m \frac{\partial V_m}{\partial t} + I_{ion}(\eta, V_m)) = \nabla \cdot (\sigma \nabla V_m), \tag{3.10}$$

where $\sigma = \sigma_i(\sigma_i+\sigma_e)^{-1}\sigma_e$. Thus, the monodomain requires the assumption that the intracellular and extracellular conductivities are related by a constant factor.

3.1.3 Numerical Solution Methods for the Monodomain and Bidomain Models

The PDEs associated with the bidomain (3.5) and monodomain (3.10) models cannot be solved analytically. To obtain approximate solutions, the computational domain on which they are solved is discretized. In cardiac electrophysiology, the Finite Difference Method (FDM), the Finite Volume Method (FVM) and the Finite Element Method (FEM) are the most commonly applied techniques. Less commonly, the Boundary Element Method (BEM), mesh free methods, and spectral methods have been used [138, 139, 140].

The FDM is based on approximating the partial derivatives of PDEs by finite difference equations, usually with the domain discretized into a regular grid. The FDM has been applied by many groups for studies of cardiac tissue electrophysiology [93, 141, 142]. Although the FDM is conceptually straightforward and easy to implement, the common grid-based formulation makes it difficult to describe smooth surfaces of the heart without steps, in turn making it difficult to implement boundary conditions [143].

For the FVM, volume integrals in a PDE involving a divergence component are converted to surface integrals using Gauss' theorem. This method preserves the flux of charge across volumes. The FVM can be applied easily to unstructured meshes and has been used in bidomain modelling by several groups [144, 145, 146]. Because of the flux conservation, a coupling

between excitation propagation in the heart and the electrical field in the body is also easy to formulate using the FVM [144].

The FEM is currently the most common numerical solution technique for the bidomain and monodomain equations. Here, the computational domain is divided into sub-domains known as finite elements. Inside each of these elements, the solution to the PDE is approximated by a suitably chosen interpolating function. The FEM is widely used because of its well-developed theory allowing for variable mesh representations and interpolation of material properties. A particular advantage is that complex, curved geometries can be modelled accurately, see for example [147]. The FEM method has been used by many groups in the field of cardiac modelling, especially for simulations with anatomically detailed geometries. The main disadvantage of the FEM compared to the FDM and FVM is the higher computational cost and more involved implementation [143].

The FDM, FVM and FEM applied to the bidomain and monodomain equations lead to large sparse, linear systems, that are typically solved using iterative numerical methods. Simulations of action potential propagation also require boundary conditions. These are usually no-flux boundary conditions and time dependent boundary conditions, representing current or voltage stimuli in the intracellular and extracellular space at given points in time. Dirichlet boundary conditions associated with applied potentials can generally be applied by some modification of the linear system once it is generated. For the FEM, which is the method used to solve the monodomain equation in Chapter 4, no-flux boundary conditions are naturally incorporated at the problem definition stage.

3.1.4 The Eikonal Equations

As we have seen, the monodomain provides a computationally lighter alternative to the bidomain model under assumptions that can be considered mild in most circumstances [93]. How-

ever, the monodomain model still relies on a full description of the reaction-diffusion like interplay between the ionic currents and the diffusion of charge within the tissue. Thus, for studies that exclusively concern the propagation of the cardiac activation wave through the tissue, further simplifications are both possible and justified.

A convenient way of describing the wavefront propagation is to define the excitation time $u(x)$ as the time at which the wavefront passes through given point x. The position of the wavefront at any time t is then given by the surface along which $u(x) = t$. A governing equation for u is referred to as an eikonal equation. Several eikonal equations have been derived specifically for describing the wave propagation in cardiac tissue. Keener [26] derived an eikonal-curvature equation by selecting a moving coordinate system such that V_m is a function of only a spatial variable normal to the wavefront and then requiring the current conservation equation to be satisfied at the wavefront. Colli-Franzone used singular perturbation theory to derive a related equation, known as the eikonal-diffusion equation [107]. The various eikonal formulations can be summarized using the following equation:

$$c_0\sqrt{\nabla t \cdot D\nabla t} - \kappa = \tau_m. \tag{3.11}$$

Here t is the activation time, c_0 is a dimensionless constant related to the upstroke velocity of the cell model, $\tau_m = C_m R_m$ where R_m is the membrane resistance, and κ is a correction term that takes different forms depending on the type of eikonal equation. D is the coupling tensor, defined according to $D = R_m\sigma$. For the eikonal curvature equation,

$$\kappa = \sqrt{\nabla t \cdot D\nabla t}\nabla \cdot (\frac{D\nabla t}{\sqrt{\nabla t \cdot D\nabla t}}). \tag{3.12}$$

This term corresponds to an anisotropic generalization of the mean curvature, lending the eikonal-curvature equation its name. The curvature term accounts for the increased speed of concave wavefronts and decrease in speed for convex wavefronts, reflecting the dependency

of tissue depolarization on diffusion of charge from nearby tissue that is already depolarized. Intuitively, a point just in front of a concave wavefront will be close to a larger amount of depolarized tissue than a point just in front of a convex wave front, and will thus depolarize faster in the concave case. The curvature term in Eq. (3.12) will be particularly important in highly isotropic settings. To see this, consider a wavefront propagating from a point source. Using the basic form of the eikonal equation, with $\kappa = 0$, the expanding wavefront will take on a strongly elliptic shape, with radii proportional to the inverse of the conduction velocities in the three principal directions. However, in the case of the eikonal-curvature equation, the significantly higher curvature in the direction of fastest propagation will result in a large penalty on conduction velocity, effectively making the elliptical wavefront more circular. Thus, the κ term introduced in Eq. (3.12) diminishes the effects on tissue anisotropy, especially during the initial phases of tissue activation, where anisotropic effects have a large impact on wavefront curvature.

For the eikonal-diffusion equation we have

$$\kappa = \nabla \cdot (D \nabla t), \tag{3.13}$$

describing the anisotropic Laplacian of the activation time, which is intimately related to diffusive phenomena. With respect to curvature, the diffusion term will have similar effects to the correction term of Eq. (3.12), but will additionally take into account effects from colliding wavefronts and domain boundaries. Finally, by setting $\kappa = 0$, we get the basic generalized eikonal equation, giving propagation velocities that are independent of the wavefront curvature. This is the eikonal formulation that will be used for the remainder of the work.

3.1.5 Numerical Solution Methods for the Eikonal Equation

A great number of methods have been published dealing with obtaining solutions to the eikonal equation, with applications spanning fields as diverse as geology, robotics, optics, computer graphics, computer vision and computational biology. Early efforts focused on obtaining numerical solutions to the isotropic version of the equation, corresponding to Eq. (3.11) with $\kappa = 0$ and D equal to the identity matrix. Here the equation was typically solved using finite difference and finite element discretizations. However, during the mid 1990s, a new approach known as fast marching methods (FMM) emerged, beginning with two independent publications by Tsitsiklis [148] and Sethian [149]. Early incarnations of the FMM relied on a regular discretization of the computational domain. However, over the past decades a number of FMM-like algorithms have been published, extending the original formulation to incorporate anisotropic and heterogeneous wave speed, effects from wave front curvature, wavefront interactions and unstructured computational meshes [150, 151, 152], resulting in several methods that have been specifically tailored for estimating the cardiac activation sequence.

The FMM relies on the use of an upwind finite difference approximation to the operators in the Eikonal equation. The structure of this finite difference approximation to the gradient yields a causality relationship that lends itself to an efficient programming approach. The Fast Marching Method is related to Huygens' principle [153], which is a construction involving expanding wavefronts, and to Dijkstra's method [154], which is an algorithm for computing first arrival times in a connected graph. It is thus instructive to review these two concepts before moving on to the FMM.

Huygens' Principle

Huygens' principle states that each point on an expanding wavefront can itself be seen as the source of a secondary wavefront expansion [153]. The wave position in a future time instant

can then be determined by constructing the envelope of these secondary wavefronts, and the process is repeated. Under isotropic conditions, these secondary wavefronts take on a circular shape, whereas under anisotropic conditions, they become elliptical. Starting from a wavefront at time t, Figure 3.1 shows a schematic illustration of the construction of a new wavefront at a later time $t + \Delta t$ using Huygens' principle.

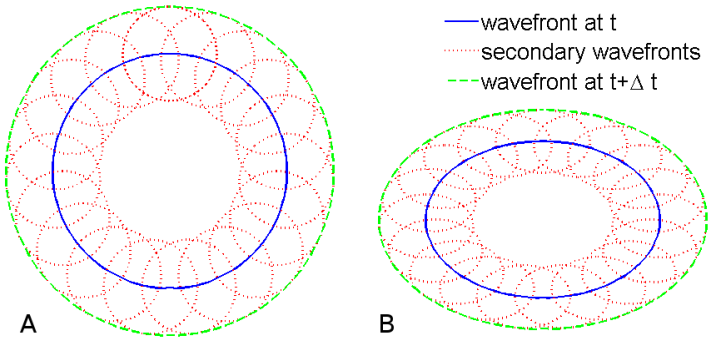


Figure 3.1: An illustration of Huygens' principle. Panel A shows the isotropic case, whereas panel B shows the transversely anisotropic case.

As can be seen in Fig. 3.1, for homogeneous conduction properties, Huygens' principle is valid both for isotropic and transversely anisotropic conduction.

Dijkstra's algorithm

As previously mentioned, Dijkstra's algorithm is a method of calculating first arrival times of a signal propagating along the edges of a connected graph [154]. It can be seen as a graph search equivalent of Huygens' principle. Formally, a graph is a set of nodes $\{n_i\}$ connected by a set of edges $\{e_{i,j}\}$. If $e_{p,q}$ is an element of the set $\{e_{i,j}\}$, we say that there is an edge from node n_p to node n_q, and that n_q is a neighbour of n_p. Every edge $e_{i,j}$ is associated with a positive cost $c_{i,j}$. A graph is said to be connected if there exists a path between any two nodes in it.

At the heart of Dijkstra's algorithm is a priority queue, $QUEUE$, commonly implemented in as a min-heap structure. The queue is used to determine the order in which to expand nodes n

to maintain a causal flow of information. In this context, expanding a node means finding its neighbours n_i, and calculating their arrival time $t(n_i)$, using information propagated from node n. The value of $t(n)$ is used to determine the position of a node n in $QUEUE$. For a start node n_0 and a goal node n_{goal}, a pseudo-code description of Dijkstra's algorithm can be found in Algorithm 1:

Algorithm 1 Dijkstra's Algorithm

$QUEUE \leftarrow QUEUE \cup n_0$
$t(n_0) \leftarrow 0$
while $QUEUE \neq \emptyset$ **do**
  $n_i \leftarrow argmin_{n \in QUEUE} t(n)$
  $QUEUE \leftarrow QUEUE \setminus n_i$
  for all $n_j \in NEIGHBOURS(n_i)$ **do**
    $t(n_j) \leftarrow min(t(n_j), t(n_i) + c_{i,j}$
    $QUEUE \leftarrow QUEUE \cup n_j$
  end for
end while

Here $\cup$ denotes set union, $\setminus$ denotes set complement, and $\emptyset$ denotes the empty set. The similarity with Huygens' principle lies in the node expansion step, corresponding to the expansion of Huygens' secondary wave fronts. However, in the case of continuous propagation, all points on a propagating wave front can be thought of as simultaneously expanding into secondary wave fronts, as illustrated in Fig. 3.1. Conversely, in the discrete case treated by Dijkstra's algorithm, the nodes need to be expanded at different discrete time points. Dijkstra's algorithm ensures a correct order of expansion by always selecting to expand the node with the lowest arrival time, efficiently implemented using the min-heap structure.

The Fast Marching Algorithm

The FMM can be viewed as a continuous space equivalent of Dijsktra's algorithm. The heart of the FMM, like that of Dijkstra's algorithm, is an identical priority queue, $QUEUE$. However, while Dijkstra's algorithm uses any acyclic connected graph as its computational domain, the

FMM uses a space divided into elements, like an FEM mesh. However, this distinction between these two types of domains is blurred by the fact that a space divided into polygonal elements can always be viewed as a connected graph. While Dijkstra's algorithm restricts expansion to occur along the graph edges, the expansion of the FMM is allowed to cut across elements, making its behaviour more similar to the workings of Huygens' principle. A pseudo-code description of the FMM can be found in Algorithm 2:

Algorithm 2 The Fast Marching Method

$QUEUE \leftarrow QUEUE \cup n_0$
$t(n_0) \leftarrow 0$
while $QUEUE \neq \emptyset$ AND $CHANGED \neq \emptyset$ **do**
  $n \leftarrow argmin_{n \in QUEUE} t(n)$
  $QUEUE \leftarrow QUEUE \setminus n$
  for all $n_i \in NEIGHBOURS(n)$ **do**
    $t_{new}(n_i) \leftarrow min_{p,q}(t(p,q) + \sqrt{v(p,q)^T D^{-1} v(p,q)}/F)$
    where
    $p, q, (p+q) \in [0,1] \times [0,1] \times [0,1]$
    $t(p,q) = t(n_j)(1-p-q) + t(n_a)p + t(n_b)q$
    $v(p,q) = v(n_i,n_j)(1-p-q) + v(n_i,n_a)p + v(n_i,n_b)q$
    if $t_{new}(n_i) < t(n_i)$ **then**
      $t(n_i) \leftarrow t_{new}(n_i)$
      if $n_i \notin QUEUE$ **then**
        $QUEUE \leftarrow QUEUE \cup n_i$
      end if
    end if
  end for
end while

Here, symbols are the same as in Algorithm 1. Additionally, D is a tensor creating anisotropy and F is a speed function, both varying spatially, and $v(n_i, n_j)$ is the vector going from node n_i to node n_j.

Extensions of the FMM

Under anisotropic conditions, a well known problem with the FMM comes from the fact that characteristic directions and gradients of the solution no longer coincide [150]. Thus, the al-

ready activated neighbourhood of a given node may not always contain the characteristic direction of propagation at the time it is computed. For certain combinations of element shape and anisotropic direction, this leads to a situation where some nodes can no longer be updated based according to the basic version of the FMM.

To address this problem, an anisotropic extension of the basic FMM was developed by Konukoglu and co-workers [155]. The basic structure of the algorithm is the same as that of the standard FMM presented in Algorithm 2. However, in order to overcome the problem introduced by anisotropic conduction, an extra step was added: Every time a node is pulled from QUEUE, a check is made to see if any of its already ACCEPTED neighbours can be updated to a lower arrival time based on the newly ACCEPTED node. If so, the updated neighbour is placed on a new queue, denoted CHANGED. Every time the main loop restarts it checks if the CHANGED queue is empty. If not, the algorithm picks the node with the lowest activation time from CHANGED instead of from QUEUE. This modification allows the possibility of correcting the errors introduced by the fact that the basic FMM does not guarantee a causal update order in the anisotropic case. A pseudo-code description of the algorithm is given in Algorithm 3 below.

Algorithm 3 has itself been extended to a multi-front formulation to account for cardiac repolarization effects [156], thereby allowing re-entry phenomena to be studied using the eikonal formulation. To facilitate this, a time counter was added to the algorithm. This allows each ACCEPTED node to be associated with an action potential duration, calculated as a function of the diastolic interval, i.e. the time elapsed since the node was previously activated. After initial activation, nodes are subsequently placed on a REFRACTORY list. While on the REFRACTORY list, nodes can no longer propagate or receive stimulus. When the time counter passes the action potential duration associated with a node, the node is removed from the REFRACTORY list and can again be placed on the QUEUE when reached by a stimulus.

Another development presented by Chinchapatnam *et al.* in [157] was aimed at extending the algorithm to treat solutions to the eikonal-diffusion equation. The initial step of the extended

Algorithm 3 Anisotropic Fast Marching

$QUEUE \leftarrow QUEUE \cup n_0$
$t(n_0) \leftarrow 0$
while $QUEUE \neq \emptyset$ AND $CHANGED \neq \emptyset$ **do**
  if $CHANGED \neq \emptyset$ **then**
    $n \leftarrow argmin_{n \in CHANGED} t(n)$
    $CHANGED \leftarrow CHANGED \setminus n$
  else
    $n \leftarrow argmin_{n \in QUEUE} t(n)$
    $QUEUE \leftarrow QUEUE \setminus n$
  end if
  $ACCEPTED \leftarrow ACCEPTED \cup n$
  for all $n_i \in NEIGHBOURS(n)$ **do**
    $t_{new}(n_i) \leftarrow min_{p,q}(t(p,q) + \sqrt{v(p,q)^T D^{-1} v(p,q)}/F)$
    where
    $p, q, (p+q) \in [0,1] \times [0,1] \times [0,1]$
    $t(p,q) = t(n_j)(1-p-q) + t(n_a)p + t(n_b)q$
    $v(p,q) = v(n_i, n_j)(1-p-q) + v(n_i, n_a)p + v(n_i, n_b)q$
    if $t_{new}(n_i) < t(n_i)$ **then**
      $t(n_i) \leftarrow t_{new}(n_i)$
      if $n_i \in ACCEPTED$ **then**
        $CHANGED \leftarrow CHANGED \cup n_i$
        $QUEUE \leftarrow QUEUE \setminus n_i$
      else if $n_i \notin QUEUE$ **then**
        $QUEUE \leftarrow QUEUE \cup n_i$
      end if
    end if
  end for
end while

algorithm is identical to the anisotropic solver described above. However, once a solution for the activation times is obtained, this is used to compute the diffusion term in Eq. (3.13). Subsequently, this diffusion term is added to τ in Eq. (3.11), and the anisotropic solver is applied to the modified equation. The solution to the eikonal-diffusion equation is obtained by continuing these so-called fixed point iterations until the magnitude of change in the computed activation times falls below a pre-specified threshold.

3.2 Parameter Estimation

In order to adapt the electrophysiological models described in the previous section to mimic the activation of a specific heart, the estimation of conduction parameters (as reviewed in Section 2.7) is a crucial step. This dissertation will focus on estimation of conduction velocities, but to put the parameter estimation efforts presented in Chapters 5 and 6 into a greater context, we will start by describing the general parameter estimation problem.

3.2.1 The Parameter Estimation Problem

Let $\mathfrak{S}$ be the physical system under study. For the purposes of this work, this system will correspond to the ventricles of the heart. In order to gain information about the properties of the system, experimentally observed functional and structural information from a large range of spatiotemporal scales is collected. After a sufficient quantity of information has been collected, a hypothesis is formed about some aspects of how the different parts of the system interact, allowing predictions to be made regarding how the system will react to new stimuli or how its behaviour will evolve over time. Such a hypothesis, when formalized, is what we will call a model of the system. In the cardiac modelling community, formalized models are usually expressed in the form of systems of differential equations. By solving the system of equations, we obtain a corresponding formal prediction of how the system behaves over time and in response to various stimuli. Based on how well the predicted behaviour of the system corresponds to the observed behaviour, we might later choose to adjust some aspects of the model in order to improve our predictions.

As a fairly coarse division of the parts of a mathematical model we may say that the qualitative understanding of the system under study defines the structure of the model, while quantitative measurements are used to assign values to the model parameters. This gives rise to a nested structure, of which the highest level is an abstract space or manifold of all conceivable models

$\mathcal{M}_1, \mathcal{M}_2 ... \mathcal{M}_N$. We will call this the *model space* and denote it $\mathfrak{M}$. Associated with each point $\mathcal{M}_i$ in $\mathfrak{M}$ is a set of parameters associated with the particular model. This again defines a manifold, but this time over conceivable parameter values for a particular model $\mathcal{M}_i$. This will be called the *parameter space* and is denoted $\mathfrak{Q}_i$[1]. (For the discussion below we will drop the indexing of $\mathcal{M}_i$ and $\mathfrak{Q}_i$, and thus assume that we have decided on a model for our system, along with a suitable set of governing parameters.)

Associated with each point in $\mathfrak{Q}$ is a set of numerical values, $\{q_1, q_2, ...q_M\}$, representing a set of *coordinates* in $\mathfrak{Q}$. Here it should be noted that the number of parameters needed to completely describe a given model may not necessarily be finite. This number is infinite, for instance, when we are interested in a property that depends continuously on the spatial position inside a volume. However, for the remainder of this work, only finite-dimensional models will be considered.

To obtain information on model parameters, data from physical experiments need to be obtained. These data will come from measurements of different features of the physical system, and like with the parameters, each species of measurements will be associated with some range of conceivable values. This prompts us to define a data space $\mathfrak{D}$, in which each point has a set of numerical values, $\{d_1, d_2, ...d_M\}$, representing a set of coordinates in $\mathfrak{D}$.

If our model structure is sound, that is if the model output will be identical to that of the physical system for some set of parameters, we say that the model constitutes a *map* between the parameter space and the data space, $\mathcal{M} : \mathfrak{Q} \mapsto \mathfrak{D}$. For each set of model parameter values, corresponding to a set of coordinates in $\mathfrak{Q}$, there will thus exist a corresponding point in $\mathfrak{D}$. Constructing the model and solving the corresponding system of equations is often referred to as the *forward problem*. Conversely, the field of parameter estimation deals with the problem of finding (or at least investigating some aspects of) a map that goes in the opposite direction, $\mathcal{I} : \mathfrak{D} \mapsto \mathfrak{Q}$, for all or certain subspaces of $\mathfrak{D}$. This is commonly referred to as the inverse prob-

[1] In some of the modelling literature, different numerical values for the same parameter(s) are referred to as different models. In this work, the term "model" will be reserved for a given set of equations, as this seems better aligned with the common use of the word.

lem, and characterizing this map constitutes a fundamental problem of parameter estimation.

A map, which for the purposes of this discussion corresponds to a model, is *injective* if every point in the data space $\mathfrak{D}$ maps to at most one point in the parameter space $\mathfrak{Q}$. A map is *surjective* if every point in $\mathfrak{D}$ maps to at least one point in $\mathfrak{Q}$. A map that is both injective and surjective is called *bijective*. Figure 3.2 shows schematic examples of a bijective and a non-injective, non-surjective map. Maps that are not bijective will simply be referred to as *non-bijective*.

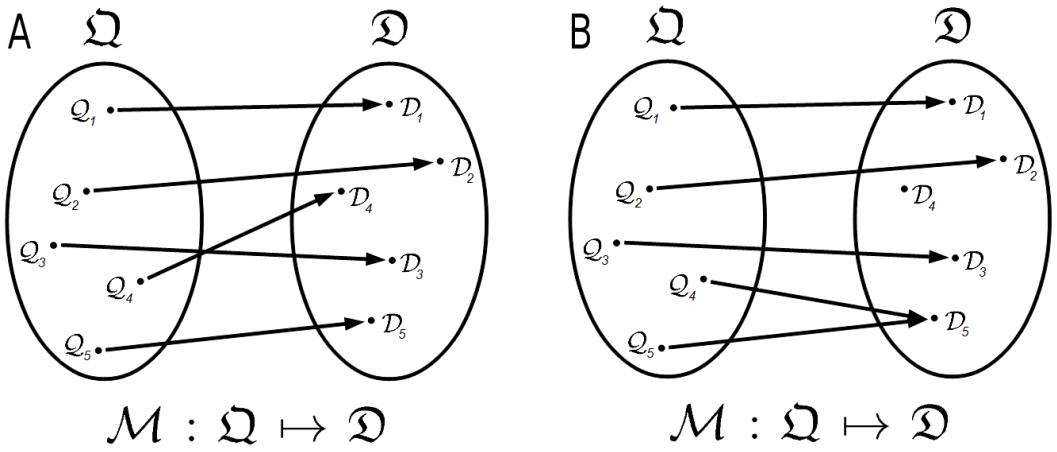


Figure 3.2: Panel A shows a bijective map between the parameter space and the data space of the model. Panel B shows a non-bijective map.

As explained above, when estimating parameters from data, we essentially want to invert the map defined by the model. In this case, non-bijective mappings pose a problem. Looking at panel B of Fig. 3.2, it is obvious that even if we were able to invert the mapping, results would be ambiguous. This is an example of an *ill-posed* problem, which is often seen in parameter estimation. Hadamard gave three classic criteria for a problem to be well-posed [158]: 1) the solution must exist, 2) the solution must be unique, and 3) the solution must depend continuously on the data, i.e. not exhibit extreme sensitivity to initial conditions.

However, even in the case of a bijective map, practical problems arise, some of which can make the problem ill-posed. Firstly, for real world applications, the model is seldom analytically in-

vertible. This is a well known difficulty that can be overcome by numerical techniques, some of which will be further described later in this chapter. Secondly, in many cases, we only have access to measurements from a low dimensional projection $\mathfrak{D}'$ of a high dimensional data space $\mathfrak{D}$. This is the case for the problem treated in this dissertation, that of estimating cardiac conduction properties from electroanatomical mapping data. Here, $\mathfrak{D}$ consists of the activation times in all points in the cardiac tissue, each of which can vary over a range of possible numerical values. However, we only have access to data from a limited set of points on the endocardium and/or epicardium. Since we have no information about the values in the remaining points, the data that we do obtain comes from $\mathfrak{D}'$ rather than $\mathfrak{D}$. This might render indistinguishable some points in $\mathfrak{D}'$ which are clearly separated in $\mathfrak{D}$, making the bijective map look non-bijective. This is illustrated in panel A of Fig. 3.3. A different problem leading to similar consequences stems from the fact that data often contain significant amounts of noise. This is a second way of making a bijective map appear non-bijective, since random chance might make two closely situated points in $\mathfrak{D}$ appear identical, even though they are mapped to by different points in the parameter space $\mathfrak{Q}$. This is illustrated in panel B of Fig. 3.3.

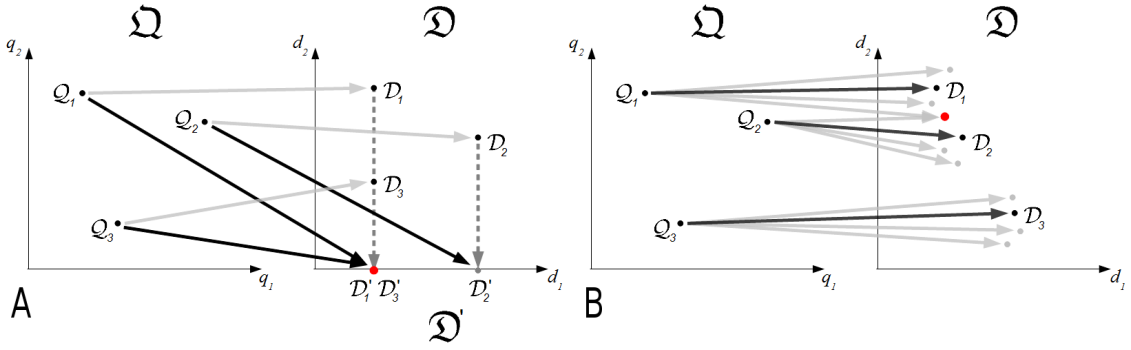


Figure 3.3: Both panels show bijective maps between a set of points $\mathcal{Q}_1$-$\mathcal{Q}_3$ in a 2D parameter space and a set of points $\mathcal{D}_1$-$\mathcal{D}_3$ 2D data space. Panel A additionally shows projected points $\mathcal{D}'_1$-$\mathcal{D}'_3$ onto the first dimension of the data space, with a larger red dot marking the coordinates where $\mathcal{D}_1$ and $\mathcal{D}_3$ appears to coincide. Panel B additionally shows each data point (black dots) represented by several points with added noise (light grey dots). A larger red dot marks a set of coordinates where noisy data from $\mathcal{D}_1$ and $\mathcal{D}_2$ appears to coincide.

All of the factors listed above contributes to the ill-posed nature of practical parameter estima-

tion problems. As we have seen, an ill-posed problem either does not have a specific single solution (Hadamard's first and second criteria are not met), or the problem of finding it is too hard to solve at all (Hadamard's third criteria is not met). Although the latter form of ill-posed problem constitutes an interesting and relevant challenge for the field of cardiac modelling (e.g. for the study of chaotic behaviours such as fibrillation), it will not be treated further in this work. Rather, we will focus on the challenge of estimating parameters from noisy data in a low dimensional projection of the data space.

3.2.2 Bayesian Inference

Bayesian inference [14, 15] provides the main tool employed in this work of handling ill-posed parameter estimation problems. As noted by Bretthorst in [159], Bayes theorem can be interpreted as a restatement of the fact that Aristotelian logic is commutative. As an example, there cannot be any difference in truth value between the following two propositions:

(Q,D) = "both Q and D are true"

(D,Q) = "both D and Q are true" .

Expressed as probabilities, this becomes $P(Q, D) = P(D, Q)$. Applying the product rule of probability, this allows us to write

$$P(Q|D)P(D) = P(D|Q)P(Q), \tag{3.14}$$

which after a trivial rearrangement of terms gives us the common form of Bayes theorem:

$$P(Q|D) = \frac{P(D|Q)P(Q)}{P(D)}. \tag{3.15}$$

In common terminology of the statistical literature, $P(Q|D)$ is called the *posterior probability* (or simply posterior) of Q, given D. $P(D|Q)$ is called the *likelihood function* (or simply likelihood), $P(Q)$ is called the *prior probability* (or simply prior), and $P(D)$ is called the *evidence* or prior probability of the data.

When applying Bayesian inference to the problem of parameter estimation, the posterior represents the probability associated with any point $\mathcal{Q}$ in our parameter space $\mathfrak{Q}$, conditional on an observed point $\mathcal{D}$ in our data space $\mathfrak{D}$. In other words, the posterior is a probability density function in $\mathfrak{Q}$, and finding it solves the parameter estimation problem. Here, it should also be noted that a consequence of posing the parameter estimation problem in this way is that it becomes "immune" to ill-posedeness stemming from Hadamard's first and second criteria; if there does not exist a point in $\mathfrak{Q}$ mapping to $\mathcal{D}$, the posterior simply becomes zero everywhere, and if several points in $\mathfrak{Q}$ map to $\mathcal{D}$, the posterior will take on non-zero values in all of these points. In both cases, the solution exists and is unique.

However, as mentioned in the previous section, our model is very seldom analytically invertible, and thus the posterior is very likely not to be directly computable. Here, Bayes theorem provides a useful tool, since it tells us how to relate the inverse problem to the forward problem. Looking at the right hand side of Eq. (3.15), the likelihood, $P(D|Q)$, tells us what points in $\mathfrak{D}$ we are likely to observe, conditional on a set of parameters corresponding to a point in $\mathfrak{Q}$. As we shall see later in this section, this term is computable by solving the forward problem using our model. The prior, $P(Q)$, represents our state of knowledge prior to observing any data. If we have reason to believe that some parameters are more likely than others, for example by knowledge of parameter data from previous studies, we can incorporate this information in the Bayesian inference procedure via the prior. If no such information is available, a non-informative prior is commonly used [14]. A non-informative prior is one that does not induce a bias towards assigning any particular point in $\mathfrak{Q}$ a higher or lower probability. A common choice is to assign all points in $\mathfrak{Q}$ equal probability, but alternatives have been proposed, such

as Jeffreys' prior [160].

The evidence, $P(D)$, can be expanded as an integral over $\mathfrak{Q}$ using the law of total probability:

$$P(D) = \int_Q P(D|Q)P(Q)dQ. \tag{3.16}$$

Regardless of how difficult Eq. (3.16) is to compute for a given problem, the structure of the equation shows that the evidence is independent of which particular point in $\mathfrak{Q}$ we choose to look at. Thus, for a particular choice of model, the evidence will never alter the shape of the posterior probability distribution, only scale it, acting as a normalizing constant.

Since the prior is by definition already known, and since the evidence reduces to a normalizing constant, the main effort of Bayesian inference, given a fixed choice of model, lies in computing the likelihood term. As we saw in the previous section, a model $\mathcal{M}$ defines a map between each point $\mathcal{Q}$ in the parameter space and a corresponding point $\mathcal{D}$ in the data space. In order to compute the likelihood, we must thus find the region of $\mathfrak{Q}$ that maps to our observed point $\mathcal{D}$ in data space. In the case of a bijective map that does not suffer from any of the noise and dimensionality problems illustrated in Fig. 3.3, the likelihood will correspond to a Dirac delta function in $\mathfrak{Q}$. However, with the presence of noise in $\mathcal{D}$, the corresponding region of $\mathfrak{Q}$ becomes less well defined.

For the rest of this discussion, we will make three key assumptions: 1) that the data space and parameter spaces are linear, 2) that the noise associated with $\mathcal{D}$ is Gaussian, and 3) that the noise distribution is independent between dimensions of $\mathfrak{D}$. This defines a multivariate Gaussian PDF in $\mathfrak{D}$, having its centre in $\mathcal{D}$. Since each point $\mathcal{Q}$ maps to some point in $\mathfrak{D}$, each point $\mathcal{Q}$ also maps to a probability defined by the PDF over $\mathfrak{D}$. This tells us that the probability associated with any point $\mathcal{Q}$ can be computed according to

$$P(\mathcal{D}|\mathcal{Q}) \propto e^{-\frac{1}{2}(\mathcal{D}-\mathcal{M}(\mathcal{Q}))^T C(\mathcal{D}-\mathcal{M}(\mathcal{Q}))}, \tag{3.17}$$

where $\mathcal{D} = \{d_1, d_2, ...d_N\}$ and $\mathcal{Q} = \{q_1, q_2, ...q_N\}$ as explained in the previous section, and C is a diagonal matrix of variances for each dimension in $\mathcal{D}$. This reduces the evaluation of Eq. (3.15) to a search problem in $\mathfrak{Q}$. By finding the regions corresponding to non-negligible probabilities, and evaluating those probabilities, we map the PDF over $\mathfrak{D}$ to a PDF over $\mathfrak{Q}$ and solve the parameter estimation problem.

Marginalization and Credibility Regions

A common tool in Bayesian inference is marginalization, which is used to reduce the dimensionality of a distribution by integrating over some of its dimensions. Marginalization can be defined as

$$P(q_i|\mathcal{D}) = \int_{q_j} P(\mathcal{Q}|\mathcal{D}) dq_j, \tag{3.18}$$

where q_i is the parameter(s) of interest and q_j signifies all other parameters. Marginalization is commonly used for two purposes: visualization and elimination of so called *nuisance parameters*, i.e. parameters that are needed for the models but of no interest from a parameter estimation point of view. The procedure is illustrated in panel A of Fig. 3.4.

Another useful concept, commonly encountered in connection with Bayesian inference is *credibility regions*. A credibility region is the Bayesian counterpart of the more traditional confidence intervals. In this work, we will use a level set definition of the credibility region:

$$CR_c = \left\{ q_i \mid 100 \int_{q_i} P(q_i|\mathcal{D}) dq_i = c \wedge P(q_i|\mathcal{D}) \geq h \right\}. \tag{3.19}$$

Here, c is the desired level of confidence and h is the unique level at which the right hand side expression integrates to c. This formulation allows the confidence region to consist of

disjoint regions of q_i, making it appropriate for analysis of potentially multimodal PDFs, such as those that will be encountered below in Chapters 5 and 6. The procedure of finding the CR is illustrated in panel B of Fig. 3.4.

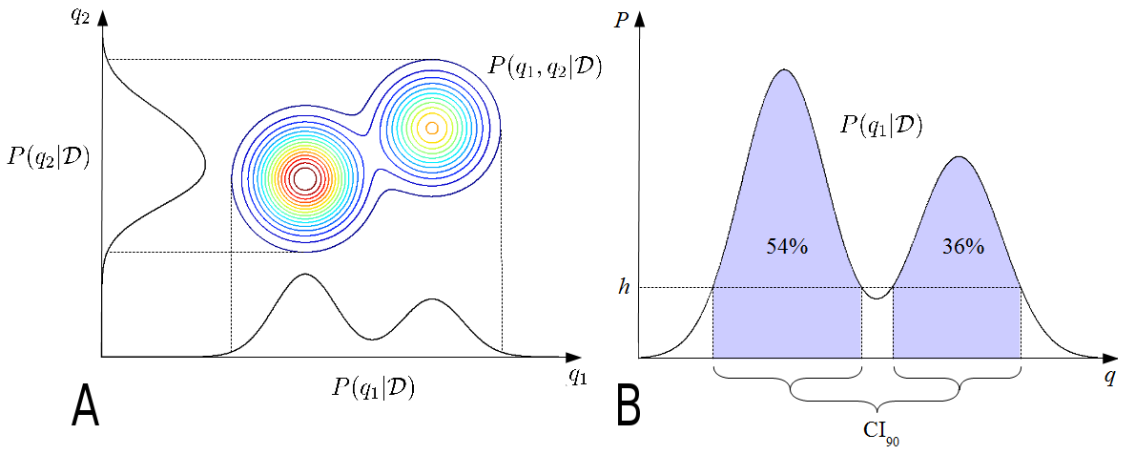


Figure 3.4: Panel A illustrates the process of marginalizing a 2D PDF in Ω over each of its two dimensions, q_1 and q_2, rendering two 1D PDFs. Panel B illustrates the process of finding the 90% CR (CR_90) for one of the marginalized PDFs from panel A. The level of the parameter c is indicated by a dashed line, and the resulting CR is marked using curly brackets ({) and blue shading. Also shown are the distribution of probability masses between the two parts of the disjoint CR.

4

Comparison between the eikonal, graph based and monodomain methods for activation time approximations

As we saw in the previous chapter, Bayesian inference provides a powerful framework for understanding the parameter estimation problem, as well as a practical means to gaining information about model parameters from measured data. However, in practice the inference procedure requires a large number of model evaluations, putting high demands on the computational efficiency of the model. In this chapter we develop a model relying on a graph search approach to estimate cardiac activation times with minimal computational cost, and compare its performance to that of monodomain and the FMM. The chapter is organized into four main sections, beginning with a brief introduction in Section 4.1. In Section 4.2 the proposed graph based method is described, along with a description of the spatial domains in which the method has been tested. Finally, details are given on the study designed to investigate the accuracy of the graph based method. Section 4.3 starts with results from the monodomain simulations performed as a gold standard by which to judge the accuracy of the two phenomenological methods. Following this, convergence properties of the graph based methods are studied, and the similarity of graph based activation times to those from monodomain and the FMM

is assessed. In Section 4.4 the presented results are discussed with respect to the fitness of the method to the purpose of model based parameter estimation. Finally, in Section 4.5, conclusions are drawn, setting the scene for the following chapters.

4.1 Introduction

As outlined in Chapter 3, solving the bidomain and monodomain equations is a computationally expensive task [161, 162, 163], making them unsuitable for situations where cardiac activation times need to be determined quickly. As stated in Section 3.1.4, phenomenological alternatives exist for applications where the main goal is to predict the activation sequence, for example estimation of cardiac conduction parameters [125] or predictive and interactive simulations for clinical interventions [100, 152]. Currently, the most popular approach is the eikonal equation. A major advantage of using the eikonal equation, rather than the bidomain or monodomain equations, is that there exists computationally efficient numerical solution methods, such as the FMM described in Section 3.1.5.

While the FMM provides a computationally efficient option for estimating cardiac activation times, further simplifications can be made for specific settings. One such setting is when the activation times at single or a small number of points need to be estimated, for example to relate model simulations to clinical data collected at those points. In this context, heuristic algorithms represent a more computationally efficient alternative to the FMM, since they are specifically designed to limit the amount of calculations to a minimum when computing activation times at individual points. However, so far heuristic methods have gained little notice within the cardiac simulation community. This chapter investigates the application of one such method, the A* algorithm [164], for finding discrete paths within a graph representing the cardiac tissue.

Specifically, the aim of this chapter is to show that graph based methods can provide a faster alternative to the FMM for estimating activation times in cardiac tissue during electrical exci-

tation, and that heuristic methods such as the A* algorithm can give an added speed advantage in situations where the activation times in only a limited number of points are required. We demonstrate this advantage by comparing the graph based method to the anisotropic FMM and the monodomain model, as described in Chapter 3, with respect to performance and accuracy. The comparison is done in a rabbit ventricular mesh, incorporating realistic fibre directions and a conduction heterogeneity, representing a transmural scar. Results demonstrate that the proposed method is able to predict cardiac activation times with good accuracy and fast computational performance. The work described in this chapter builds on methodology and results previously described by the author and coworkers in [16] and [17].

4.2 Methods

4.2.1 Graph Based Activation Time Approximations

As we have seen in Chapter 3, the bidomain, monodomain and eikonal equations are commonly solved on domains tesselated into triangles or tetrahedra. Conversely, the graph based method presented in this section is based on describing the cardiac tissue as a connected graph. A tesselated domain, such as a mesh designed for FEM computations, can always be viewed as a connected graph. Recapitulating the definitions related to graphs given in Section 3.1.5, we define a graph as a set of nodes $\{n_i\}$ connected by a set of edges $\{e_{i,j}\}$. If $e_{p,q}$ is an element of the set $\{e_{i,j}\}$, we say that there is an edge from node n_p to node n_q, and that n_q is a neighbour of n_p. Every edge $e_{i,j}$ is associated with a positive cost $c_{i,j}$. Further to these definitions, a *path* from n_1 to n_k is an ordered set of nodes $\{n_1, ..., n_k\}$ with each n_{i+1} being a neighbour of n_i. The cost of a path is obtained by adding up the individual costs of each edge $e_{i,i+1}$ of that path. A graph is said to be *connected* if there exists a path between any two nodes in it. The *connectivity* of a node is equal to its number of neighbours.

Under completely isotropic conditions, any edge cost $c_{i,j}$ will be always be proportional to the distance between n_i and n_j. However, in the case of using the graph to model cardiac tissue, the costs also need to depend on the local fibre direction. Formally, this has been taken into account by associating each edge with a fibre direction, after which its cost has been calculated according to

$$c_{i,j} = \sqrt{v(n_i, n_j)^T D^{-1} v(n_i, n_j)}/\tilde{F}, \tag{4.1}$$

where $v(n_i, n_j)$ is the vector corresponding to edge $e_{i,j}$, going from node n_i to node n_j and $\tilde{F} = F/\delta$. Here, δ is a scaling factor compensating for the added cost associated with the discrete nature of using a graph representation of the domain, while F is the speed function. $\tilde{F}$ effectively acts as a scaling parameter for the tensor D that describes the anisotropic conduction properties of the cardiac tissue. D, in turn, is defined according to

$$D = \begin{bmatrix} x_l & x_t & x_n \\ y_l & y_t & y_n \\ z_l & z_t & z_n \end{bmatrix} \begin{bmatrix} cv_l^2 & 0 & 0 \\ 0 & cv_t^2 & 0 \\ 0 & 0 & cv_n^2 \end{bmatrix} \begin{bmatrix} x_l & x_t & x_n \\ y_l & y_t & y_n \\ z_l & z_t & z_n \end{bmatrix}^T, \tag{4.2}$$

where $[x_i, y_i, z_i]$ defines vectors in the longitudinal, transverse and normal direction of the local myocyte orientation, while cv_i defines the conduction speed in these three directions. Thus, the values represented by D and F are used to describe the spatially varying conduction properties of the cardiac tissue.

When viewing the cardiac tissue as a connected graph, the problem of approximating the activation time for a given point in the cardiac tissue, given an initial stimulation, is transformed into finding a path through the corresponding graph based description of the tissue, going from an initial node n_s to a goal node n_g. As we saw in Section 3.1.4, this problem can be solved using Dijkstra's algorithm. However, if we are only interested in a approximating the activation

time in single nodes, a more computationally efficient solution is available, in the form of the A* algorithm [164].

The A* algorithm is structurally very similar to Dijkstra's algorithm, but incorporates the use of prior information to find the fastest path between n_0 and n_{goal} in a more efficient manner. As with Dijkstra's algorithm (Algorithm 1), the heart of the A* is a priority queue $QUEUE$. However, instead of ordering $QUEUE$ according to the ascending computed arrival times of the nodes in $QUEUE$, the A* algorithm employs a function $g(n_j)$ for the ordering. $g(n_j)$ is calculated as the sum of the computed cost (identical to that of Dijkstra's algorithm) and a heuristic function $h(n_j)$. $h(n_j)$ is an estimate of the cost of the lowest cost path from n_j to the specified goal node n_{goal}, and $g(n_j)$ thus becomes an estimate of the total cost of a path from n_0 to n_{goal}, passing through n_j. This greatly reduces the number of nodes that need to be expanded, generally leading to much improved computational performance compared to Dijkstra's algorithm when single goal nodes are considered. For the present application, the heuristic estimate has been calculated as $h(n) = ||v(n, n_{goal})||cv_l^{-1}$, where cv_l^{-1} is the cost (in time per distance) of conduction in the fibre direction.

Algorithm 4 The A* Algorithm

$QUEUE \leftarrow QUEUE \cup n_0$
$t(n_0) \leftarrow 0$
$f(n_0) \leftarrow h(n_0)$
while $n_i \neq n_{goal}$ **do**
  $n_i \leftarrow argmin_{n \in QUEUE} f(n)$
  $QUEUE \leftarrow QUEUE \setminus n_i$
  for all $n_j \in NEIGHBOURS(n_i)$ **do**
    $t(n_j) \leftarrow min(t(n_j), t(n_i) + c_{i,j}$
    $g(n_j) \leftarrow t(n_j) + h(n_j)$
    $QUEUE \leftarrow QUEUE \cup n_j$
  end for
end while

4.2.2 Convergence Properties of the Graph Based Method

The convergence properties of several versions of the FMM have been examined in a number of publications, see for example [149, 150, 151, 165]. These and other publications demonstrate that the FMM converges towards the viscosity solution of the eikonal equation when decreasing the element size of the mesh. Conversely, the graph based approximations depend on both decreasing element size and increasing node connectivity for convergence. This has been noted in for example [166]. To investigate the convergence properties of the graph based approach, a cuboid geometry measuring 3x3x3 arbitrary space units was constructed. The cuboid consisted of an unstructured network generated by randomly perturbing the nodes of a structured cubic grid consisting of 8000 nodes. Each direction (x,y, and z) was associated with a cost of travelling (1, 2, and 3 arbitrary time/space units, respectively). Next, the analytical arrival times were calculated according to

$$t_i^{analytical} = \sqrt{x^2 + (2y)^2 + (3z)^2}, \tag{4.3}$$

thus assuming that the propagation was initiated at the point (0,0,0). These analytical arrival times were subsequently compared to corresponding times computed using Dijkstra's algorithm in the graph under 9 different levels of node connectivity. The levels of connectivity were defined by introducing a parameter r, defining a radius of connectivity around each of the 8000 nodes. The parameter r defines a radius for each node, and every node within that radius is connected to the central node. Thus, by increasing r, the connectivity of each node also increases. Results from this comparison are presented in Section 4.3.2, where the value of r is varied from 0.2 to 1.0 space units in steps of 0.1.

For comparison against a reaction-diffusion setting, the cube was then further subdivided into 64000 nodes to facilitate solving the corresponding problem using a finite difference discretization of the monodomain equation. Simplified cellular level dynamics were described using the

Mitchell-Schaeffer model [77]. The conductivities in the x,y, and z directions were set proportional to 9, 4 and 1 respectively, to match the relation between the conduction velocities of the path finding and analytical solutions described above. The cube was again activated from the corner (0,0,0), using intracellular stimulus of a small cubic region with a side of 3 nodes. From the monodomain solution, activation times for each node were subsequently defined as the time that membrane potential state reached 90% of its peak value with a positive time derivative.

4.2.3 Anatomical Model

As a basis for the main part of the study presented here, an MRI-derived rabbit ventricular geometry was used (see Fig. 4.1), representing a computational domain that was deemed complex enough to be able to highlight potential shortcomings of the proposed method. The geometry was discretized into nine different resolutions, rendering nine tetrahedral meshes, comprising around 16000, 20000, 25000, 30000, 37000, 48000, 64000, 80000 and 4000000 nodes respectively. For these resolutions, average element sizes were around 750, 700, 650, 600, 550, 500, 450, 400 and 150μm respectively. The seven lowest resolution meshes were created by downsampling the 80000 node mesh, using the following procedure. First, the surface of the 80000 node mesh was used to produce a voxel stack, utilizing a procedure based on triangle-line intersections described in [167]. Subsequently, the voxel stack was used as input to the proprietary meshing tool Tarantula (www.meshing.at), which was applied to generate a coarser tetrahedral mesh. A combination of voxel size and input parameters for Tarantula was used to produce the desired mesh resolution. This procedure resulted in some minor structural differences between the meshes, mainly located at the base of the left ventricle. The 4000000 node mesh was used for the monodomain simulations to ensure an accurate numerical solution of the monodomain equation, while the coarser resolution meshes were used for the FMM and pathfinding algorithms.

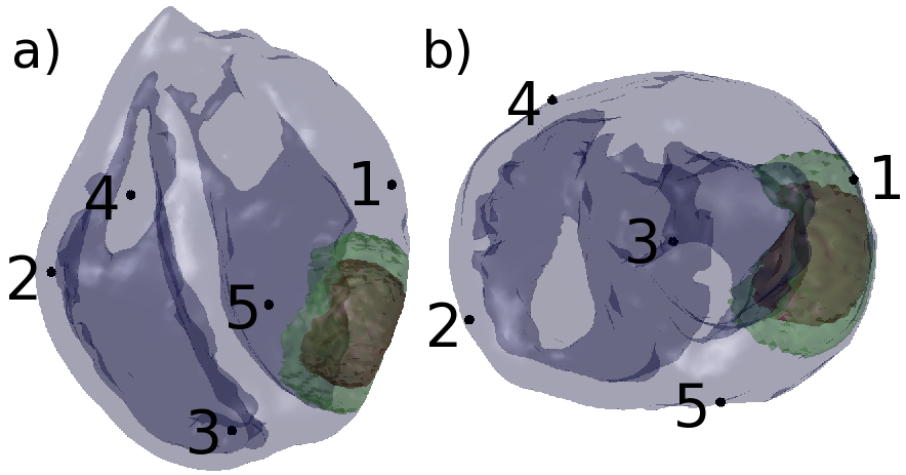


Figure 4.1: The rabbit ventricular model. Panel a) shows the geometry from the side, and panel b) shows the geometry from above. The numbers correspond to different stimulation points, while the darker area of the left ventricle corresponds to a transmural scar.

With each element of the fine mesh was associated a conductivity tensor, describing the cardiac fibre orientation. This tensor field was algorithmically generated according to the Streeter model [93, 168], using the CHASTE simulation package [169].

Additionally, a transmural conduction heterogeneity, corresponding to a ventricular scar, was incorporated. The scar consisted of border zone with low conductivity, and a passive core zone, mimicking the two zone configuration reviewed in Section 2.4. The conduction in the border zone was modelled as anisotropic but lowered in agreement with experimental results from [34]. Side and top views of the mesh (panels a and b respectively), incorporating the scar region as well as 5 stimulation locations, are depicted in Fig. 4.1.

All monodomain simulations presented in this work were performed using the Mahajan cell model [71], characterizing the cellular electrophysiology of the rabbit cardiac ventricular myocyte. Conductivity values were 0.132, 0.066 and 0.033 S/m in the longitudinal, transverse and normal directions respectively, lying in the range reported in literature [110, 111, 112, 119]. The border zone of the transmural scar was modelled using a uniform decrease in conductivity by a factor of 8, while the core zone was modelled as completely passive with conductivities

the same as in the border zone. The simulations were performed using the CHASTE simulation package.

4.2.4 Stimulation Protocols

A total of seven different stimulation protocols were employed, all relying on the five stimulation points depicted in Fig. 4.1. Out of these seven stimulation protocols, five initiated electrical activation from a single point, while two initiated activation at points 1 and 3 simultaneously. Of these two last protocols, one was performed in the presence of the transmural scar depicted in Fig. 4.1. The single point stimulation protocols were aimed at spanning the principal directions of current propagation for purposes of robustness testing, corresponding to Section 4.3.3 below, while the two dual point stimulations mimic the more clinically relevant case of biventricular pacing [170], presented in Section 4.3.5 below.

4.2.5 Error Calculation and Data Fitting

For the monodomain simulations, one activation time was obtained for each node. The activation time at a given node was defined as the time at which its membrane potential crossed 0 mV with a positive time derivative. For comparison with corresponding activation times from the FMM and the graph based method, activation times obtained using the monodomain simulations on the 4000000 node mesh were mapped onto the coarser meshes using the nearest neighbour method. The relative errors could then be calculated as $E_i = (\hat{t}_i - t_i)/t_{\text{tot}}$, where E is relative error, i is the node index, $\hat{t}$ is the approximate activation time obtained from the FMM or the graph based method, t is the correct activation time obtained from the monodomain simulation, and t_{tot} is the time it takes for all the nodes in the considered mesh to become activated.

Based on these relative errors, RMS errors were calculated according to

$$RMS = \sqrt{\frac{\sum_{i=1}^{n}(E_i)^2}{n}} \tag{4.4}$$

where n is the number of nodes in the mesh. The RMS errors were subsequently used for the estimation of conduction costs, corresponding to the inverses of cv_l, cv_t and cv_n in Eq. 5, in the cross validation procedure presented below. Specifically, the RMS error was used as an objective function in a minimization procedure based on the Nelder-Mead simplex method, as implemented in MATLAB.

4.2.6 Cross Validation Procedure

To facilitate the part of the study presented in Section 4.3.3 below, a cross validation procedure was employed. Five single point simulations were performed on the model, placed according to Fig. 4.1, as described in section 4.2.4. Using the resulting activation times, conduction costs for both the FMM and the graph based method were fitted as described in the previous section. This resulted in five sets of conduction costs for each method. Using each of these sets of conduction costs, both methods were used to estimate the activation sequence resulting from each of the five stimulation points. This resulted in 25 sets of activation times for each method. Each of these 25 sets of activation times was subsequently compared to the corresponding activation times obtained from the monodomain simulations, giving 25 RMS errors for each method, again calculated as described in 4.2.5.

4.3 Results

The following section is divided into five different parts. The first part is a brief description of the results from the monodomain simulations used as a basis for the study, while the remaining five sections each investigates some aspect of the graph based method. The first of these, Section

4.3.2, investigates how the graph based method converges with increasing node connectivity in a simplified geometry. The second part, Section 4.3.3, consists of a cross validation of the graph based method and the FMM, comparing the accuracy and robustness of the two methods when using one monodomain simulation for fitting the conduction costs and another one for estimating activation times. Using the optimal cost functions thus obtained, a study of the accuracy of the method at different graph resolutions is performed in Section 4.3.4. Section 4.3.5, investigates the performance of the graph based method and the FMM for a set of known cost functions in a more complex setting, incorporating multiple points of initial activation, as well as a transmural scar. Finally, in Section 4.3.7 the computational performance of the graph based method is investigated, and again compared to the FMM.

4.3.1 Monodomain Simulations

As described in section 4.2.4, seven monodomain simulations were performed as a basis for validation of the proposed method for activation time approximation. The data resulting from these simulations were exclusively used to calculate activation times for each node. Total activation time for the model was around 63 ms. As an example of the resulting activation patterns, Fig. 4.2 shows activation isochrones on the anterior and posterior of the epicardium, following simultaneous stimulations at points 1 and 3 (see Fig. 4.1). Panel a) clearly shows the elliptical shape of the activation wave propagating from stimulation point 1, illustrating the effect of the anisotropic conduction properties of the tissue. In panels c) and d), the scar is visible as a reduction in the space between consecutive isochrones, corresponding to a decrease in propagation speed of the activation wavefront in the border zone. The passive core zone can be seen as a dark area in the middle of the scar, corresponding to an infinitely late activation. Panels a) and b) show the area of the right ventricle where the wavefronts from the two stimulation points meet, resulting in sharp corners on the 40ms isochrones.

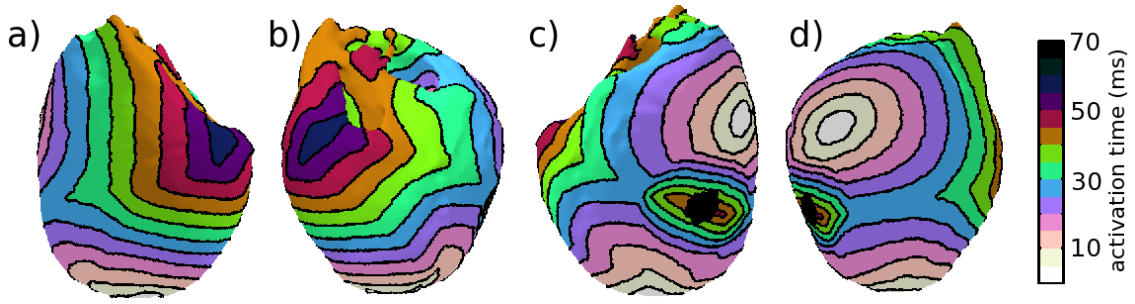


Figure 4.2: Example of epicardial activation isochrones for the ventricular mesh, obtained from a simulation incorporating the transmural scar with activation initiated at points 1 and 3 (see Fig. 4.1). Every isochrone corresponds to a 5 ms interval. The mesh rotates 270 degrees clockwise from panel a) to d).

4.3.2 Connectivity vs Convergence in a Cuboid Geometry

Figure 4.3 depicts activation isochrones on the surface of an unstructured cuboid geometry (3x3x3 arbitrary space units), calculated using Dijkstra's algorithm. Conduction velocities were 3, 2 and 1 (arbitrary space/time units) in x, y and z directions respectively, and the point of initial activation (t=0) was placed in a single node in the corner (0,0,0). As described in Section 4.2.2, the connectivity for each of the 8000 nodes in the cube is determined using the parameter r. The parameter defines a radius for each node, and every node within that radius is connected to the central node. Thus, by increasing r, the connectivity of each node increases. Fig. 4.3 also depicts the true, analytical isochrones (calculated using Eq. (4.3)) as thinner lines, towards which the computed isochrones converge.

As can be seen from Fig. 4.3, the analytical and computed isochrones become increasingly similar with increasing connectivity, while keeping the resolution fixed. Also visible is that the isochrone similarity is greater in directions where the isochrone spacing is denser, which corresponds to a smaller conduction velocity. However, as the connectivity is increased, corresponding to higher values of r, the isochrone similarity rapidly increases and the RMS errors decrease. At the highest r value, the isochrones are virtually indistinguishable, and the RMS error has decreased with 2 orders of magnitude compared to the lowest connectivity.

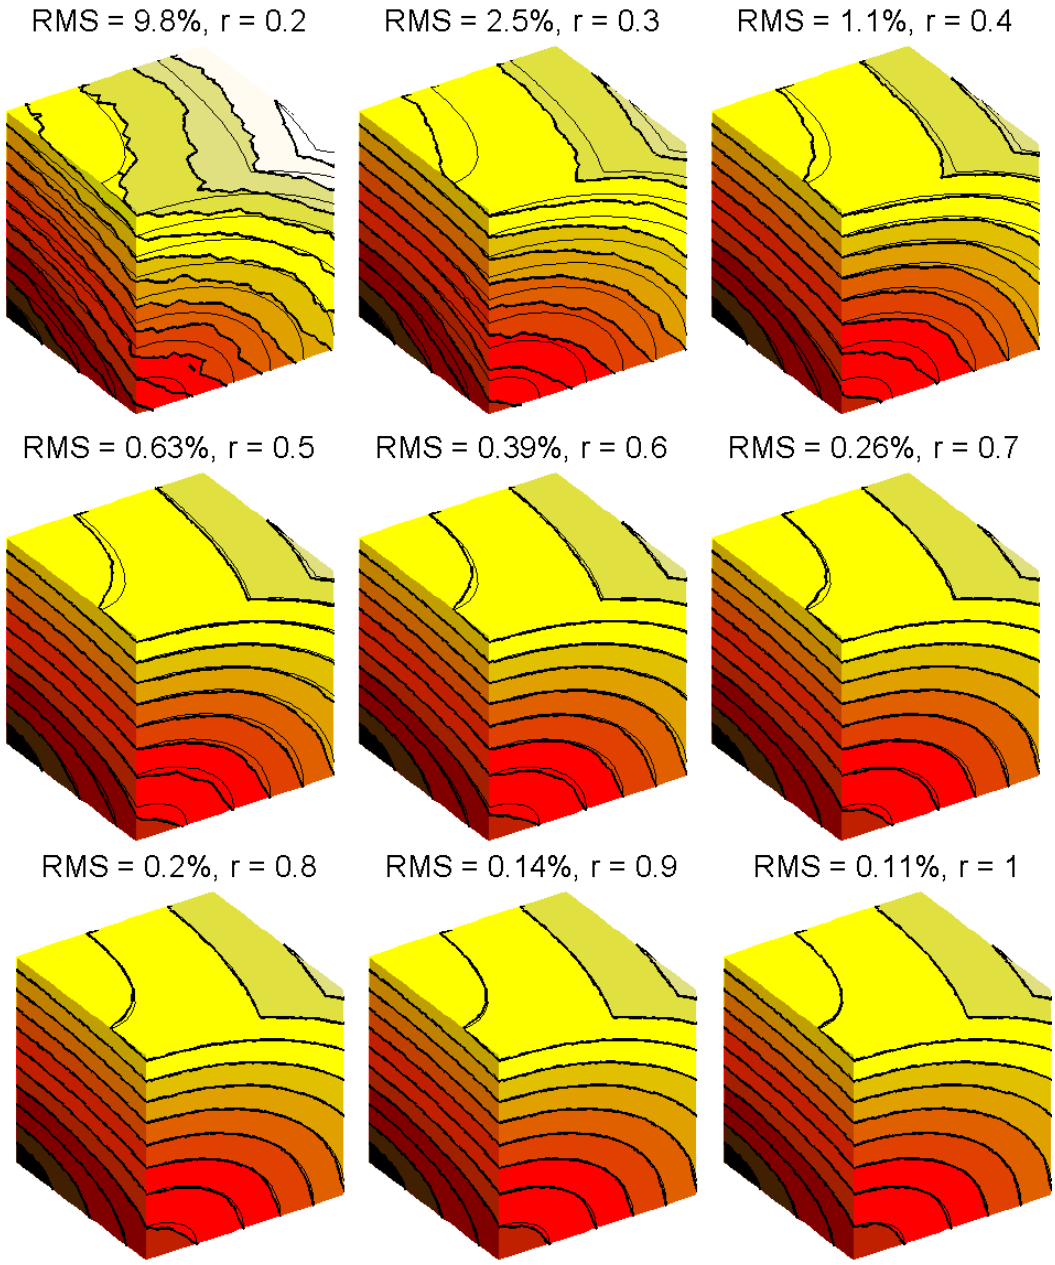


Figure 4.3: Activation isochrones for a 8000 node cuboid geometry (3x3x3 arbitrary space units) under successive increase of the connectivity. Isochrones in thin lines show the correct activation time isochrones, as computed according to Eq. (4.3), while isochrones in thick lines are calculated using Dijkstra's algorithm. Propagation was initialized from a single node in the lower left corner. The parameter r corresponds to the radius within which all nodes are connected to a given, central node. Conduction velocities are 3, 2 and 1 arbitrary space/time units in x, y and z directions, respectively. Also shown are the root-mean-square (RMS) errors in activation time, given as a percentage of the correct total activation time.

Next, as described in Section 4.2.2, the corresponding reaction-diffusion problem was studied by solving the monodomain equation on the cuboid geometry, using the Mitchell-Schaeffer model for the reaction part of the equation. Conductivites in the x,y and z directions were set proportional to 9,4 and 1 respectively.

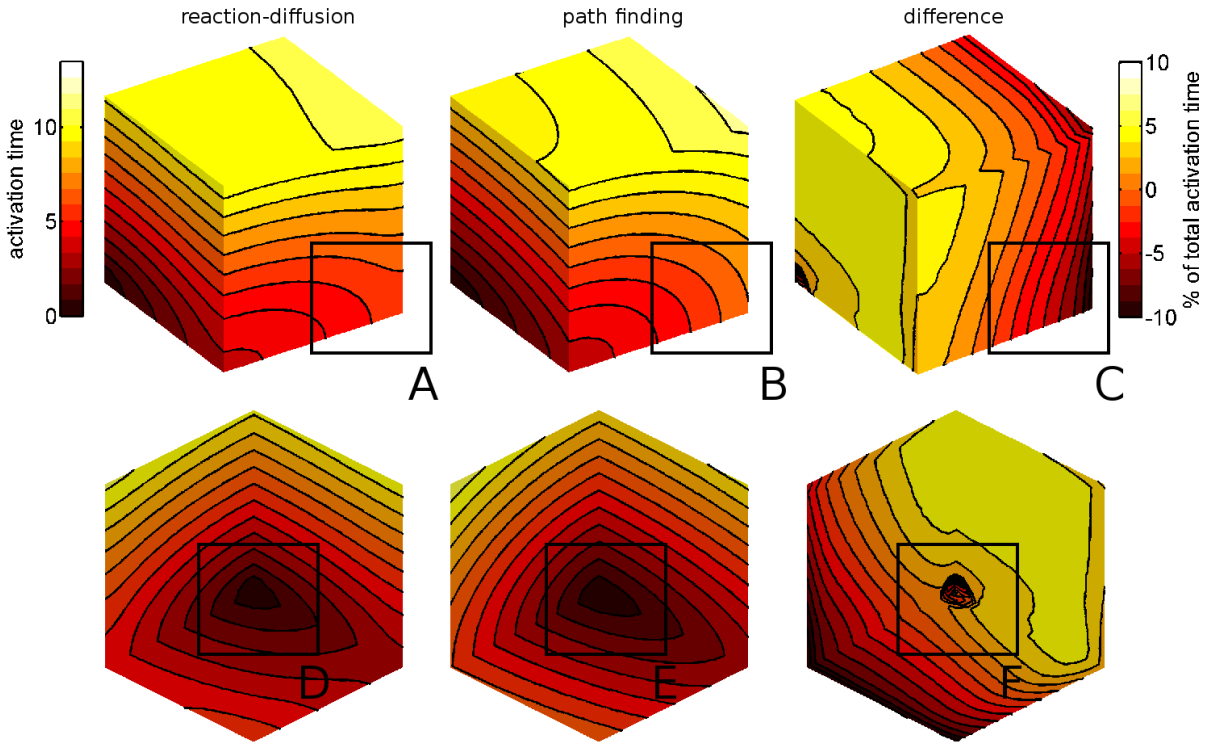


Figure 4.4: Activation isochrones obtained from the reaction-diffusion system (left column) and from the highest connectivity solution (r=1 in Fig. 4.3) of the path finding model (middle column). Also shown is differences between path finding and reaction-diffusion activation times (right column). Two different views are shown, corresponding to upper and lower rows. Activation times, relating to the left and middle columns, are in arbitrary time units. Differences in activation times, relating to the right column, are expressed as percentages of the latest activation time of the isochrones from the path-finding method. Rectangles marked A-F mark regions of interest, which are further explained in the text.

The differences between the path-finding and reaction-diffusion isochrones, as visible in Fig. 4.4 stem from two major factors: the activation of the tissue and the wavefront behaviour near boundaries. For the reaction-diffusion system, no-flux boundary conditions lead to faster propagation near the boundaries of the computational domain, making the isochrones perpendicular to the boundary, while no such effects are present in the case of the path finding method, where

the boundary is simply defined by having a conduction velocity equal to 0 outside. Comparing the isochrones in Fig. 4.4, these boundary effects are most clearly visible in the lower right corners of the upper row of cubes, marked A-C in the figure. Region A clearly shows that the isochrones from the reaction-diffusion system are perpendicular to the domain boundary, whereas the isochrones from the path-finding solution (region B) are not. This leads to the large negative differences in activation time visible in region C. Regions D-F show the site of initial activation. The isochrones from the path finding method (region E) are evenly spaced, indicating a uniform propagation speed. Conversely, the corresponding isochrones for the reaction-diffusion model (region D) are denser close to the point of initial activation, indicating an initially slower propagation due to the explicit treatment of tissue excitability incorporated in the monodomain equation. This difference is also clearly visible on the difference map shown in the rightmost column (region F).

Figure 4.5 shows the relationship between errors, CPU time and connectivity for the cuboid geometry. The left panel of Figure shows the CPU time as a function of node connectivity, while the right panel shows the RMS error (in $\%$ of the total activation time) as a function of the CPU time (ms).

As shown in the left panel of Fig. 4.5, there is a strong linear relationship between the connectivity and the CPU time needed to compute the arrival times. This is not surprising, when considering the structure of Algorithms 1 and 4; the number of nodes are the same, but the computation time devoted to expanding each node increases linearly with the number of neighbours, giving rise to the linear relationship. The right panel of Fig. 4.5 shows that the gain in accuracy with increasing connectivity is initially faster than exponential, making increasing accuracy a fairly inexpensive investment in terms of CPU time.

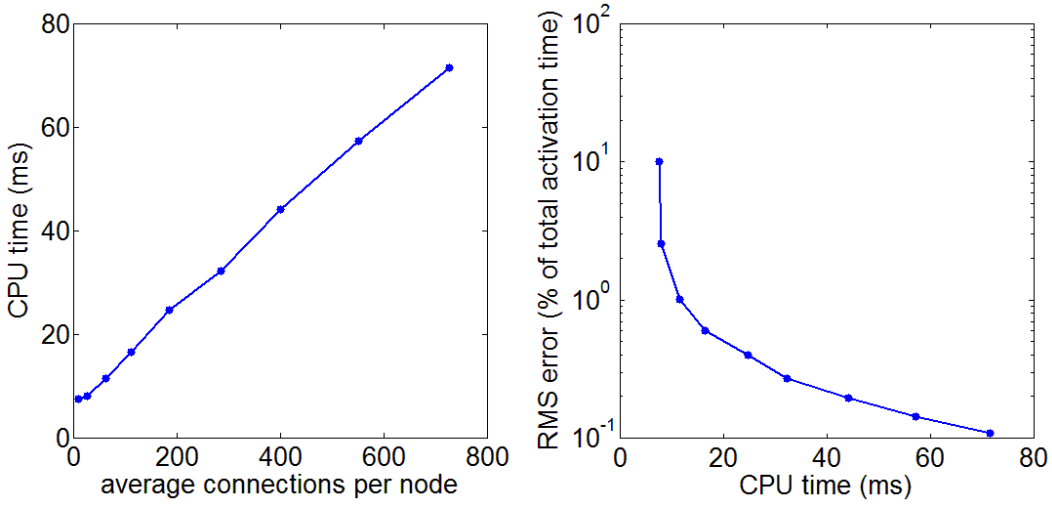


Figure 4.5: Left panel shows the CPU time as a function of node connectivity for the 8000 nodes cuboid geometry. Right panel shows the RMS error (as % of the total activation time) as a function of the CPU time (ms).

4.3.3 Comparison with Independently Estimated Cost Functions

Next, using the monodomain activation times from the previous section as the gold standard, a cross validation procedure was employed as described in Section 4.2.6. The aim was to study the robustness against variations in stimulation point of the graph based method and the FMM, as well as to compare the magnitude of the errors between the two methods. The resulting errors for both methods are presented in Table 4.1. Columns in Table 4.1 correspond to varying stimulation points, whereas rows correspond to varying conduction costs.

As can been seen from Table 4.1, the RMS errors from the respective cross validation of the two methods have a number of similar features. Quantitatively, no individual error exceeds 3.01 ms and the majority are below 2 ms for both methods. The diagonal elements represent the lowest error in each row for both methods. This is to be expected, since the diagonal corresponds to the situation where the same point has been used for both estimating the conduction cost and finding the RMS error. Additionally, it can be observed that the highest and lowest errors correspond to the same respective positions in both tables. It should also be noted that the highest and lowest

cross validation errors of the graph based method / FMM (ms)						
		cost				
		1	2	3	4	5
p	1	1.58/1.47	2.45/2.23	1.81/2.29	1.82/1.79	1.80/2.00
o	2	2.46/2.05	1.43/0.97	3.01/2.70	1.69/1.11	1.93/1.01
i	3	0.97/1.93	1.61/2.00	0.88/0.86	1.43/1.66	1.45/1.80
n	4	1.79/1.69	1.58/1.19	2.04/2.01	1.41/1.07	1.56/1.10
t	5	1.65/2.07	1.97/0.99	1.95/2.16	1.59/0.99	1.47/0.97

Table 4.1: RMS errors for cross validation of the graph based method (left value in columns) / FMM (right value in columns), using five different stimulation points for estimating conduction costs and calculating errors on the 80000 node mesh. Columns correspond to a fixed conduction cost over different stimulation points, whereas rows correspond to a fixed stimulation point over different conduction costs.

errors actually correspond to the same training point (point 3), and that the lowest value is on the diagonal, indicating an over-fitting of the estimated conduction costs for that particular point of stimulation. In spite of this, the remaining results show fairly consistent errors both along columns and rows, indicating that there is no great difference in robustness between the two methods under the conditions investigated.

4.3.4 Mesh Resolution and Quality

Using the estimated conduction costs found in the previous section for the 5 points of stimulation illustrated in Fig. 4.1, 5 relative RMS errors were calculated. This was done for each of the 8 different mesh resolutions, rendering a total of 40 errors.

Fig. 4.6 shows the relative error in activation time from the graph based method with respect to corresponding results from the monodomain as a function of the number of nodes in the mesh. As can be seen from the figure, the graph based algorithm has a generally higher RMS error than the FMM. However, both methods exhibit minor fluctuations in error over the different resolutions. Notably, for the highest resolution, one of the FMM errors is significantly higher than the others and one of the graph based errors is significantly lower that the others. These correspond to stimulation points 1 and 3 of Fig. 4.1, respectively. In spite of the fluctuations,

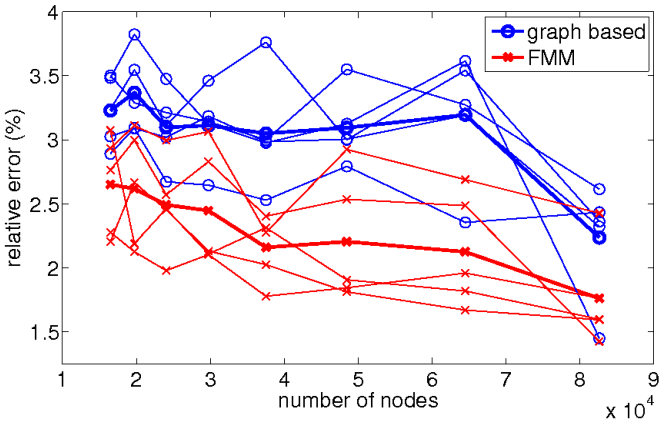


Figure 4.6: RMS errors in activation time for five different stimulation points over 8 mesh resolutions. The upper set of curves (blue line, circular makers) shows results for the graph based method, while the lower set of curves (red line, X-shaped markers) shows corresponding results for the FMM. The two thick lines correspond to average RMS errors, while the thin lines correspond to individual stimulation points.

the general trend in the error is decreasing for both methods with increasing mesh resolution, and all errors of the graph based method are below 4%.

As noted in Chapter 3, the quality of numerical solutions to the bidomain and monodomain equations is highly dependent on the mesh resolution. Conversely, the quality of numerical solutions to the eikonal equation obtained from the FMM are more sensitive to the tetrahedral angles of the mesh. In order to study the quality of the generated meshes from this point of view, the distributions of tetrahedral angles of the meshes were calculated. The resulting histograms are shown in Fig. 4.7.

Fig. 4.7 shows that the distributions of angles are very similar between figures, with every histogram exhibiting two major peaks at around 45 and 90 degrees respectively. This seems to be a consequence of the fact that the algorithm used by the Tarantula meshing software to generate the meshes largely relies on dividing voxels (cubes) into tetrahedra. The general trend is that these two peaks, along with two lower peaks at around 36 and 56 degrees, tend to get more pronounced when increasing resolution. However, the rough distribution and range does not change drastically between resolutions, and would thus not be expected to have a great

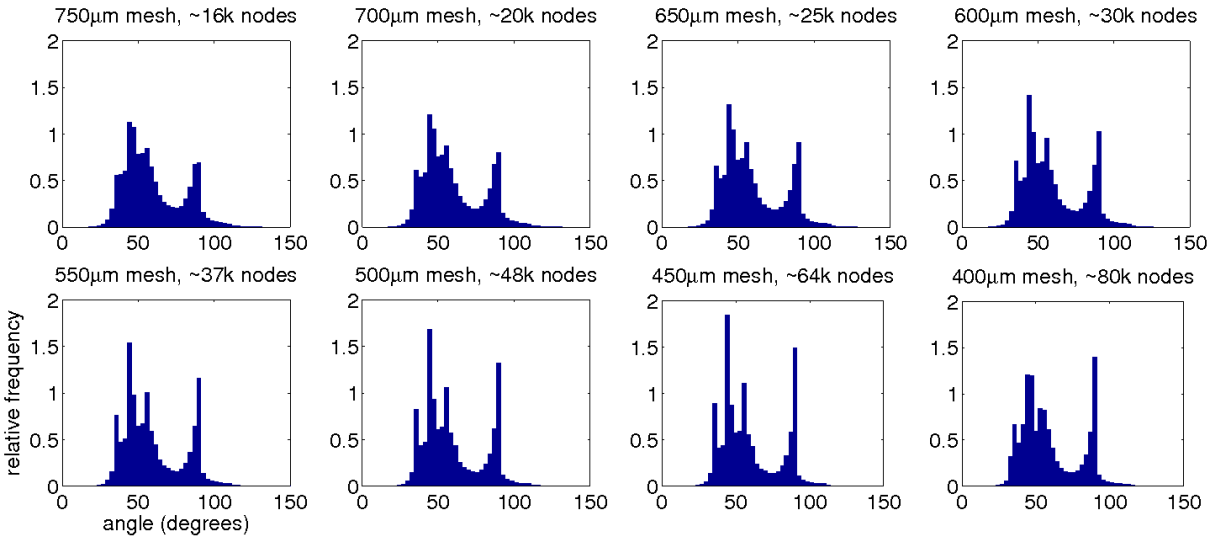


Figure 4.7: Histograms of tetrahedral angles in the 8 meshes. The resolution of the meshes increases from left to right and from top to bottom as indicated in the figure.

impact on the error rate variability between resolutions shown in Fig. 4.6.

4.3.5 Comparison with Known Cost Functions

Next, a new set of activation times obtained from monodomain simulations was compared to corresponding activation times obtained using the graph based method and the FMM. Here, tissue activation was initiated in points 1 and 3 simultaneously, roughly mimicking the typical lead positions of biventricular pacing [170]. This was done both in the presence and absence of the transmural scar, as illustrated in Fig. 4.1. However, this time the conduction costs were not independently estimated for each of the methods, but instead obtained using the conductivity tensors used for the monodomain simulations.

Specifically, the wave propagation speeds (i.e. the inverses of the conduction costs) for the three principal directions were taken to be $c_0[d_l^{0.5}d_t^{0.5}d_n^{0.5}]$, where d_i are the coupling coefficients used in the monodomain simulations, and c_0 is a constant that depends on the action potential upstroke velocity. c_0 was estimated using minimization of the RMS error as described in Section

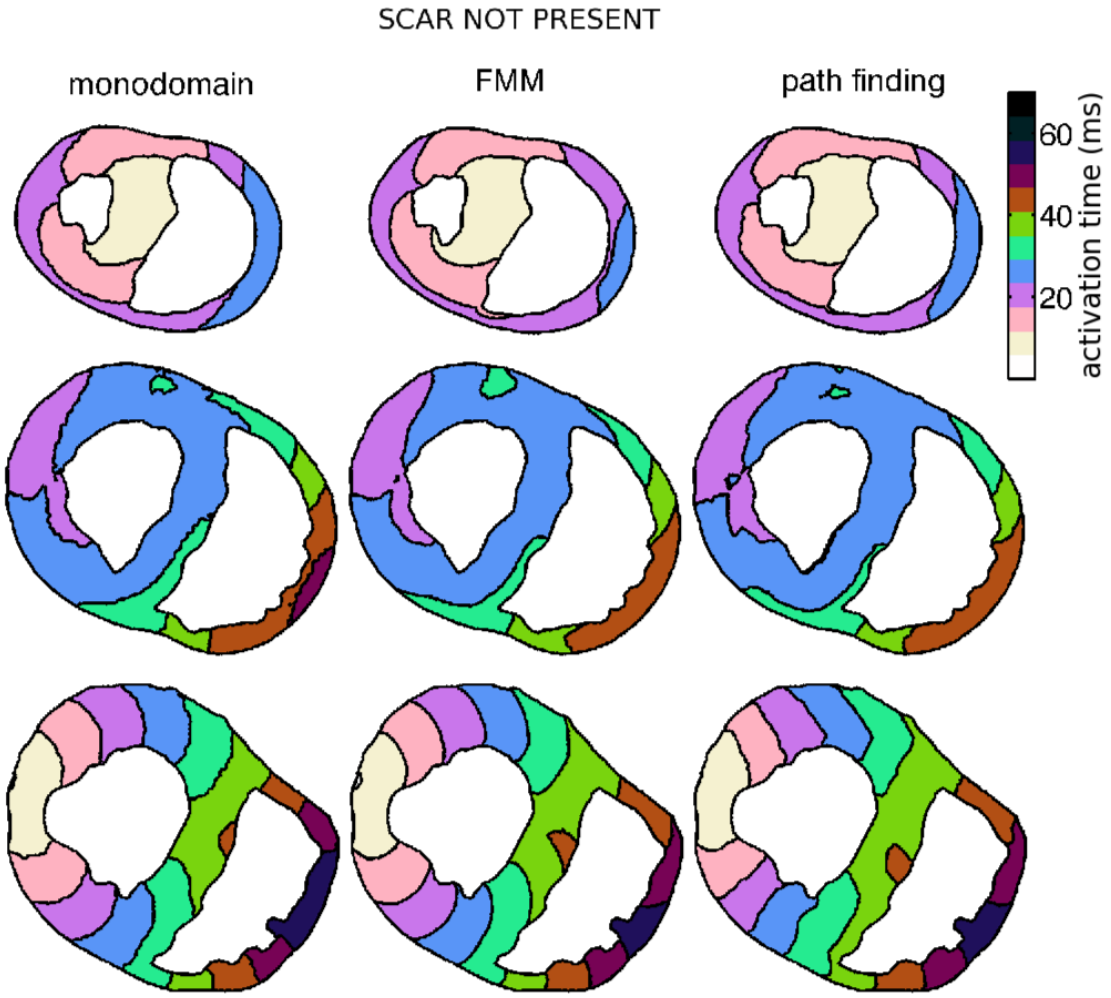


Figure 4.8: Activation isochrones for cross sections of the mesh at the z planes 0.5, 1.1, and 1.6 cm. The columns correspond to the three different methods, and the rows correspond to the three different cross sections. Values on the color bar are in ms.

4.2.5. Results are shown in Figs. 4.8, 4.9, 4.10, 4.11, 4.12 and 4.13.

Figs. 4.8 and 4.9 shows activation isochrones obtained using monodomain, FMM and the graph based method on three different cross sections of the ventricular mesh. Fig. 4.8 shows activation times obtained without the transmural scar, while Fig. 4.9 shows the activation times obtained in the presence of the transmural scar. Figs. 4.10 and 4.11 shows the corresponding epicardial activation isochrones for the same sets of activation times and for three different rotations, without and with the scar respectively. Corresponding epicardial error maps for the same three

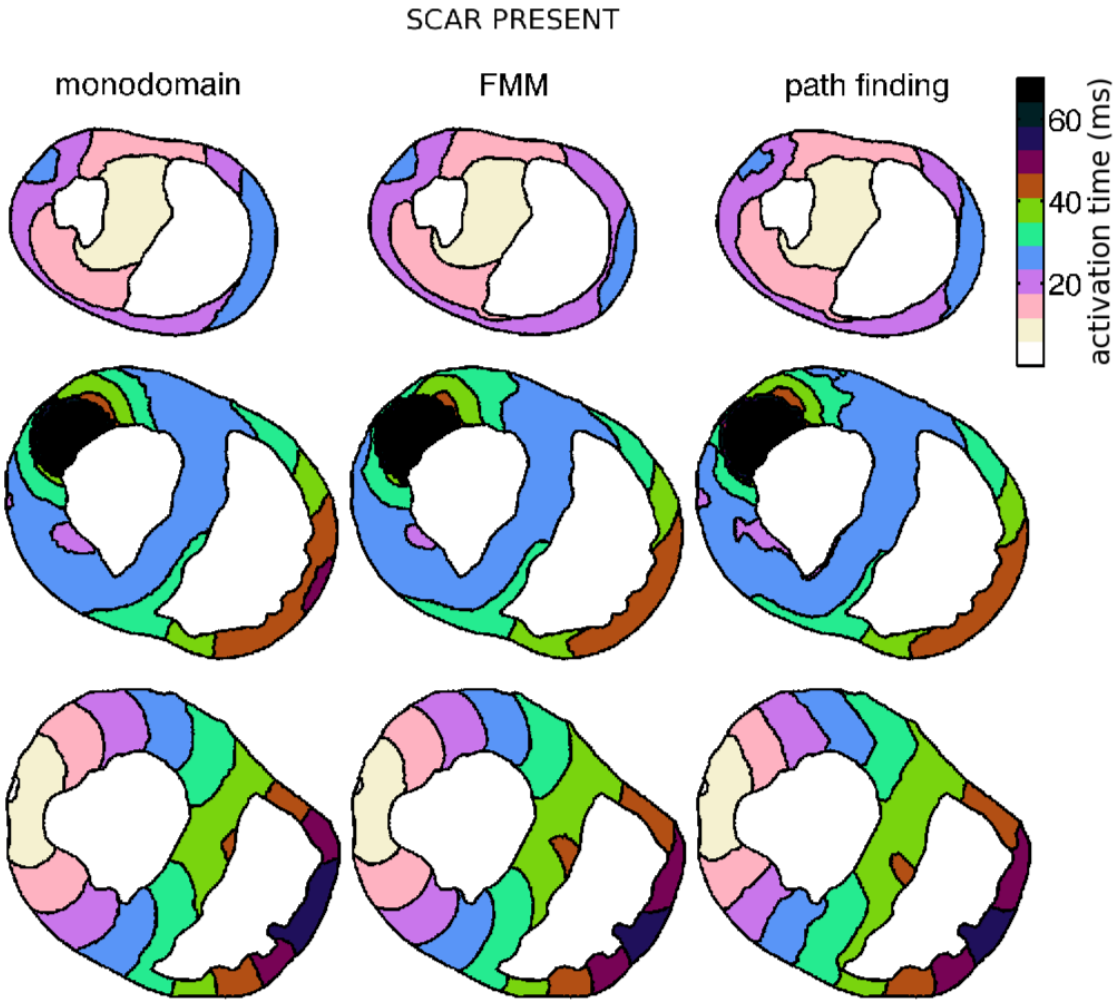


Figure 4.9: Activation isochrones for cross sections of the mesh at the z planes 0.5, 1.1, and 1.6 cm, in the presence of the transmural scar. The columns correspond to the three different methods, and the rows correspond to the three different cross sections. Values on the colour bar are in ms.

rotations are shown in Fig. 4.12. Results are shown for simulations both in the presence and absence of a transmural scar, as indicated by the labels in the figures.

As can be seen from the isochrones, the three methods give very similar activation times. The area where the difference is most pronounced is in the right ventricular free wall, where the wave fronts from the two stimulation points 1 and 3 (see Fig. 4.1) collide. The wave front originating from point 3 appears to arrive slightly early compared to corresponding results from the monodomain simulations, while the wave front initiated at point 1 arrives slightly late. This

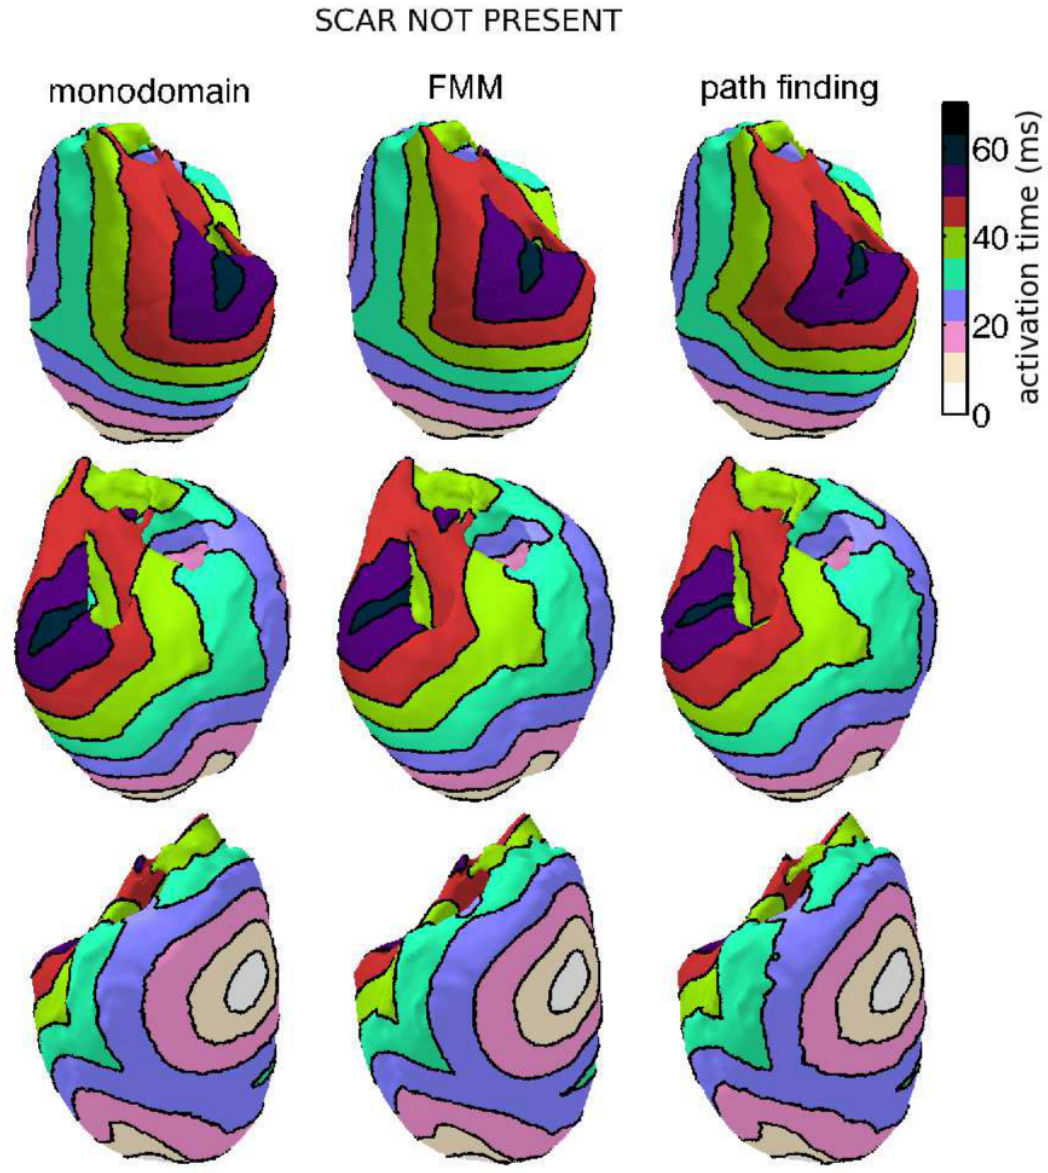


Figure 4.10: Activation isochrones for the epicardium of the mesh. The columns correspond to the three different methods, and the rows correspond left right ventricular, anterior and left ventricular views. Values on the colour bar are in ms.

difference is clearly visible on the two upper rows of Fig. 4.12 as one darker and one lighter region with a sharp border between them. This also is visible in the bottom row of Figs. 4.8 and 4.9, where the collision of the wave fronts, corresponding to the latest activation in that region, is slightly displaced towards the posterior part of the right ventricular free wall. Additionally, the FMM error maps in Fig.4.12 show a region of early activation within the inner parts of

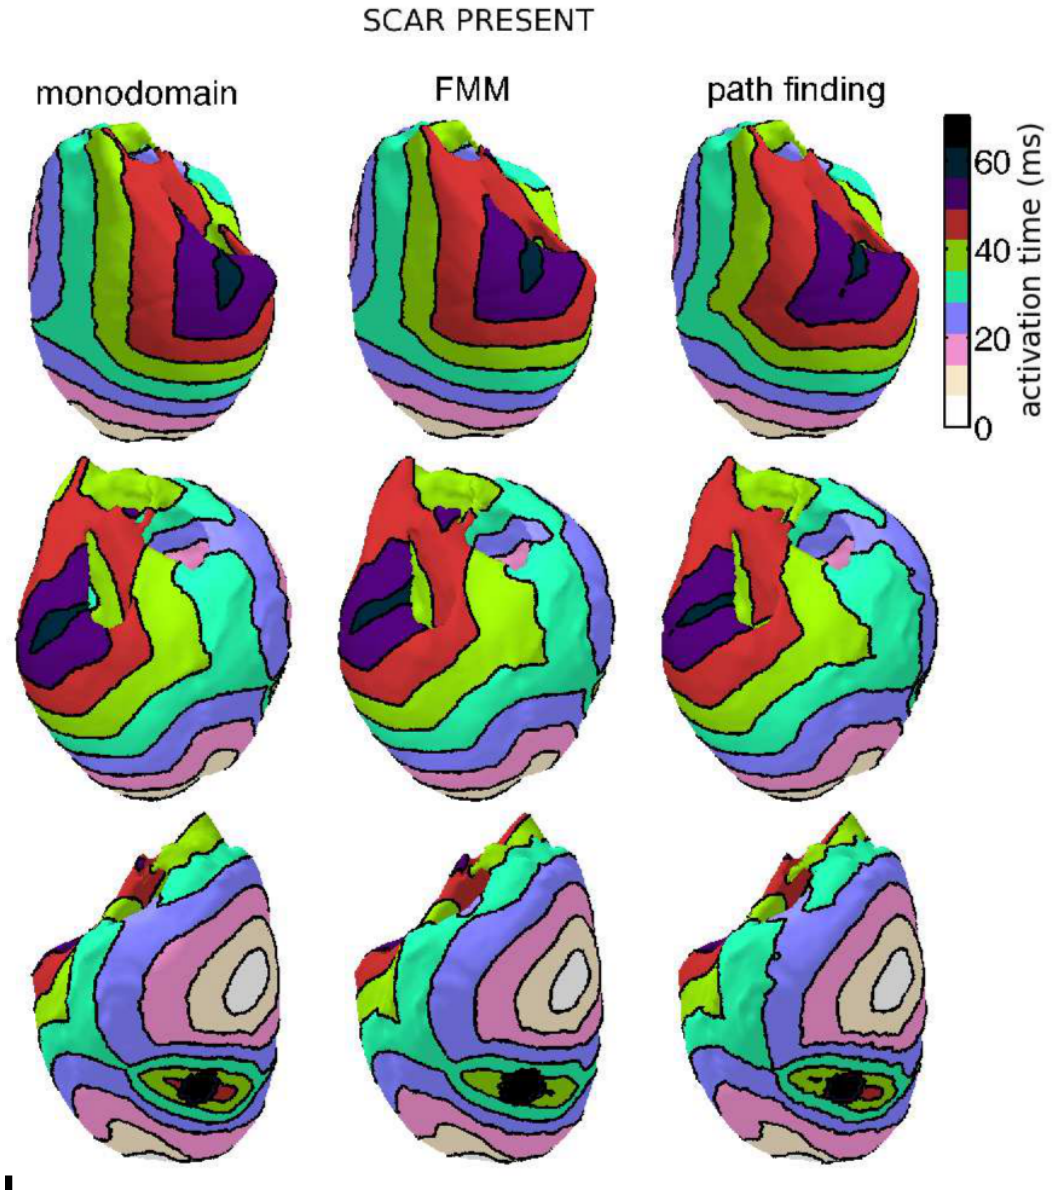


Figure 4.11: Activation isochrones for the epicardium of the mesh in the presence of the transmural scar. The columns correspond to the three different methods, and the rows correspond left right ventricular, anterior and left ventricular views. Values on the colour bar are in ms.

border zone of the scar, possibly indicating that the propagation speed in the border zone is too high. Conversely, the graph based method exhibits no clear pattern for that region.

The histograms in Fig. 4.13 show the distribution of errors for the FMM and for the path finding method, using a range of values for the δ parameter, spanning from 0.4 to 1.2 in steps of 0.1.

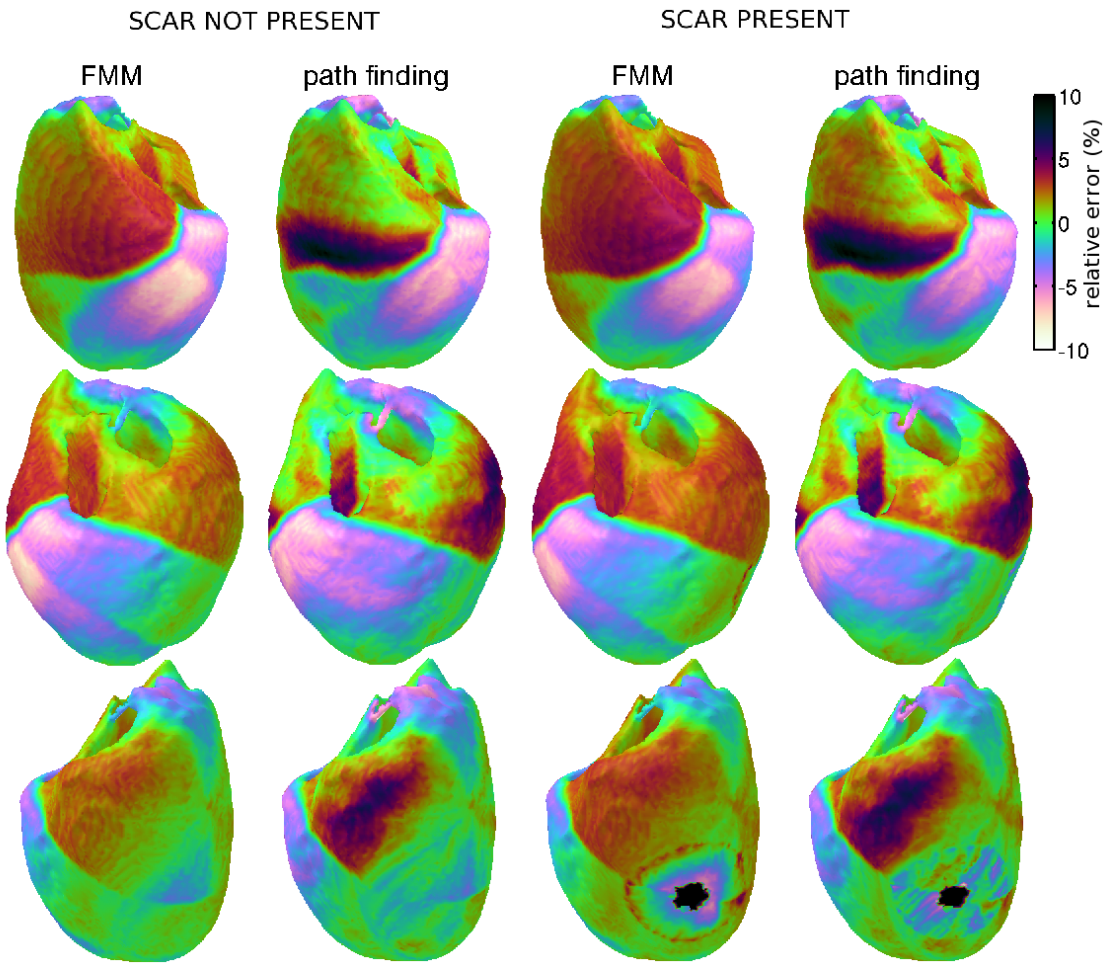


Figure 4.12: Epicardial maps of the relative error. The columns correspond to the FMM and graph based, and the rows correspond to left right ventricular, anterior and left ventricular views. Values on the colour bar are in $\%$ of the total activation time.

The histograms exhibit a clear effect from the choice of the value of δ on both the bias and the variance of the error, with values deviating from the optimum leading to higher errors as well as broader distributions. This indicates that the adverse effects of describing the tissue as a graph can be compensated for using a simple linear scaling of the cost function. A δ value around 0.83 appears to give the best results, suggesting that the average propagation speed between two points in the graph under study is around 83% of the propagation speed of the activation wave obtained from the monodomain equation.

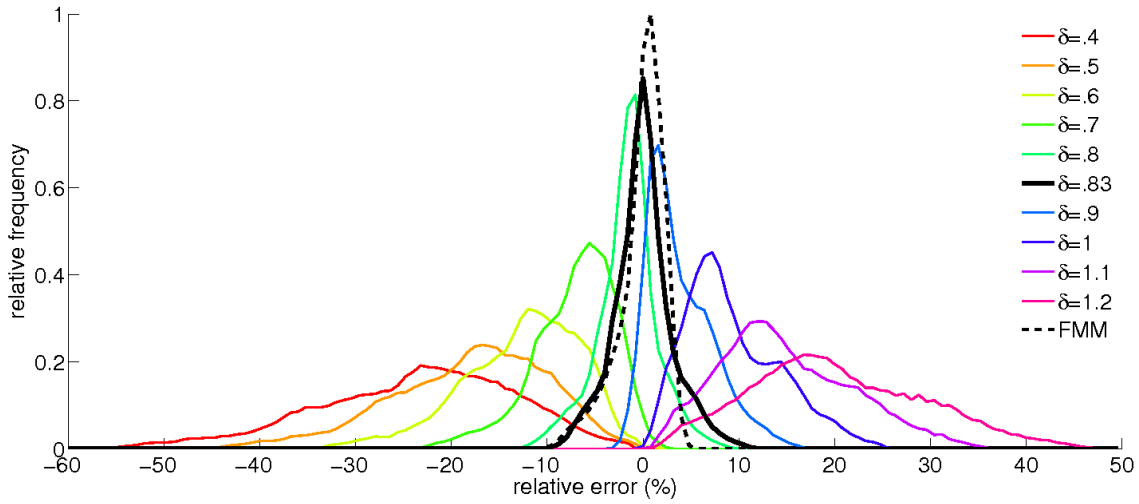


Figure 4.13: Histograms showing the distribution of relative errors in activation time approximations from the FMM, and from the path finding method for a range of values of the δ parameter. Histograms were constructed form 138 bins, equally spaced between -60% and 50% relative error. Coloured lines represent histograms of errors from suboptimal delta values, with high values leading to positive errors and low values to negative errors. The solid black line represents the histogram of errors from the optimal fit of the path finding algorithm ($\delta = 0.83$), and the dashed black line represents the histogram of errors form the FMM.

4.3.6 Fibre Orientation

As an additional step in investigating the ability of the model to adapt to anatomical information, the sensitivity of the predicted activation times to fibre directions was examined and compared to that of the FMM.

Starting from the correct fibre directions, (i.e. those used to perform the monodomain simulations described above), the fibre direction defined in each element of the mesh was rotated around an axis pointing in the transmural direction (i.e. approximately perpendicular to the endocardial surface), leading to an offset in the transmural fibre rotation. This was done in steps of 9° for a total of 20 steps, corresponding to a final rotation of 180°. Activation times were estimated using the graph based method and the FMM for each fibre offset, rendering 20 relative errors for each of the methods relative to the monodomain results, shown in Fig. 4.14.

Fig. 4.14 shows the impact of varying the fibre direction on the predicted activation times for the

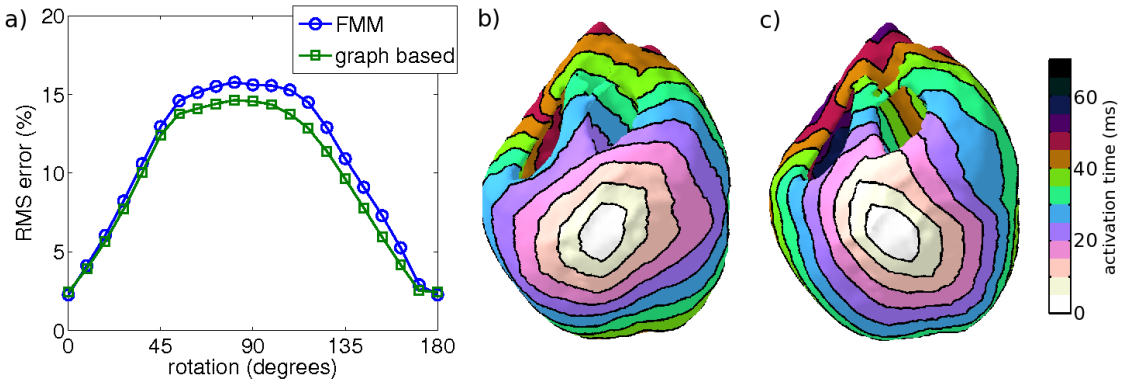


Figure 4.14: Panel a) shows errors resulting from an offset in transmural fibre rotation. Offset ranges from 0° to 180° in steps of 9°. Panels b) and c) show epicardial activation isochrones for the graph based method, with fibre rotation offsets of 0° and 90° respectively, corresponding to the lowest and highest errors in panel a).

FMM and the graph based method. Panel a) shows the RMS error for the FMM and the graph based method as a function of offset in transmural fibre rotation, while panels b) and c) show epicardial activation isochrones obtained using the graph based method. Panel b) corresponds to the lowest RMS error (no offset), while panel c) shows the highest RMS error (90° offset). As can be seen from panel a) of 4.14, the errors for the FMM and the graph based method behave very similarly in response to an offset in fibre rotation. For both methods, the errors monotonically increases from 0° to 90°, after which they monotonically decrease from 90° to 180°. The minima occur at the endpoints, corresponding to alignment with the fibre directions used for the reference monodomain simulations. The graph further shows that the maximum of the error curve for the graph based method is about 1% lower that that for the FMM, indicating that the graph based method is slightly less sensitive to the offset in fibre rotation. The lack of exact symmetry is likely caused by variations in mesh/graph structure.

4.3.7 Computational Performance

As a final step, the computational performance of the graph based method was investigated and compared to that of the FMM. The 80000 node mesh was used, representing a resolution where

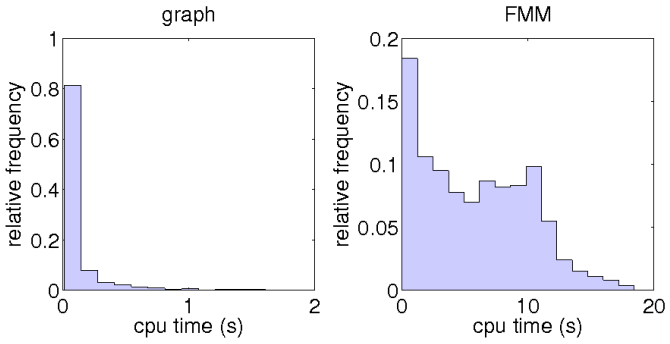


Figure 4.15: Histogram of CPU timings for the graph based method (left panel) and the modified FMM (right panel) when applied to 1000 randomly chosen paths in the realistic rabbit geometry.

sub-millimeter anatomical detail can be incorporated while still being easily manageable on a modern PC. The A* algorithm and the FMM were used to find the paths between 1000 randomly chosen start and end points in the cardiac tissue model described in previous sections.

In order to make the comparison as unbiased as possible, an early stopping criterion was implemented for the FMM, making it terminate as soon as the desired goal node is found. For each path, the computation time was recorded, yielding a set of 1000 CPU timings for each of the methods, from which a histogram was subsequently constructed as shown in Fig. 4.15.

Fig. 4.15 shows histograms for the CPU time taken for estimating the activation time at 1000 randomly chosen nodes in the ventricular mesh, using the graph based method (left panel) and the FMM (right panel). As can be seen from Fig. 4.15, the majority of the investigated CPU timings are well under 0.1 second. In fact, a closer examination reveal that 60% of the times are below 50 ms, 72% are below 100 ms and 92% are below 500 ms. Corresponding times for the FMM are 1% below 50 ms, 3% are below 100 ms and 10% below 500 ms.

It should be noted that these timings represent a relatively naive way of using the graph based algorithm; if the goal or start points are held fixed (corresponding to a fixed point of activation or measurement), which is very likely to be the case in a realistic application of the method, information can be reused, rendering a further decrease in the computation times. The times presented here should thus be viewed strictly as a worst case scenario. Monodomain simulations

were intentionally omitted from this section, since their significantly higher CPU timings would add little to the possible conclusions from this result.

4.4 Discussion

In this chapter, a heuristic, graph based method for approximation of cardiac activation times has been presented. An evaluation of the accuracy and computational performance of the method has been performed, comparing results of the proposed method to results from both the monodomain equation and the FMM.

The results show a consistently high degree of similarity between the activation times obtained using the three methods for all investigated settings. Activation times estimated using the graph based method are typically within 4% of the corresponding activation times obtained using the monodomain, and CPU times on a standard desktop PC are often close to real time when activation times for individual nodes are estimated in an 80000 node mesh.

At the time of writing, to the knowledge of the author, only two previous studies have used path finding to estimate cardiac activation times. Here follows a brief summary of the work presented in them.

Recently, van Dam et al. [171] used a graph based method to estimate activation times on a ventricular surface. Although superficially similar, the method differs from the one presented here in several important ways. While the method of van Dam only considers nodes on the cardiac surface, our method also takes into account nodes that are situated within the cardiac walls, allowing description of for example transmural scars. Additionally, while van Dam et al. use a fully connected graph, our method uses a sparser graph with a comparatively small number of connections for each node, allowing it to be applied to larger graphs, which in turn allows more geometrical detail to be incorporated in the model. Furthermore, our method treats fibre

directions in a more explicit way, in that it uses information about the the principal conduction directions defined at each node when determining the conduction costs for the graph.

Another related approach by Pashaei et al. [172] relies on a mesh-less, closed form solution of the paths through the cardiac tissue. The method is compared to the FMM in a 2D model of ventricular tissue, and is shown to give a better approximation of cardiac activation times when low resolution meshes are used for the FMM, and similar performance for high resolution meshes. While the idea of considering paths in the cardiac tissue is common between Pashaei's method and the method presented here, the underlying methodology is very different between the two approaches. Although an attractive idea, it is not obvious how to extend the mesh-less approach to more complicated geometries and settings, which would indeed be needed to facilitate a fair comparison to the work presented here.

However, in spite of the similarities between the approaches reviewed above and the method presented in this work, it seems that the FMM is the best candidate for comparison. Similarly to the FMM, the graph based method exclusively concerns itself with the activation sequence of the cardiac tissue, and is thus unable to provide information concerning more detailed cellular dynamics. This prevents the application of the method to study of phenomena such as reentry and restitution. Additionally, neither of the methods take into account the effects on the activation sequence of diffusion-related phenomena, such as the wavefront curvature and interaction with other wave fronts or domain boundaries.

However, the proposed method is well suited for conduction parameter estimation and activation time prediction under (pseudo) steady state conditions. This includes estimation of regional conduction parameters as in the Bayesian inference procedure described in Chapters 5 and 6.

In this study, the graph based method has been compared against results from the FMM, as well as monodomain, in several settings. Results from a cross validation procedure were presented in Section 4.3.3. The RMS errors in Table 4.1 indicate that conduction costs which give low

average errors in activation time for one stimulation point tend to give low errors for other stimulation points as well. It can also be seen that the errors are generally low, and that their magnitudes correlate with corresponding errors from the FMM. While Section 4.3.3 does not constitute a full study of parameter estimation under realistic conditions, it does indicate that the proposed method is well suited for predictive applications, given sufficient data for conduction cost estimation. How much data is sufficient will depend on both the number and placement of measurements, as well as on the method of parameter estimation used. This problem will be given further attention in Chapters 5 and 6.

Results for a fixed set of conduction costs, presented in Section 4.3.5, add to the conclusion that the proposed path finding method is able to mimic activation time results from monodomain with a high degree of accuracy. This is demonstrated both in the presence and absence of a transmural scar, as shown in Figs. 4.10, 4.11, 4.8 and 4.9. As shown in Fig. 4.13, the δ parameter in Eq. (4.1) plays an important role when using previously known conduction costs as parameters for the graph based method. However, in situations where the method is used for both parameter estimation and activation time approximation, as in the cross validation study presented in Section 4.3.3, the δ parameter is implicitly incorporated into conduction costs estimated from data, and does not need to be estimated separately. For other situations, the parameter can be estimated using a single run of the FMM. Results in [166] show that for a fully connected graph, the error in time cost between discrete and continuous paths can be made arbitrarily small, while still keeping the number of nodes limited. Thus, by increasing node connectivity and resolution simultaneously, we expect the δ parameter to approach unity. This is also indicated in Section 4.3.3, where Figs. 4.3 and 4.5 clearly demonstrate that increasing connectivity leads to a rapid decrease in RMS error and to an increasing similarity between isochrones.

The sensitivity of the method to variations in fibre orientation was examined by rotating the fibres around a roughly transmural axis, creating an offset in angle from the fibre orientation

used in the monodomain simulations. Results presented in Fig. 4.14 show a strong effect of offsets in fibre orientation on the error of the graph based method. Together with results presented in Section 4.3.5, this suggests that the proposed method is able to react to both positional and directional changes in conduction velocity in a way that mimics the behavior of the monodomain solutions. This is important from a parameter estimation perspective, since it suggests that parameters governing regional and directional changes in conduction velocity will have an impact on the error between measured activation times and activation times approximated using the graph based method. This, in turn, indicates that such parameters could be estimated using the proposed method.

The study of computational performance presented in Section 4.3.7 shows that the CPU times for activation time approximations of single nodes are likely to be of the order of 10-100 milliseconds for the investigated graph resolution, compared to 1-10 s for the FMM. For graphs with coarser resolutions, the CPU time will naturally decrease, due to the smaller number of nodes that need to be expanded in order to find the path between the two nodes of interest. It should also be noted that even relative CPU timings are always implementation dependent, and that the differences in efficiency between the considered algorithms may exhibit slight variations for different programming languages, systems architectures etc.

The main purpose of the study presented in this chapter has been to investigate the accuracy and computational speed of the proposed method, by comparing the resulting activation time estimates to corresponding results from the FMM and monodomain. For this reason, the protocols and results of the monodomain simulations are not primarily intended to correspond to a particular experimental setup or data modality, but to provide a reasonably realistic basis of comparison of the methods. The monodomain simulations were all performed on a high resolution ($\sim$150 μm), MRI derived finite element mesh, assuring an accurate solution of the governing equations in a realistically shaped computational domain. The fibre orientation and scar tissue, although not derived from measured data and thus less realistic, have provided an

added layer of complexity that has increased the chances of revealing any weaknesses in the methods.

4.5 Conclusions

A method for approximation of cardiac activation times, based on path finding in a connected graph, has been presented, along with validation results from several different settings. The validation has been based on comparison between activation times obtained from monodomain simulations, representing a functional model of electrical conduction, and from solutions to the eikonal equation, representing a phenomenological description of electrical conduction. The method has been tested with respect to accuracy of the predicted activation times, robustness against varying conduction properties and stimulation points, and computation time required. It has been shown to be able to accurately mimic activation times obtained from monodomain simulations, with computational costs close to real time for individual nodes. We thus conclude that the graph based method is well suited for use as a part of the Bayesian inference procedure presented in the following chapter.

5

Bayesian Inference of Conduction Properties

After having developed and tested the fast graph based method for cardiac activation time approximations in the previous chapter, in this chapter we incorporate it into a Bayesian inference procedure with the aim to estimate conduction properties from electroanatomical mapping data. The chapter is divided into four main parts. In Section 5.1, additional context is given to the Bayesian inference study presented in this chapter. This is followed by Section 5.2, where we give an account of the theoretical and computational tools utilized in the study. This section is further divided into three subsections. In Section 5.2.1, an image-derived human anatomical model is presented, along with an account of the conduction parameters that will be estimated. In Section 5.2.2, the Bayesian methodology previously described in Chapter 3 is applied to the current parameter estimation problem, including an account of how errors in modelling and measurements are incorporated. Based on this, Section 5.2.3 describes how the data from the in silico *experiments, which serve as a basis for the study, were generated. Next follows Section 5.3, where results are given, quantifying the posterior PDF for a number of different simulated accuracies in modelling and measurements. Finally, in Section 5.4, the study is discussed with respects to its implications, limitations and relation to previous work*

in the area.

As described in Section 2.5, electroanatomical mapping data are associated with various degrees of inaccuracy in the spatial location of the measurement points, in particular after having mapped the data onto a geometry generated using image data from CT or MRI. Naturally, this will translate into a corresponding uncertainty in activation time for any given point in the tissue. Most studies report spatial errors of the order of a few millimeters, but as described in Section 2.5.7, many of these studies are performed in the atria using error measures that have been put into question. It is thus not clear what are the distributions of errors in ventricular electroanatomical maps. For the study presented in this chapter we take a conservative approach in that we compare a range of accuracies for the data used to validate the developed inference framework. The work described in this chapter corresponds to methodology and results presented by the author and coworkers in [18].

5.1 Introduction

It is the goal of this chapter to develop a methodology for Bayesian inference of conduction parameters based on combined structural and electrophysiological data, such as obtained using the electroanatomical mapping systems reviewed in Section 2.5. The combination of structural and electrophysiological information has benefits both as a direct tool for the clinician, for the experimental researcher, and for the growing efforts of the computational modelling community. For the clinician, it can give direct insights into the complex spatial EP features arising during various disease conditions, such as the location of ischemic zones, scars and ectopic foci. This information, in turn, can be used to optimize clinical treatment such as ablation therapy [173] and lead positioning for pacing devices [174]. For the experimentalist, successful analysis can reveal information about quantities that are not directly measurable, such as conduction anisotropy ratios, regional cell orientations and conduction velocity [6, 175]. For the compu-

tational modeller, integration of EP model parameters, such as regional conductivity values, which in turn enables simulation of the electrical activity of the heart, provides a powerful tool for understanding the complex mechanisms involved in many cardiac pathologies and interventions [91, 94, 97, 100, 122, 176]. However, without reliable values for model parameters, there is little guarantee that the model output will bear any resemblance to reality, which severely limits the use of modelling in a clinical context.

In this chapter, we will focus on parameters governing the conduction velocity of the electrical activation wave, which are crucial to estimate accurately in order to construct a personalized cardiac model. As mentioned above, these conduction parameters are estimated from activation time data, typically obtained using electroanatomical mapping systems. The data are always associated with a significant degree of uncertainty, mainly caused by inaccuracy in catheter location [13] and in co-registration with detailed geometry from for example MRI [50]. Because of this, estimated conduction parameter values are also associated with uncertainty. Thus, it is important to be able to quantify the uncertainty in extracted parameters, to prevent estimation errors from being translated to errors in therapy.

The framework of Bayesian inference [14, 15] is well suited to describe this uncertainty. As we have seen in Section 3.2.2, the Bayesian methodology generates a joint probability distribution in the parameter space of the model, rather than estimating single parameter values. This distribution encompasses information about the probability of any possible combination of parameter values. Thus, it provides both the means of selecting the most probable set of parameter values, and gives information about the probability of a given set of parameter values in relation to alternative sets. This addresses the first challenge posed in Section 1.1, that of quantifying the uncertainty in estimated parameters. However, Bayesian inference applied to this type of problem generally requires a large number of simulations. This makes computationally intensive simulation using frameworks such as the monodomain and bidomain models unsuitable. In the study presented in this chapter, we tackle the problem by utilizing the graph based model

developed in Chapter 4 to perform computationally efficient simulations of the cardiac electrical propagation.

The present study concentrates on describing a few significant conduction parameters, in order to focus the study on the impact of the number and location of the electroanatomical measurements on the uncertainty in estimated parameters. By keeping the dimensionality of the parameter space lower, we can perform a detailed study of the impact on each new measurement. This detailed study will subsequently be used in Chapter 6 to tackle the second challenge posed in Section 1.1, that of finding how measurements should be placed in order to optimize the amount of information contained in the measured data. For the present study, simulated measurements are generated using a human ventricular model, incorporating realistic fibre orientation and a conduction heterogeneity corresponding to a ventricular scar. This gives data as spatially located activation times, which are used to probabilistically estimate regional and directional conduction velocities, expressed as dimensionless ratios. Performing the study *in silico* gives us complete knowledge of the correct values of the estimated parameters, as well as control over errors levels in estimation and modelling, which is crucial to validate the presented methodology.

5.2 Methods

In this section we give an overview of the components used to perform the study. We begin by describing the human ventricular model, along with the conduction parameters that will be estimated. Next, we describe the Bayesian inference procedure used to estimate parameters for the conduction model. Finally, we describe how synthetic data were generated for the validation of the inference procedure.

5.2.1 Anatomical Model

The anatomical model used for this study is a human ventricular mesh, derived from CT data [177]. The model includes an algorithmically defined fibre orientation derived from experimental data [93], generated using the CHASTE simulation package [169, 178]. To mimic a clinically relevant setting, the model also includes a manually defined, transmural scar on the left ventricle, consisting of a core zone and a border zone. The scar border zone is defined as all tissue within a given distance of the closest point on the core zone. We denote this distance BZ_{ext}. Conduction velocity outside the scar is assumed to be anisotropic but homogeneous throughout the ventricular tissue. To describe the anisotropy, each point in the tissue is associated with three conduction velocity values, describing conduction speeds longitudinal, transverse and normal to the local myocyte orientation, denoted cv_l, cv_t and cv_n respectively. Conduction inside the border zone is modelled as the outside conduction velocity multiplied by a constant factor, expressed as $cv_l/cv_{l,BZ}$. The scar core zone is modelled as having conduction velocity equal to zero, approximating a dense scar region with very low conductivity [179]. Thus, a set of five effective parameters were defined: cv_l, cv_t, cv_n, $cv_l/cv_{l,BZ}$ and BZ_{ext}. For the purposes of the parameter estimation procedure, cv_t and cv_n were normalized by cv_l, resulting in four parameters to estimate. The geometrical model and the fibre orientation of a cross section of the mesh are shown in Fig. 5.1.

5.2.2 Bayesian Inference

Using the graph based method described in Chapter 4, the Bayesian inference described in Chapter 3 provides a tool to infer cardiac tissue conduction properties based on available measurements of activation times. Using the conduction velocity parameters described above, we now define a point $\mathcal{Q} = \{cv_l/cv_t, cv_l/cv_n, cv_l/cv_{BZ}, BZ_{ext}\}$ to be estimated from a data set consisting of measurements of the activation time at a number of points on the LV endocardium,

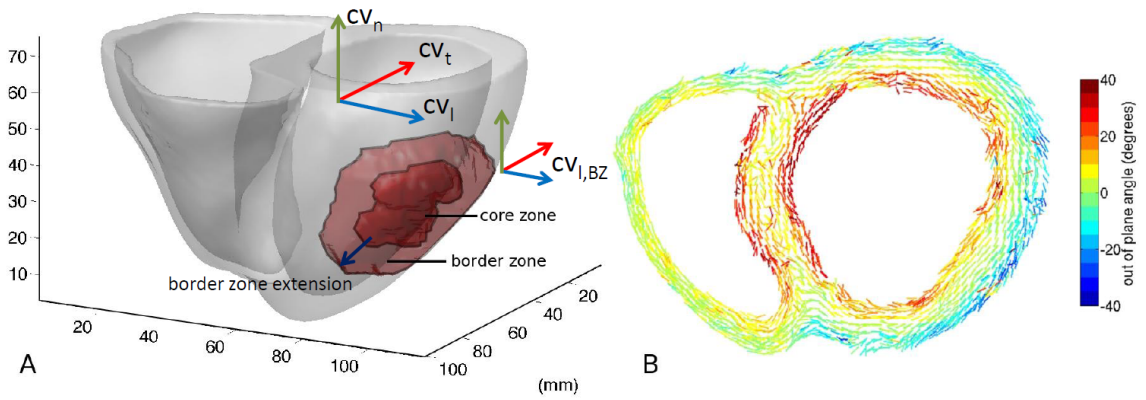


Figure 5.1: Panel A shows the CT-derived ventricular model used in this work. The transmural scar located on the LV anterior wall, is shown in red. The panel also illustrates the quantities used to construct the parameters that are estimated. Panel B shows a cross section of the fibre orientation, approximately half way between apex and base. Lines indicate in-plane fibre directions, and colours indicate out-of-plane fibre directions.

such as obtained using the CARTO system [13]. This data set defines a single point in the data space, and will be denoted $\mathcal{D} = \{d_1, d_2, ..., d_N\}$, with N being the total number of measurements obtained.

Any given choice of values for the four parameters in $\mathcal{Q}$ corresponds to a point in a four dimensional parameter space. As described in Section 3.2, the goal of the inference procedure is to compute the Bayesian posterior PDF in the parameter space of the model, $P(\mathcal{Q}|\mathcal{D}) = P(cv_l/cv_t, cv_l/cv_n, cv_l/cv_{BZ}, BZ_{ext}|d_1, d_2, ..., d_N)$. The posterior quantifies our rational uncertainty about what values the model parameters should take in order to make the simulation output as similar as possible to the activation sequence of the heart from which the measurements were obtained.

As described in Section 3.2, for any point $\mathcal{Q}$ in the parameter space, a set of activation time estimates, defining a point $\hat{\mathcal{D}}$ in data space, can be calculated using the model. Since the graph based propagation model used in this work is deterministic, $\hat{\mathcal{D}}$ is fully determined by $\mathcal{Q}$. For any point $\mathcal{Q}$ there is likely a discrepancy between the resulting model output $\hat{\mathcal{D}}$ and the corresponding measured values $\mathcal{D}$. The discrepancy comes from two sources: noise inherent in the

measurements, and errors inherent in the propagation model. This is to say that there exists a set of underlying, true, activation times for all points in the tissue, from which measured activation times deviate because of imperfections in the measurement techniques employed, and from which the model predictions deviate because of imperfections in the model. Denoting this true, underlying set of activation times in the cardiac tissue $\mathcal{D}^* = \{d_1^*, d_2^*, ..., d_N^*\}$, we describe the measured activation times as $d_i = d_i^* + \epsilon_m$ and the estimated activation times as $\hat{d}_i = d_i^* + \epsilon_{est}$. This allows us to express the total error at a given measurement point as

$$\epsilon = \hat{d}_i - d_i = (d_i^* + \epsilon_{est}) - (d_i^* + \epsilon_m) = \epsilon_{est} - \epsilon_m. \tag{5.1}$$

The PDF for the errors in modelling can be assumed to be fairly constant, since we have no particular reason to believe that contact measurements obtained at some points of the endocardium will be more or less accurate than measurements obtained at some other points. Conversely, the errors in estimation should depend on how far away from the point of initial stimulation the estimates are obtained. Close to the point of initial stimulation, the errors are small, since the activation times are small, but further away from the point of initial stimulation, the errors should grow larger. Here we approximate the behaviour by making the errors in estimation dependent on the underlying, true, activation times d_i^*. With the further assumption that both types of errors have a Gaussian shape, we get $\epsilon_m \sim N(0, \sigma_m)$ and $\epsilon_e \sim N(0, \sigma_{est}(d_i^*))$, and it follows that $\epsilon \sim N(0, \sqrt{\sigma_{est}(d_i^*)^2 + \sigma_m^2})$. Specifically, we model the dependence of σ_{est} on d_i^* as according to

$$\sigma_{est}(d_i^*) = \sigma_0 d_i^* \tag{5.2}$$

where σ_0 is a constant describing the accuracy of the propagation model. Assuming mutual independence between the errors we can rewrite the equation for the posterior (Eq. 3.15) as

$$P(\mathcal{Q}|\mathcal{D}) \propto P(\mathcal{Q}) \prod_{i=1}^{n} e^{-\frac{1}{2}\frac{(d_i-\hat{d}_i)^2}{\sigma_{est}(\hat{d}_i)^2+\sigma_m^2}} \tag{5.3}$$

The assumption about Gaussian error distributions is very common, and is included in most published cardiac parameter fitting procedures. This is done either implicitly, in the shape of a least squares minimization function, or explicitly, included in the Kalman filters and other Bayesian approaches. For the modelling errors, Fig. 4.13 also shows a roughly Gaussian distribution for the errors between the activation times calculated using the monodomain model and corresponding activation times computed using the graph based method presented in the previous chapter, giving further indications that the assumption is reasonable.

5.2.3 Data Generation

As a basis for the study, the graph based method was used to obtain 80000 sets of activation times for all nodes in the tissue model, each corresponding to a different point in parameter space. The points were picked based on a grid discretization of the parameter space. Three dimensions, corresponding to the parameters cv_l/cv_t, cv_l/cv_n, and $cv_l/cv_{l,BZ}$, sampled from 1 to 8 in 20 steps of around 0.37 each, while the fourth, corresponding to the BZ_{ext} parameter was sampled in 10 steps of around 1.5 mm each. This corresponds to a grid of 10x20x20x20 points in parameter space. Subsequently, these 80000 sets of activation times were used to generate synthetic data from measurements and simulations by adding various degrees of noise to them, in a way that ensures that Eq. (5.3) holds. Each of the model evaluations took around 0.1 seconds on a single core of an AMD Athlon II X4 640 3.0 GHz Processor. This allowed the required points to be computed in about 1.5 hours, utilizing the fact that $cv_l/cv_t < cv_l/cv_n$. It should however be noted that an almost linear speed-up can be achieved when utilizing more processors for this task, since the individual model evaluations, which take up the main computational effort, are fully independent. The details of the data generation procedure are described

below.

We designate one of the points in parameter space, which we denote $\mathcal{Q}^* = \{cv_l/cv_t^*, cv_l/cv_n^*, cv_l/cv_{l,BZ}^*, BZ_{ext}^*\}$, to be the one corresponding to $\mathcal{D}^*$. Given $\mathcal{Q}^*$, and the corresponding activation times $\mathcal{D}^*$, we now want to ensure that our assumptions from the previous section about the error distributions are true:

1. $\hat{d}_i - d_i^* \sim N(0, \sigma_{est}(d_i^*))$, and
2. $d_i - d_i^* \sim N(0, \sigma_m)$.

Requirement 1 says that using the parameter set $\mathcal{Q}_\rangle$, the difference between estimated activation times, $\hat{d}_i$, and true underlying activation times d_i^* should be normally distributed with a standard deviation $\sigma_{est}(d_i^*)$. This is achieved by sampling a set of random variates $\mathbf{n}^{est}$ from a normal distribution $N(0, \sigma_{est}(d_i^*)$ and adding them to d_i^*, that is to say

$$\hat{d}_i = d_i^* + n^{est}. \tag{5.4}$$

In a similar manner, we ensure that requirement 2 is fulfilled by adding another set of random variates $\mathbf{n}^m \sim N(0, \sigma_m)$ to the true underlying values, that is to say

$$d_i = d_i^* + n^m. \tag{5.5}$$

Generating values for $\hat{d}_i$ and d_i according to equations 5.4 and 5.5 ensures that Eq. 5.3 is correct, and thereby allows us to study the effects of varying accuracies, pacing sites and measurements sites on the outcome of the inference procedure described in Section 5.2.2.

5.3 Results

Here we present results for the validation of the proposed inference procedure. The following section gives details on the credibility regions of the estimated parameters over a range of noise levels in modelling and measurements.

5.3.1 Validation of Inference under Different Levels of Measurement and Model Accuracy

Measurement data d_i and simulation data $\hat{d}_i$ were generated as described in Section 5.2.3 for four different levels of measurement accuracy, corresponding to σ_m taking the values 1, 2, 4 and 8 ms, and four levels of model accuracy, corresponding to σ_0 taking the values 0.05, 0.1, 0.2 and 0.4. These values were chosen to cover a wide range of possible levels of error. For the errors in measurement (σ_m), reported accuracy in electrode positioning, covered in Section 2.5.4 and 2.5.6, lie in the range of a few millimeters. However, reports from positional errors in co-registration with geometrical models derived from imaging, covered in 2.5.7, vary from around 2 mm to around 18 mm in average positional errors. To get an estimate of the time scales involved, we note that at a propagation rate of 100 cm/s, the electrical activation front would need 18 ms to travel 18 mm, and for 2 mm the corresponding value would be 2 ms. Judging from this, the investigated range of error levels seems relevant, although the standard gaussian shape of the associated distributions should be seen as a first approximation. Regarding the errors in modelling, this will likely be just as dependent on the parametrization of the model as on the mathematical formulation used, and is currently not well investigated. Thus, our aim was to cover a broad range.

Subsequently, the inference procedure detailed in Section 5.2.2 was used to obtain a 4D posterior PDF $P(\mathcal{Q}|\mathcal{D})$, describing our uncertainty about the value of the four conduction parameters

$cv_l/cv_{l,BZ}$, cv_l/cv_t, cv_l/cv_n and the BZ_{ext} as described in Section 5.2.1. These posteriors were obtained by evaluating Eq. (5.3) at the 80000 discrete grid points in parameter space, described in Section 5.2.3. In order to visualize the behaviour of the posterior PDF under the addition of subsequent measurements, the marginal distributions for each parameter were calculated. Next, the 95%, 85% and 50% confidence regions for each of the marginal distributions was calculated. This procedure was repeated after the incorporation of each of 100 measurements, for each level of model and measurement accuracy.

Fig. 5.2 shows the evolution of the confidence regions of the marginal distributions for each parameter under variation of the model and measurement accuracy, as subsequent measurements are added. Also shown is the true value for each of the four parameters, marked by a dashed line and an arrow on the right side of each panel. Here, adding a measurement corresponds to incorporating a new term in the right hand side product of Eq. (5.3) as described in the Section 5.2.2. In this case, the measurement points were chosen randomly.

As can be seen from Fig. 5.2, some features are shared by all panels. The confidence regions are typically broad for low numbers of measurements, reflecting a high level of uncertainty about the true parameter value. As more measurements are included, the confidence regions exhibit a clear tendency to become narrower, indicating that the underlying PDF is concentrating its probability mass to a smaller and smaller region, or equivalently that our uncertainty about the true parameter values has descreased. However, while this behaviour represents the general trend, some panels (for example cv_l/cv_{BZ} on the top row, cv_l/cv_{BZ} and BZ_{ext} on the two bottom rows) do exhibit occasional increases in confidence region width. This is to be expected in situations where conflicting information is obtained, such as in the case of a particularly inaccurate measurement. In the present study, this corresponds to drawing a large value for one (or both) of the random variates in Eqs. (5.4) and (5.5).

Further, examining the behaviour of the cv_l/cv_n and $cv_l/cv_{l,BZ}$ parameters, we see very minor reductions in the width of the credibility regions for high levels of modelling uncertainty (bot-

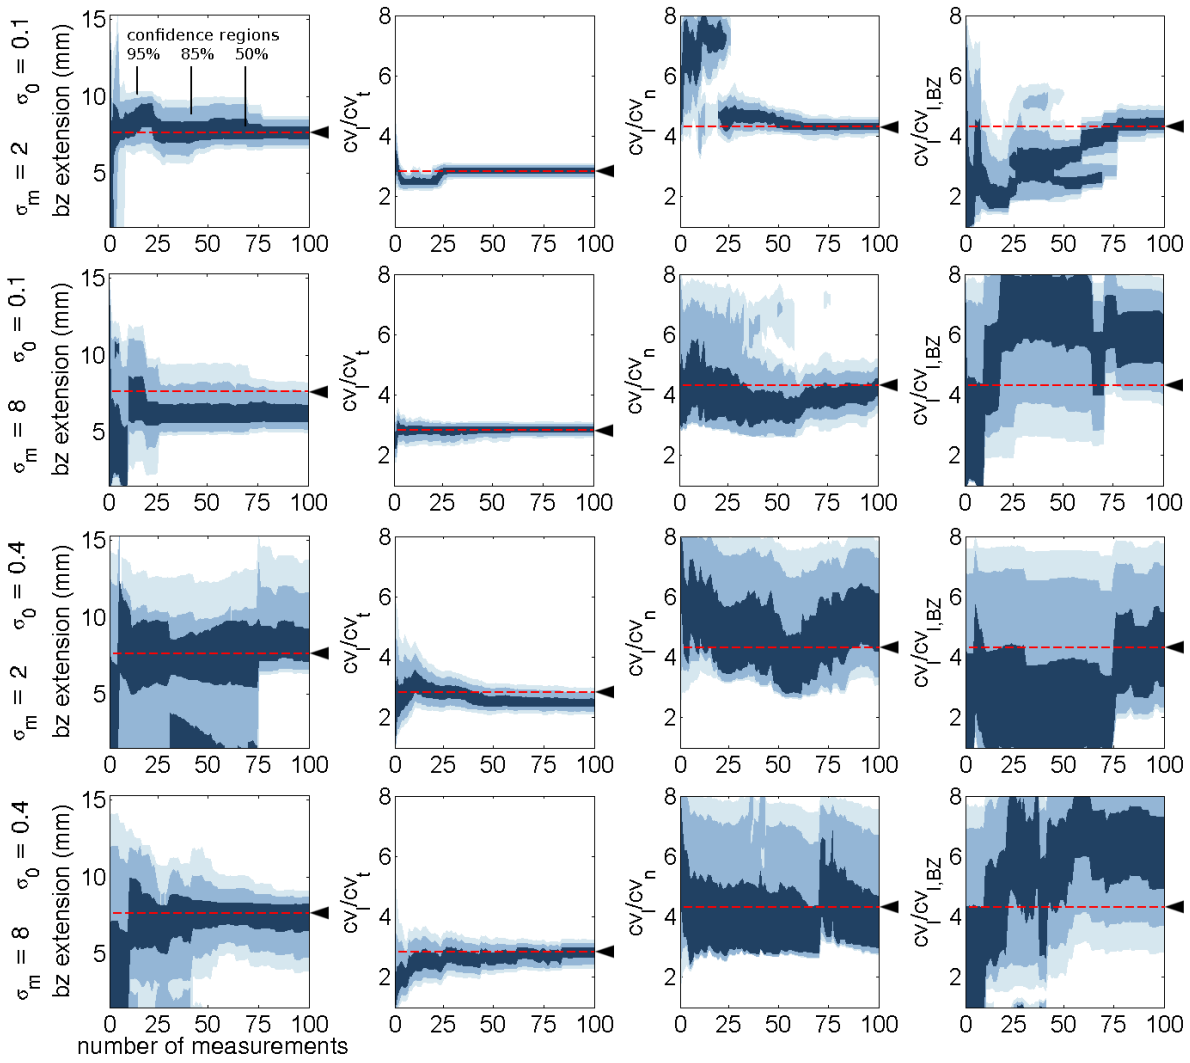


Figure 5.2: Figure shows the evolution of marginal distributions of the four conduction parameters $cv_l/cv_{l,BZ}$, cv_l/cv_t, cv_l/cv_n and $BZ\ extension$ as the number of measurements increases. As indicated in the figure, panels on the same row correspond to the same accuracy, and panels in the same column correspond to the same parameter. The colored fields correspond to the 95%, 85% and 50% credibility regions respectively, as indicated in the upper left panel. True underlying values of the estimated parameters are marked by a dashed red line and an arrow on each panel.

tom two rows), especially at the 50% level. We believe this is due to two main reasons. Firstly, these two parameters have the smallest impact on activation times, and are thus the ones that are most strongly affected by errors. Secondly, especially in the case of $cv_l/cv_{l,BZ}$, relatively few points on the endocardium contain useful information about the value of the parameters. The data in Fig. 5.2 were obtained using a randomized placement of measurements. Since this strat-

egy makes no special effort to optimize the information content of the measurement locations, locations with high information content might be missed, and the reduction in uncertainty can be expected to be slow for high levels of noise. Both of these reasons will be further elaborated in the next Chapter, particularly in the context of Fig. 6.6.

An interesting feature is the occurrence of disjoint confidence regions, corresponding to multimodal marginal distributions. In the examples shown in Fig. 5.2, disjoint confidence regions are exhibited at some set of measurements for 8 of the 16 examples shown in the figure. Only the marginal distributions corresponding to the cv_l/cv_t parameter appears completely free of multimodality, due to their rapid convergence towards the underlying value of the parameter. Running the experiment a number of times with different instantiations of the random numbers that are used to emulate inaccuracies in modelling and measurements, we have observed confidence regions separated into two or even three disjoint parts (data not shown).

Looking at the confidence regions of the different parameters, it is visible that the estimates of cv_l/cv_{BZ} and the BZ extension (columns 1 and 4 of Fig. 5.2) tend to give the most uncertain estimates, or equivalently the widest marginal distributions. As will be further elaborated in the next section, a likely reason for this is that the area of the endocardium where information about these parameters can be obtained is much smaller than the corresponding area for cv_l/cv_t and cv_l/cv_n (columns 2 and 3 of of Fig. 5.2). It is expected that only measurements that are upstream of or inside the border zone will have a significant impact on the estimates of the cv_l/cv_{BZ} and BZ extension parameters. Considering the confidence regions of the cv_l/cv_t parameter, it can be seen that the marginal distribution of the parameter reaches a high accuracy after 20 measurements, after which the confidence regions stay fairly constant.

Comparing different levels of accuracy from the top to bottom panels of Fig. 5.2, it is clear that higher values of σ_{est} and σ_m lead to broader confidence regions. However, the estimation of different parameters exhibits different robustness against these effects. The cv_l/cv_t parameter is almost unaffected at the lowest accuracy, corresponding to the bottom row of Fig. 5.2,

whereas the remaining parameters exhibit significantly broader confidence regions for the lower accuracies.

To further test the validity of the CRs, the marginalizations shown in Fig. 3.4 were calculated 100 times using different sets of random variates. This was done for all combinations of model and measurement errors as described at the beginning of this section. Out of these 100 trials, the number of times that the true parameter value actually fell within the respective CRs (95%, 85% and 50%) was counted. The aim of this procedure was to ensure that the CRs actually captured the true value the appropriate fraction of times. The results of this experiment are presented in Tab. 5.1.

		$\sigma_m = 1$			$\sigma_m = 2$			$\sigma_m = 4$			$\sigma_m = 8$		
	CR	95%	85%	50%	95%	85%	50%	95%	85%	50%	95%	85%	50%
cv_l/cv_t	$\sigma_0 = 0.05$	98	98	97	98	98	94	97	87	78	93	85	56
	$\sigma_0 = 0.1$	96	90	79	94	85	69	95	85	57	97	90	51
	$\sigma_0 = 0.2$	95	86	64	99	87	69	98	92	56	98	91	49
	$\sigma_0 = 0.4$	98	91	58	99	92	56	97	89	52	98	92	45
cv_l/cv_n	$\sigma_0 = 0.05$	100	100	100	100	100	100	100	100	100	97	94	88
	$\sigma_0 = 0.1$	100	100	99	100	97	96	99	98	95	97	90	83
	$\sigma_0 = 0.2$	97	89	83	94	87	72	97	92	74	99	91	77
	$\sigma_0 = 0.4$	98	93	62	98	86	59	96	88	66	93	86	59
$cv_l/cv_{l,BZ}$	$\sigma_0 = 0.05$	100	89	75	99	82	57	97	84	56	95	88	55
	$\sigma_0 = 0.1$	96	90	56	93	84	60	97	84	45	92	82	57
	$\sigma_0 = 0.2$	94	85	45	90	81	47	97	84	54	98	81	46
	$\sigma_0 = 0.4$	98	89	62	100	85	53	97	89	56	99	93	58
BZ_{ext}	$\sigma_0 = 0.05$	94	82	51	92	83	53	93	83	52	89	80	47
	$\sigma_0 = 0.1$	94	86	49	96	83	43	90	73	35	97	83	45
	$\sigma_0 = 0.2$	96	79	32	98	85	45	98	85	41	97	88	34
	$\sigma_0 = 0.4$	100	97	36	100	97	24	100	95	28	100	98	31

Table 5.1: Validation data for the confidence regions of the cv_l/cv_t, cv_l/cv_n, $cv_l/cv_{l,BZ}$, and BZ_{ext} parameters for the 16 combinations of values for σ_0 and σ_m as described at the start of this section. Values represent the number of times the true parameter value was contained within each CR out of a hundred trials. The values in the CR row represents corresponding expected percentages after a very large number of trials. Marked in green and red are anomalous values that will be further discussed in the text.

As can be seen from Tab. 5.1, the true parameter values are contained within the CRs close to or in excess of the expected percentages. Since only 100 trials were made, we do not expect an exact match between the confidence level and the number of times the true value lies within the respective CR. However, some anomalies are visible. Some of the values for low noise levels are significantly higher than should be expected from the CRs. The cv_l/cv_n parameter shows a 100% hit rate at the lowest level of model error (marked in green) for all but the highest levels of measurement noise. However, from studying Fig. 3.4, we see that a likely explanation for this is the limited discretization of our posteriors. The posterior is as narrow as it can get after 75 incorporated measurement points, and the CRs have likely all collapsed to their resolution limit. Conversely, the BZ_{ext} parameter shows some values (marked in red) where the number of times the CRs contain the true value is less than we would expect. This seems to occur for very high levels of model errors, and only for the 50% CR. The reason for this is unclear, although some of the discrepancy might simply be due to random fluctuation caused by the relatively low number of samples used to compute the values.

5.4 Discussion

In this chapter, a probabilistic method for electrophysiological parameter estimation from electroanatomical mapping data has been presented. An *in silico* validation study conducted using an image-derived anatomical model shows that the proposed method is capable of accurately estimating the most likely parameter values as well as giving a quantitative measure of their associated uncertainty. Results presented in Fig. 5.2 indicate that the Bayesian inference procedure presented here produces reliable parameter estimates and confidence regions, given a correct model of the inaccuracies in measurements and modelling. Subsequently, the method has been used to study how the accuracy of measurements and modelling affects the uncertainty in estimated parameters. By studying the behaviour of the Bayesian posterior PDF when incor-

porating successive measurements for various degrees of simulated model and measurement errors, the impact of the inaccuracies has been quantified. Building on this approach, a strategy for the improved placement of endocardial point measurements will be presented in the next chapter.

As previously mentioned in Section 2.7, Konukoglu and coworkers recently used Bayesian inference to estimate a conductivity function for the LV endocardium in [125]. To the knowledge of the author, this is the only published study that gives a full Bayesian treatment to the estimation of cardiac conduction properties from electroanatomical data. Thus, we begin this section with further discussion and comparison of the work presented by Konukoglu *et al.* in relation to the study presented in the current chapter.

To describe current propagation, the authors used the eikonal-diffusion equation, presented in Section 3.1.4. Using this as a model for electrical propagation, a polynomial chaos representation of the system was constructed. The polynomial chaos formalism works by describing the relation between uncertain parameters values and system output (in this case activation times) using a weighted sum of orthogonal polynomials. However, the analysis required to find the proper weights can become very computationally expensive for large numbers of parameters. To speed up this computation, and thus allow for more considered parameters, the authors employed a compressed sensing approach. Once computed, the authors used the polynomial chaos expansion of their model to infer a 20 dimensional posterior using 30000 samples obtained using an efficient implementation of MCMC [180]. Konukoglu *et al.* report results showing a good match between estimated and underlying parameter values for the simulated data, and prediction RMS errors of the order of 10% of the total activation time for clinical data.

The basic principle of the estimation method presented in [125] and the one presented in this study is the same, in that they both ultimately rely on a mathematical model of the cardiac electrical propagation, coupled to a method for point-wise evaluation of the Bayesian posterior PDF. A further similarity is that both rely on a pre-computation step, based on a fixed parametrization.

In the case of the previous study, this step consists of the polynomial chaos expansion, while in the present study, it consists of the grid of 80000 evaluations of the model. A major difference between the methods lies in the parametrization of the conduction model. While Konukoglu et al. use a flexibly parameterized representation of the LV endocardium to represent intrinsic activation through the purkinje system, our findings focus on a setting where purkinje activation is blocked by pacing of the tissue, as often done in electrophysiological studies. Also, while Konukoglu's parametrization allows for scar tissue in any part of the endocardium, but nowhere in the remaining tissue, we assume the scar position to be approximately known from imaging modalities such as late enhancement MR [10, 181].

It is well established that MCMC methods provide a means to evaluate characteristics of high dimensional distributions with comparatively low computational requirements. However, for the present work, an important goal has been to study the behaviour of the posterior PDF under a great number of conditions, both in terms of different levels of errors in modelling and measurement, and in terms of measurement placement. For this purpose, we believe that the grid based evaluation employed here is better suited than an MCMC approach. This is both due to the fact that makes it easy to obtain a detailed description of the posterior, and because it allows the posterior to be sequentially evaluated (as described in Section 5.2.2) for each new measurement. In principle, this could be achieved also with MCMC methods, but we expect the computational costs to be significantly larger.

In the work presented here, the results from the simulation study indicate that the developed method gives accurate credibility regions for parameter values, when provided with a correct description of the noise inherent in measurements and modelling. This can be observed in Fig. 5.2 by noticing how the true value is almost always contained within the 95% confidence region, seemingly independent of the actual level of noise. This conclusion is further strengthened by Tab. 5.1, showing that the true parameter values seem to be contained within their respective CR the correct fraction of times. Thus, for an appropriately selected parametrization of the model,

we believe that the method has great potential to be useful in any setting where conduction parameters need to be estimated from a small amount of measured data.

Figure 5.2 also illustrates the common occurrence of disjoint confidence regions, corresponding to multimodal posterior distributions. This further highlights the need for probabilistic treatment of the parameter estimation problem considered here. In the presence of several local maxima of the posterior distribution, a naive approach such as optimizing the square difference between measurements and model predictions is not guaranteed to find the global maximum. Additionally, even when employing MCMC-based approaches to sample the posterior distribution, summary statistics such as mean and standard deviation are likely to give deceptive results when the posterior is multimodal.

The study presented here has been conducted completely *in silico* for validation purposes. An *in silico* study gives access to the underlying, true values of the estimated parameters, which is crucial when validating the method. However, it is also possible that the setting presented here is less complex than an *in vivo*, clinical setting. The conductive properties of the tissue have been modelled as anisotropic but homogeneous, apart from the inhomogeneity represented by the scar region. Specifically, the author is aware of the complex functional [31, 32] and structural [29] abnormalities associated with infarcted cardiac tissue, and a further refinement of the corresponding parts of the current model could constitute an interesting future development, as will be further discussed in Chapter 7.

Further, the method for Bayesian inference used in this work gives relative rather than absolute values for the posterior PDF, since the evidence term of Eq. (5.3) is not explicitly computed. This makes the method used here unsuitable for model comparison procedures, which could otherwise be used to optimize the parametrization of the model [182]. As described in Section 5.2.2, the errors resulting from inaccuracies in modelling and measurement techniques have been approximated as normally distributed and unbiased. While this is a fairly common approach for measurements, the errors in modelling electrical conduction are largely unknown,

and will furthermore differ between different models and parameterizations. Thus, the description presented here should be viewed as a first approximation, useful for studying the behaviour of the inference procedure in a setting where the error distribution is known. In principle, a more accurate description of the errors in modelling could be derived from repeated comparisons between computer models and large sets of clinically obtained data, but this is outside the limits of the current study.

6

Application of Bayesian Inference to Optimal Measurement Placement

In the previous chapter, we saw how Bayesian inference can be utilized to produce probabilistic estimates of parameter values, naturally incorporating information about the associated uncertainty. The inference procedure was made computationally efficient using the graph based conduction model developed in Chapter 4. In this chapter, we build on the these results to address the challenge of how to place electroanatomical point measurements on the endocardial surface in a way that optimizes their information content. Like previous chapters, this chapter is divided into four main sections. In Section 6.1, we give a brief context to the work presented in subsequent sections. In Section 6.2 we give a theoretical analysis of the problem of optimizing the placement of measurements (in Section 6.2.1), as well as present a practically useful measure for selecting the measurement locations (in Section 6.2.3). This is followed by Section 6.3, where results from different strategies for measurement placement are compared. Finally, in Section 6.4, these results are discussed and conclusions are drawn.

6.1 Introduction

It is the goal of this chapter to devise a strategy for placing electroanatomical measurements on the cardiac surface in order to optimize the amount of information that can be extracted from the collected data. Optimizing the amount of information that can be extracted from a given number of measurements has the potential to reduce the time and cost of clinical invasive procedures, and to provide a better means of optimizing treatment strategies. Here it is worth noting that the advent of electroanatomical mapping systems has in itself had a great impact on the amount of information that can be extracted during an EP study, and has been shown to reduce risk factors, such as exposure to ionizing radiation [36].

To address the challenge of how to optimize the way electrophysiological measurements are placed on the cardiac surface, we require a measure with which to compare distinct estimates of the same parameters. By looking at the process of parameter estimation from a probabilistic point of view, such as that developed in Chapter 5, such a measure can be obtained, namely the uncertainty associated with the estimated parameter values. In this chapter, we build on the results and methods presented so far in this dissertation to devise a strategy for measurement placement. The work described in this chapter corresponds to methodology and results presented by the author and coworkers in [18].

6.2 Methods

In this section we present the methods and theory on which the rest of the chapter is built. The goal is to pose the problem of optimal measurement placement in terms of the probabilistic description of parameter values developed in previous chapters, in a form that allows for mathematical treatment.

6.2.1 A Combinatorial View of Measurement Placement

As discussed in Chapters 3 and 5, different choices of measurement points correspond to different subspaces of the data space $\mathcal{D}$ of our model. Strictly speaking, the data space should be considered infinite-dimensional, since the endocardial surface is continuous. However, for both practical and computational reasons, we will approximate this space as finite dimensional. From a practical standpoint, the reason is that the data from two measurement points sufficiently close together will be indistinguishable, while computationally we are dealing with a geometrical model that is discretized into a finite number of points. The approximation of the data space as finite-dimensional is explicitly or implicitly used in practically all contemporary work on parameter estimation from EP measurements [183].

Thus, we consider the endocardial surface to consist of N separate regions r_i, each corresponding to a dimension of $\mathcal{D}$. Each of these regions can be thought of as having two states, visited or unvisited, depending on whether or not a measurement has been placed in the region. This, in turn, allows us to enumerate 2^N different states for the tissue as a whole, corresponding to every conceivable combination of states (visited/unvisited) for the regions. Below, we will refer to the states of individual regions as *region states* and the states for the tissue as a whole as *tissue states*.

Viewed in this way, adding successive measurements corresponds to moving from one tissue state to another. However, not all tissue states are accessible from any other tissue state. Firstly, the movement between tissue states is constrained by placing one measurement at a time, meaning that two successive tissue states can only differ by one region changing state. Secondly, moving between tissue states is constrained to a region state going from unvisited to visited, and not the other way around. This defines a directed acyclic graph (DAG), with 2^N nodes corresponding to tissue states, and $2^{N-1}N$ edges corresponding to allowed changes in region states.

Additionally, each node in the DAG is associated with a cost. In order to understand how these costs are obtained, we consider the structure of Eq. (5.3). The right hand side of the equation consists of a product, where each of the terms (except the prior $P(\mathcal{Q})$) corresponds to the effect that a single measurement has on the posterior PDF. Evaluating the right hand side can be done iteratively, by starting from the prior $P(\mathcal{Q})$, multiplying it by the first term of the likelihood $P(d_1|\mathcal{Q})$, multiplying the result by the second term of the likelihood $P(d_2|\mathcal{Q})$ and so forth, until all terms are included. Each new term will modify the shape of the resulting, intermediate, posterior PDF in some way, typically by making it narrower, i.e. concentrating its probability mass to a smaller region of parameter space. But, going back to the graph description of measurement placement, each term also corresponds to moving from one tissue state to another in the DAG by making a region state change from unvisited to visited. Thus, each possible movement within the DAG is connected to a change in the shape of the posterior PDF. This lets us associate each node of the DAG with a cost of visiting it. For the purpose of this work, these costs are taken to be the minimum number of points on the discretized, intermediate posterior PDF, associated with a particular node, that contains 95% of the probability mass. A schematic illustration, showing a simplified domain made up of three regions, is shown in Fig. 6.1.

Following this analysis, we realize that the optimal sequence and placement of electroanatomical measurements reduces to a path finding problem in a DAG, such as that described above and shown in Fig. 6.1. Specifically, with the given formulation of costs, the aim is to find the least expensive path to each node in the DAG, starting from the initial tissue state (no regions visited). Here, the cost of a given path is found by adding up the costs of all the nodes it contains. A solution to this problem gives information about the sequence of measurements leading to the fastest decrease in uncertainty for any desired level of final uncertainty and number of measurements. However, although the problem of finding optimal paths in a DAG is well known and solvable in linear time [184], the very large number of nodes and edges make the computational requirements prohibitively large.

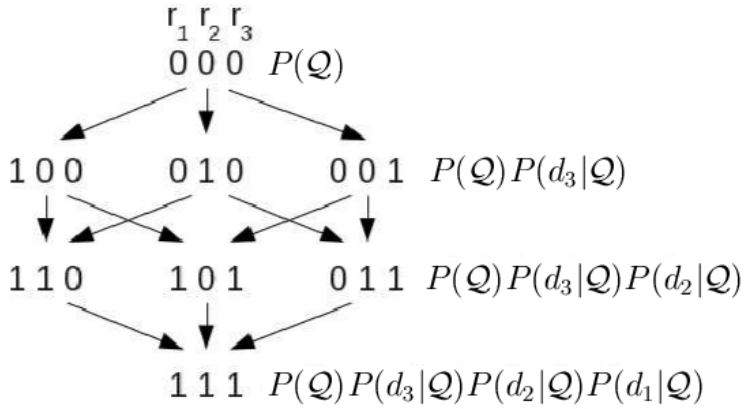


Figure 6.1: The connected graph corresponding to measurement placement in a three region domain. r_1, r_2 and r_3 are the three regions, each having one of the two possible regions states: unvisited (0) or visited(1). Each group of region states represents a specific tissue state. Arrows indicate allowed movements within the graph. Also shown are four intermediate posteriors, corresponding to iterative evaluation of the right hand side of Eq. 5.3 for the measurement sequence $\{r_3, r_2, r_1\}$.

6.2.2 An Approximate Greedy Solution for Measurement Placement

To obtain an approximate solution to the path finding problem, we employ a greedy algorithm. The aim of this is to provide a gold standard for the rate of decrease in width of the posterior PDF of Eq. 5.3, against which other strategies for measurement placement can be compared. The greedy algorithm works by always selecting as next measurement the region leading to the greatest decrease in uncertainty of the intermediate posterior PDF of Eq. 5.3. This is repeated until all regions have been visited. However, because the greedy algorithm requires knowledge about the results of measurements in all accessible regions of the cardiac tissue, it is not practically useful for predicting how to optimally place measurements. Thus, we need to develop a new heuristic, that does not rely on prior knowledge of the outcome of the measurements.

6.2.3 A Heuristic for Measurement Placement

Considering Eq. 5.3, we see that the interface between measurements and model is fully contained in the likelihood term $P(\mathcal{Q}|d_i)$. The right hand side will take small values for large

absolute values of the difference between the model output and the measured data, $\epsilon_i = d_i - \hat{d}_i$, thus modifying the resulting posterior PDF. Intuitively, this is equivalent to testing different hypotheses about the parameter values contained in $\mathcal{Q}$, and assigning a low probability to ones that lead to large discrepancies between measured and simulated activation times. If, for some measurement point, ϵ_i does not vary at all over the range of parameter values investigated, it cannot have any impact on our uncertainty about the true parameter values. Conversely, if ϵ_i varies greatly, the structure of Eq. 5.3 ensures that it will reduce the probability for certain parts of the parameter space, and thus likely contribute towards decreasing the width of the posterior PDF, describing our uncertainty about the true parameter values. This suggests that the degree of variation in activation time at a given point in the tissue will be important for how much information can be gained from measuring the activation time at that point. This leads to a measurement selection criterion akin to the Fisher information measure [185]. The degree of variation can be described as

$$S_i^2 = \sum_{\mathcal{Q}} (\epsilon_i - E(\epsilon_i))^2 P(\epsilon_i), \tag{6.1}$$

or, since the dependency on d_i will cancel out,

$$S_i^2 = \sum_{\mathcal{Q}} (\hat{d}_i - E(\hat{d}_i))^2 P(\hat{d}_i), \tag{6.2}$$

where $E(\hat{d}_i)$ is the expected value of $\hat{d}_i$, and where the index i spans the N regions on the endocardium as described in Section 6.2.1.

With this analysis in mind, measurements on the LV endocardium were placed according to three strategies. The first strategy consisted of randomly choosing points on the LV endocardium, until an approximately uniform coverage was achieved. The second strategy consisted of selecting every measurement point as the one that lead to the greatest decrease in uncertainty about the true parameter values, corresponding to the approximate solution to the path finding

problem described in Fig. 6.1. Finally, the third strategy consisted of visiting the points with an order according to their standard deviation in activation time over all investigated sets of parameters, starting with the point of highest standard deviation. These three strategies will be termed the *random*, *greedy* and *heuristic* strategies, respectively.

6.2.4 Visualization of the LV Endocardium

The anatomical model employed for the study presented in this chapter is the same as that described in Section 5.2.1 of the previous chapter. As an aid for visualizing results on the LV endocardium, a 2D projection based on a cylindrical coordinate transform has been employed. The projection is illustrated in Fig. 6.2.

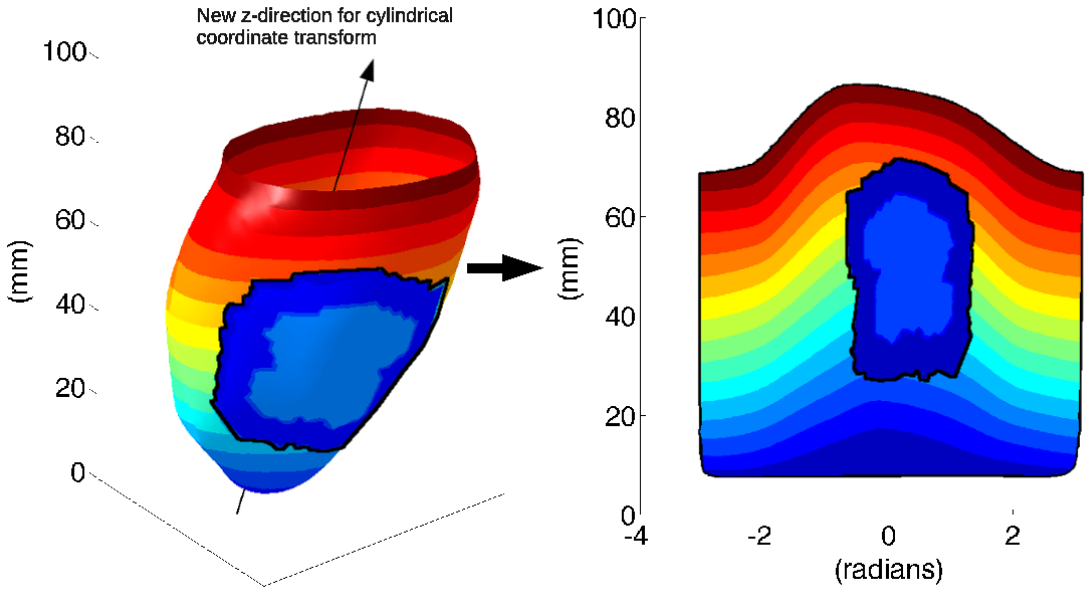


Figure 6.2: Illustration of cylindrical coordinate transform. Left panel shows the left ventricular endocardium. Right panel shows the same surface transformed into cylindrical coordinates. The central area corresponds to the endocardial side of the transmural scar, with border zone in dark blue and core zone in light blue.

6.3 Results

In this section we present results demonstrating and comparing the effects of the three measurement strategies detailed above. We illustrate the distribution of measurement points across the accessible tissue, and give details on the resulting reduction in the Bayesian posterior distributions for the parameter values. For reference, in Fig. 6.3 we show the 2D projection of the activation times that was used as ground truth for the results on measurement placement presented below.

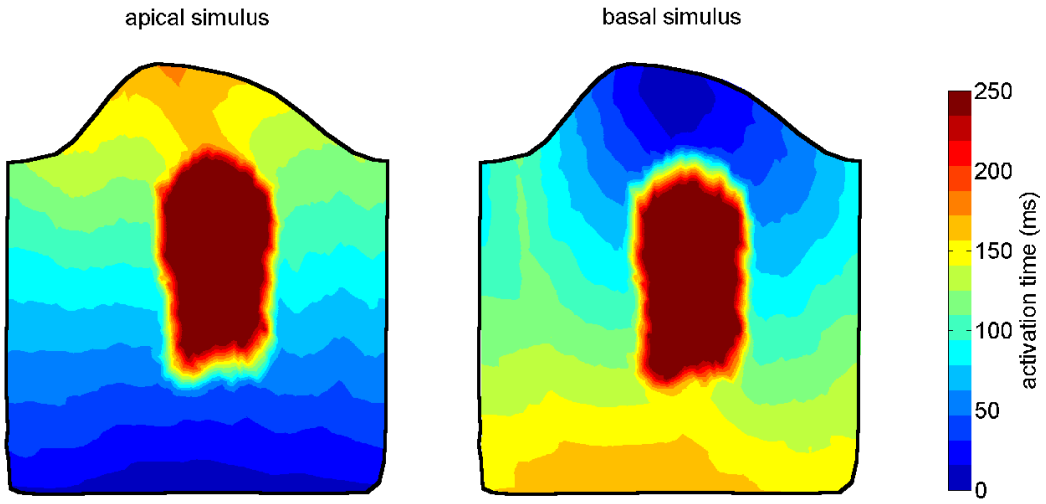


Figure 6.3: Ground truth activation times as projected using the 2D transform described in the previous section. Activation times shown are limited to 250 ms and below in order to clearly depict the activation pattern outside the infarcted tissue.

6.3.1 Effects of Measurement Placement

After having investigated the effects of accuracies in simulations and measurements on the posterior PDFs resulting from the inference procedure described in the previous chapter, we now turn our attention to the challenge of finding an efficient strategy for placing endocardial EP measurements. As a setting for this, data from virtual measurements were generated as

described in Chapter 5 for two different points of initial activation of the human ventricles. Both points were placed on the endocardium of the LV, one close to the apex and the other close to the base, simulating pacing of the tissue with a catheter.

The inference procedure described in the previous chapter was again applied to computing the simultaneous posterior PDF of the parameter values, but now using the *greedy* strategy for measurement placement, as described in Section 6.2.2. Rather than dictating a uniformly distributed placement of measurements, this strategy resulted in well defined patterns of measurements forming on the endocardial surface. Additionally, for each node in the mesh, the standard deviation in activation time across all parameter values was calculated. The standard deviation and point placement resulting from the *greedy* strategy are shown in Figs. 6.4 and 6.5 for apical and basal stimuli respectively. The points selected by the *greedy* stategy are color coded according to their order, with colors progressing from blue (indicating points that were picked early) to red (indicating points that were picked late) over the standard color spectrum. Also shown are the exact placements the stimulation points, which are marked with an X in the top left sub-panels of the left and right panels respectively. As discussed in Section 6.2.3, the standard deviation gives an indication of the sensitivity of each of the nodes to variation in parameters. By comparing the pattern of standard deviation across the nodes with the pattern of measurements from the *greedy* strategy, we verify that the standard deviation gives an indication of where to place measurements to obtain a fast decrease in uncertainty about the true parameter values.

Before moving on to analyzing the results, we need to clarify our use the concepts of ”upstream” and ”downstream”, since it will serve to simplify the discussion. By ”x is upstream of y” we will mean that y is not part of the path that leads from the point of initial activation to x *for any set of parameters*. Conversely, by ”x is downstream of y” we mean that y *is* part of the path that leads from the point of initial activation to x *for at least one set of parameters*.

Inspecting Figs. 6.4 and 6.5, we see that the distribution of points selected according to the *greedy* strategy varies with the level of model and measurement noise, while the background

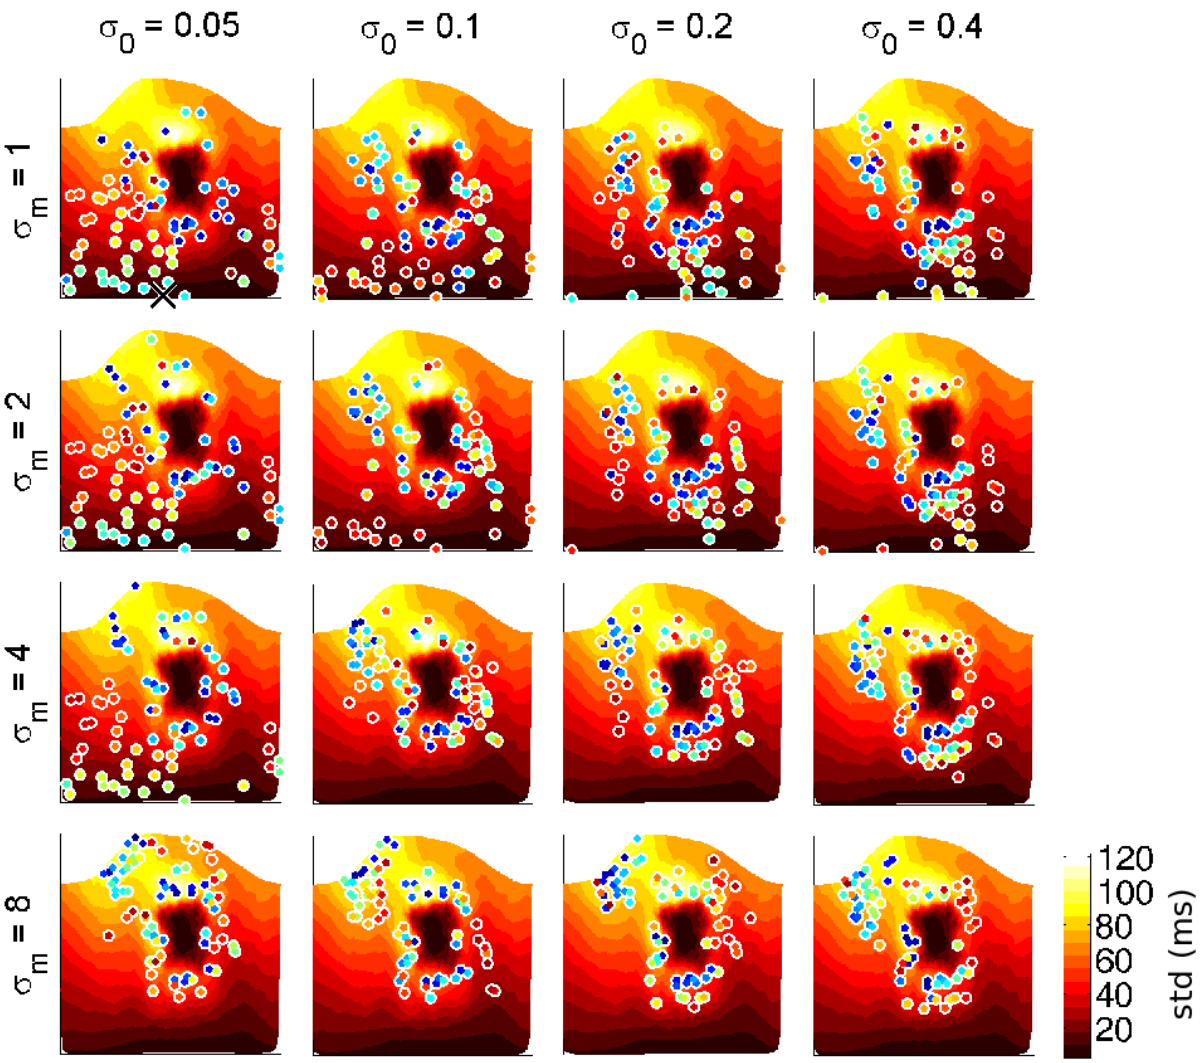


Figure 6.4: Map of measurement points selected according to the *greedy* strategy on the LV endocardium, given initial activation of the tissue close to the apex, as indicated by the X on the upper left subpanel. Colored circles correspond to the measurement points as chosen using the greedy algorithm. The circle color indicates the order in which the points are chosen, with blue being the earliest and red the latest. The background color gradient shows a map of the standard deviation of activation time (S) over all parameter sets.

map of standard deviation remains constant, being independent of the noise. As indicated in the figures, the level of measurement noise increases in the downwards direction, while the level of modelling noise increases when moving to the left. Both figures exhibit the same trend: for low levels of noise, the points selected by the *greedy* strategy tend to be distributed across a large area of the tissue, but as the level of noise increases, the points begin to cluster around areas

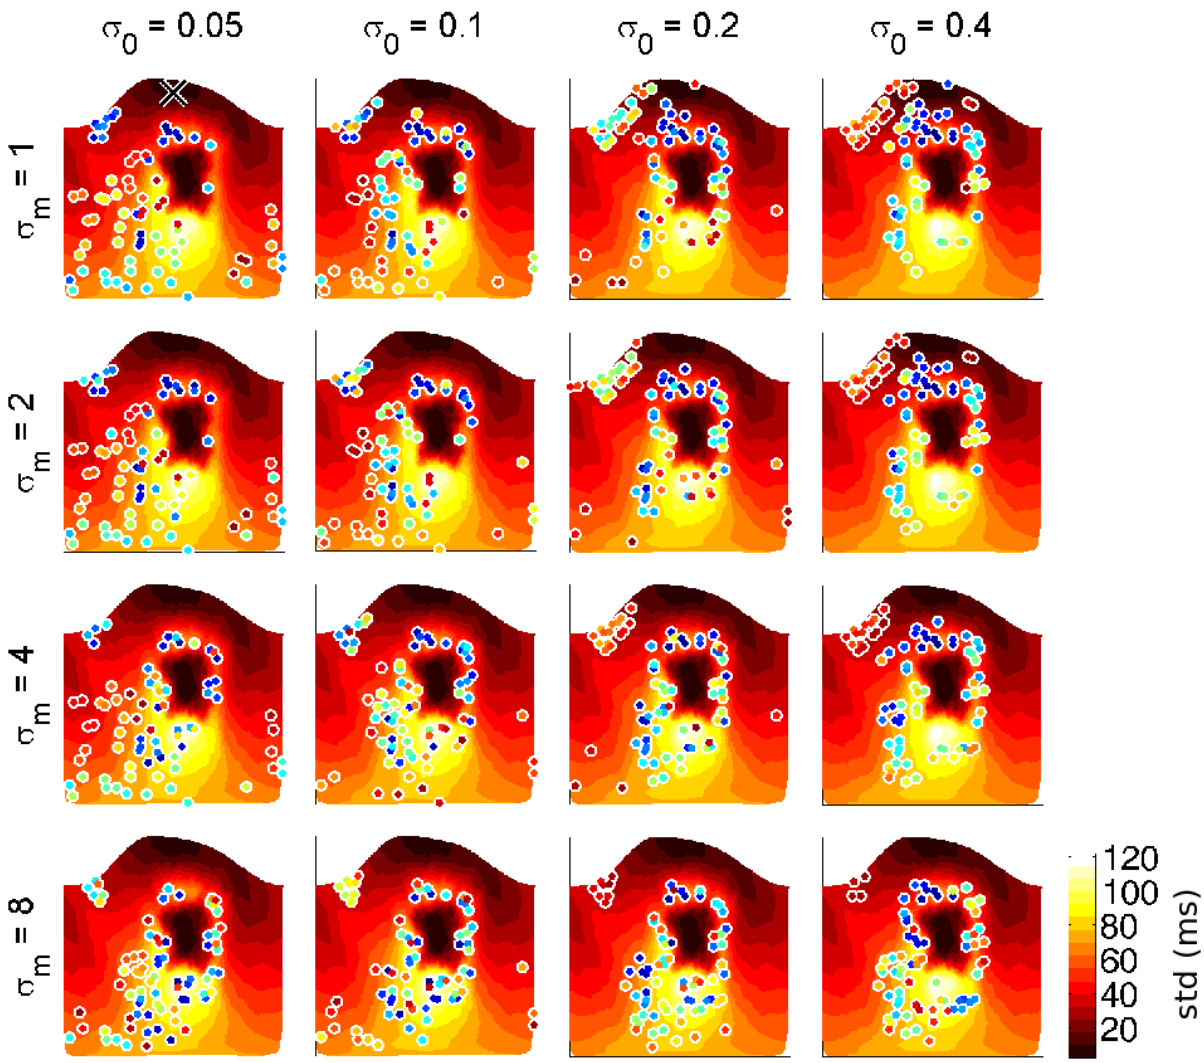


Figure 6.5: Map of measurement points selected according to the *greedy* strategy on the LV endocardium, given initial activation of the tissue close to the base, as indicated by the X on the upper left subpanel. Colored circles correspond to the measurement points as chosen using the greedy algorithm. The circle color indicates the order in which the points are chosen, with blue being the earliest and red the latest. The background color gradient shows a map of the standard deviation of activation time (S) over all parameter sets.

with a high standard deviation. A likely reason for this behaviour is that useful information can be extracted almost anywhere for low noise levels. Conversely, at higher noise levels, it becomes increasingly important to choose points where the variation in activation time is large, since the points with only minor variations will be ”masked” by the noise.

The color values of the points adds more detail to the picture. Blue points, which are picked

early by the *greedy* algorithm, seem to exhibit a less variation in their position than red points, which are picked late. In both Fig. 6.4 and 6.5, the earliest selected points tend to cluster around the infarct, while the points selected later, as indicated by green, yellow, orange and red colors, tend to be more spread out. Another interesting trend is that certain areas of the tissue, although remaining occupied by a number of points for most noise levels, decrease in "popularity" with increasing noise. Initially, these areas populated by blue or green points, but as noise levels increase, these are replaced by orange and red points. Perhaps the most clear example of this phenomenon can be seen at the left side of the basal area in Fig. 6.5, where a small colony of blue points at the lowest level of noise (upper left panel) gradually gets replaced by other colors as noise levels increase, finally ending up as a small colony of red points for the highest noise level (lower right panel). This, together with the observations regarding the over all distribution of points listed above, strengthens the previous conclusion that increasing noise levels increases the preference of the *greedy* strategy for parts of the tissue that exhibit large variations in activation time.

Fig. 6.6 shows the pattern of standard deviation in activation time individually for each of the four parameters. Four values of standard deviations per node were calculated by keeping three of the parameters fixed to the centre of the range of investigated values, and computing the standard deviation in activation time over the range of the remaining parameter. Subsequently, for each node these four values were divided by the average activation time of that node, compensating for the larger variation in nodes that are further away from the point of initial activation.

The cv_l/cv_t and cv_l/cv_n panels exhibit inverse correlation in standard deviation for both points of initial activation. Nodes with a high standard deviation for the cv_l/cv_n parameter show low values for cv_l/cv_n and *vice versa*. The actual pattern is strongly dependent on the fibre orientation, since the parameters essentially describe the conduction speed in two perpendicular directions. For the apical stimulation site (upper row in the figure), the pattern manifests itself as an almost vertical band of lower standard deviation of cv_l/cv_t just to the right of the scar. It

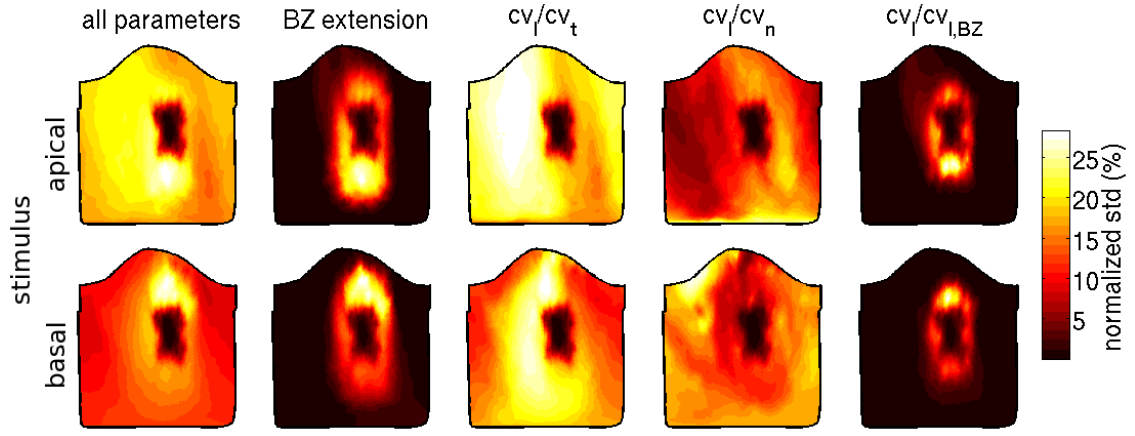


Figure 6.6: Maps of standard deviation of activation time per node, normalized by average activation time per node, when varying each parameter individually, as indicated by the lables in the figure. Lighter color corresponds to higher standard deviation.

is matched by a corresponding band of low standard deviation on the cv_l/cv_n panel. Here, it is likely that the current propagation is roughly aligned with the normal direction of the fibres, and thus almost perpendicular to the transverse direction. Consequently, the parameter cv_l/cv_t, governing the propagation speed in the transverse direction, has a lesser impact on the activation time in this region.

For the basal stimulation site (lower row in the figure), the cv_l/cv_t and cv_l/cv_t panels show a similar pattern to those arising from the apical stimulation site. For the cv_l/cv_t parameter, a region of high standard deviation can be seen just to the left of the scar, partly enclosing it at the top and bottom. Again, the cv_l/cv_n panel contains a roughly corresponding band of low standard deviation. Interestingly, for the basal stimulation site, an area of high standard deviation in cv_l/cv_n can be seen at the left hand side of the base. Going back to Fig. 6.4, it can be seen that this area corresponds to a cluster of *greedy* measurement points that is conserved for all levels of noise on the basal stimulation panel. This strongly indicates that this area is an important source of information about the cv_l/cv_n parameter.

The cv_l/cv_{BZ} and BZ extension panels of Fig. 6.6 display high values of standard deviation in the border zone. This is to be expected, since both parameters govern the conduction prop-

erties connected to the border zone. It can also be seen that the regions of high standard deviation around the scar are larger above and below the scar than to either side. This effect is the most pronounced on the panels corresponding to the BZ_{ext} parameter, but also visible on the cv_l/cv_{BZ} panels. The effect might partly be due to the geometrical distortion introduced by the coordinate transform, but is also likely to be influenced by the direction of conduction. Interestingly, when the standard deviation pattern is normalized by the average activation time per node, as in Fig. 6.6, the highest values actually occur at the same side of the scar as the stimulation is situated. The reason for this is not obvious, but it might be connected to the fact that the current probably reaches these areas along the same path irrespective of the chosen set of parameters.

Finally, we examine the effects on the posterior PDF from the three different measurement strategies described in Sections 6.2.2 and 6.2.3: 1) the *random* strategy, consisting of randomly picking the next point, 2) the *greedy* strategy, consisting of always picking the point leading to the largest decrease in probability mass of the posterior PDF, and 3) the heuristic strategy, consisting of visiting the points in order of decreasing standard deviation in activation time. A rough measure of the relative posterior width was computed by counting the minimum number of discrete points needed to cover 95% of the total probability, and dividing this number by the total number of discrete points (80000). This was done for each combination of model and measurement noise, and for 100 successive point measurements, ordered according to the three strategies. Average values of the relative posterior width for 100 runs of this experiment are shown in Fig. 6.7. Results in the figure correspond to the case of apical stimulation. The graphs in the figure give a picture of how the different measurement strategies impact the uncertainty in parameter estimates during the course of data collection. Results for the basal stimulation are almost identical and not shown here.

As can be seen from Fig. 6.7, for all investigated levels of noise, the *greedy* and *heuristic* strategies lead to a faster reduction in uncertainty than the *random* strategy. For low to moderate

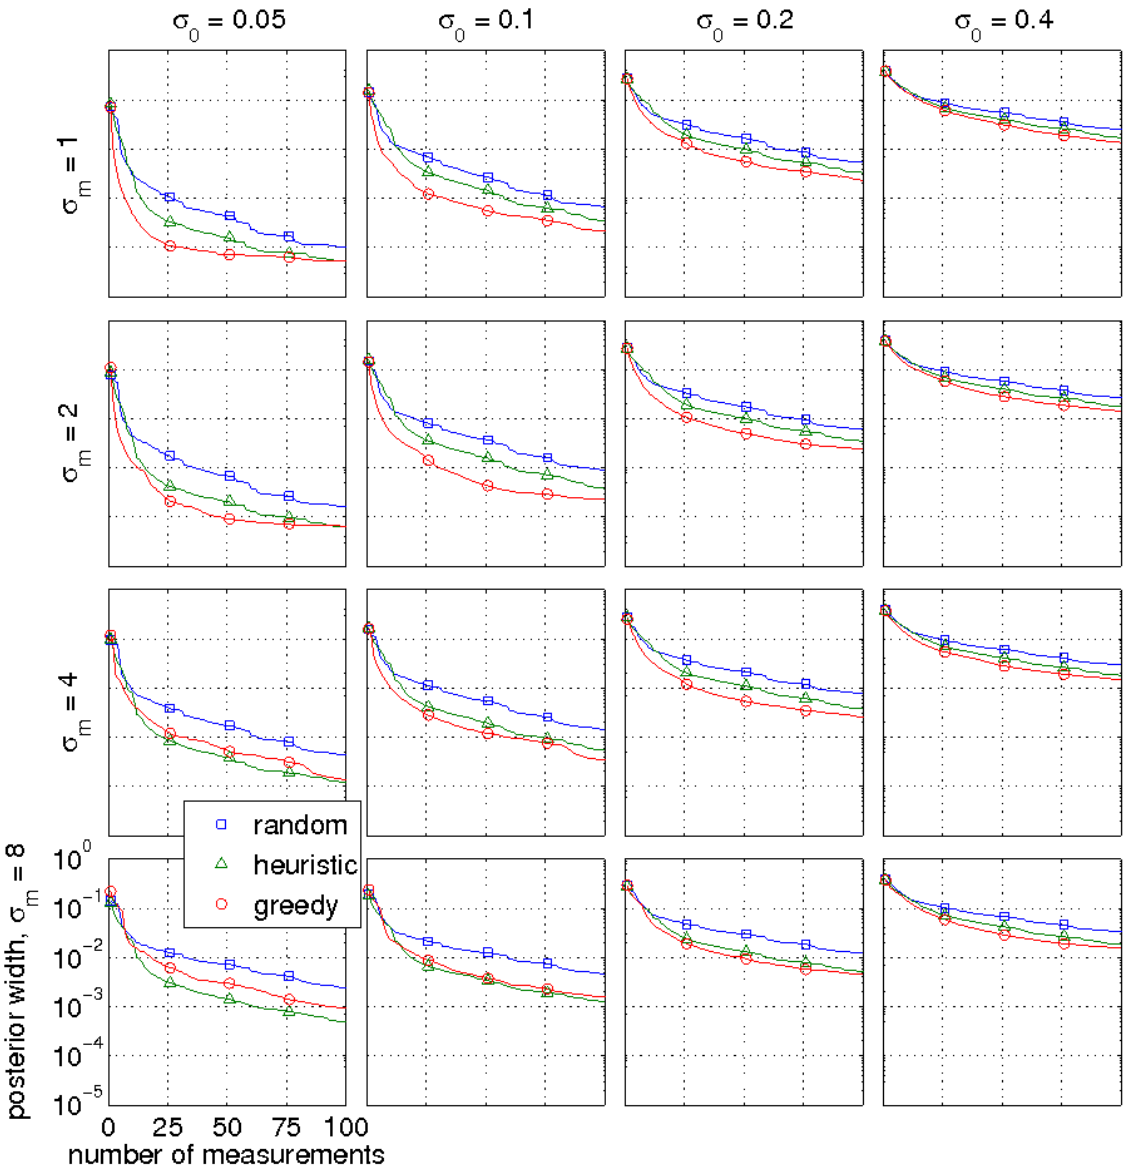


Figure 6.7: Relative width of the posterior PDF for different combinations of model noise ($\sigma_e(\mathbf{q})$) and measurement noise (σ_m). Lables on the top and left hand side of the figure show the correspondence between noise components and figure panels. Each panel shows three curves, corresponding to the three measurement strategies *random* (blue lines, square markers),*heuristic* (green lines, triangular markers), and *greedy* (red lines, circular markers).

levels of modelling noise, this effect is fairly pronounced, with results for the *greedy* strategy after 25 measurements often being comparable to the *random* strategy after 75 measurements. For high modelling noise, the advantage of the *greedy* and *heuristic* strategies are still present, but less pronounced.

As for the two lowest levels of modelling noise, corresponding to the two leftmost columns of Fig. 6.7, the figure shows that as measurement noise increases, the results for the *heuristic* strategy approaches those arising from the *greedy* strategy. At the three panels closest to the bottom left corner of the figure, corresponding to high measurement noise and low modelling noise, the *heuristic* strategy actually does better than the *greedy* strategy. This is both surprising and significant, since the *greedy* strategy actually relies on having seen all data in advance, while the *heuristic* strategy only relies on information available before any data have been obtained.

Additionally, the data from the 100 virtual experiments described above were analysed with respect to average errors and credibility regions, expressed in terms of conduction velocities. Here, we assumed a conduction velocity of 0.7 m/s in the longitudinal direction (cv_l), and ratios between directional velocities corresponding to those in Fig. 5.2. Using these values, the average absolute error and the average length of the 95% credibility regions were computed for all parameters and for each strategy. The average errors were calculated as the absolute difference between the true parameter value and the value indicated by the highest probability point of the normalized posterior PDF, while the credibility regions were calculated from the marginal distributions of each parameter according to Eq. (3.19). Results after 50 and 100 simulated measurements are shown in Tables 6.1 and 6.2.

Looking at Table 6.1, we see that the *greedy* and *heuristic* strategies both give lower average errors than the *random* strategy in almost all cases, with the *greedy* strategy generally giving the lowest errors. Where exceptions occur, such as for rows C and D for the cv_t parameter, the discrepancy is small. The advantage of the *greedy* and *heuristic* strategies appears more pronounced for cv_n and $cv_{l,BZ}$ than for BZ_{ext} and cv_t, and more pronounced for higher levels of uncertainty (B,C and D) than for low levels (A). This is particularly visible in the B row, corresponding to a high uncertainty in measurements but fairly accurate modelling, where the *heuristic* strategy results in significant improvements in accuracy, even compared to the *greedy* strategy, as also visible in Fig. 6.7. Contrastingly, the cv_t parameter seems the most insensitive

		$BZ_{ext}(mm)$			$cv_t(cm/s)$			$cv_n(cm/s)$			$cv_{l,BZ}(cm/s)$		
	n	R	H	G	R	H	G	R	H	G	R	H	G
A	50	1.32	1.08	1.02	0.33	0.35	0.27	1.76	1.50	1.37	6.28	5.66	1.32
	100	1.03	0.98	0.99	0.05	0.03	0.09	1.01	0.87	0.75	2.56	2.54	0.91
B	50	1.38	1.25	1.14	0.89	0.35	0.37	3.71	1.55	3.24	11.29	7.62	2.72
	100	1.19	1.11	0.97	0.61	0.12	0.12	2.57	1.16	1.52	6.16	3.57	2.45
C	50	2.11	1.80	1.83	1.58	1.58	1.73	3.79	3.22	3.72	19.63	15.86	8.00
	100	1.75	1.66	1.42	0.94	1.06	1.53	3.29	3.68	3.02	13.29	11.44	6.86
D	50	2.23	1.58	1.75	1.63	1.84	1.59	4.53	3.02	4.59	21.30	15.26	10.38
	100	1.80	1.61	1.45	1.45	1.22	1.64	4.25	3.39	3.37	13.78	11.51	6.71

Table 6.1: Average absolute errors over 100 trials. The table shows average absolute errors for the estimated variables (separated triplets of columns) after 50 and 100 measurements respectively (separated pairs of rows). A triplet of columns contain values for the *random* (R), *heuristic* (H) and *greedy* (G) strategy respectively. These data are shown for different levels of uncertainty in modelling and measurements, corresponding to the letters A, B, C and D. The letter A signifies the combination $\sigma_m = 2, \sigma_0 = 0.1$, B signifies $\sigma_m = 8, \sigma_0 = 0.1$, C signifies $\sigma_m = 2, \sigma_0 = 0.4$ and D signifies $\sigma_m = 8, \sigma_0 = 0.4$.

to the positioning of measurements. Considering results from Fig. 6.6, this is not surprising, since useful information about the parameter can be found in almost any part of the tissue.

Considering the lengths of the 95% credibility regions shown in Table 6.2, a similar trend to that in Table 6.1 can be observed: the *greedy* and *heuristic* strategies generally produce shorter credibility regions. Here, the *greedy* strategy gives an especially pronounced advantage for the BZ_{ext} parameter, with reductions in uncertainty sometimes in excess of 50% compared to competing strategies. The *heuristic* strategy, on the other hand, has its strongest effects in the B row, where it results in large reductions of the credibility regions of cv_l/cv_n and $cv_l/cv_{l,BZ}$.

6.4 Discussion

In this chapter we presented a heuristic for determining how to optimally place endocardial point measurements during electroanatomical mapping. The heuristic is based on the degree of variation in activation time for individual points on the cardiac surface when varying the parameters. The heuristic strategy for measurement placement has further been compared to two

		$BZ_{ext}(mm)$			$cv_t(cm/s)$			$cv_n(cm/s)$			$cv_{l,BZ}(cm/s)$		
	n	R	H	G	R	H	G	R	H	G	R	H	G
A	50	3.41	3.51	1.70	5.73	5.43	5.35	7.92	5.92	6.52	21.16	20.45	6.80
	100	2.33	2.35	1.64	5.02	4.99	5.04	4.55	4.35	4.35	13.07	11.63	5.82
B	50	5.11	4.09	2.86	7.57	5.96	6.58	12.20	6.47	12.92	33.00	24.71	13.88
	100	3.72	2.93	2.38	6.09	5.07	5.52	9.95	5.10	8.67	23.99	16.93	12.05
C	50	6.75	6.76	5.03	11.24	11.84	11.60	14.12	13.53	14.11	42.85	40.56	27.52
	100	5.92	5.93	4.22	9.07	8.62	10.00	13.40	12.82	12.18	36.99	34.77	24.66
D	50	6.84	6.74	5.93	12.21	12.34	10.72	14.00	13.42	14.73	44.47	39.11	30.71
	100	5.92	5.87	4.49	9.61	8.76	9.88	13.68	12.99	13.44	36.30	32.93	24.76

Table 6.2: Average credibility regions lengths over 100 trials. The table shows lengths of CR for the estimated variables (separated triplets of columns) after 50 and 100 measurements respectively (separated pairs of rows). A triplet of columns contain values for the *random* (R), *heuristic* (H) and *greedy* (G) strategy respectively. These data are shown for different levels of uncertainty in modelling and measurements, corresponding to the pairs of rows. The letter A signifies the combination $\sigma_m = 2, \sigma_0 = 0.1$, B signifies $\sigma_m = 8, \sigma_0 = 0.1$, C signifies $\sigma_m = 2, \sigma_0 = 0.4$ and D signifies $\sigma_m = 8, \sigma_0 = 0.4$.

other strategies: the random and the greedy strategy. The random strategy consists of randomly distributing the measurements across the accessible parts of the cardiac surface, whereas the greedy strategy consists of always choosing the point leading to the largest reduction in posterior width.

Results presented in Figs. 6.4 and 6.5 show the pattern of 50 measurement points as determined by the greedy strategy as described in Section 6.2.2, superimposed on a map of the measure of activation time variability S as described in Section 6.2.3. This illustrates how, as noise levels increase, the distribution of the points according to the greedy strategy becomes more concentrated to the regions where the S is high. This indicates that the *a priori* measure S provides a good approximation of the *a posteriori*, greedy measurement selection criterion, especially for high uncertainties in modelling and measurements. Here, the greedy stategy was based on the reduction of the 95% posterior confidence region, but other choices of PDF dissimilarity measures, such as the Kullback-Liebler distance, are also possible. A more in-depth investigation of the impact of the choice of measure might form an interesting line of future investigation.

Figure 6.7 and Tables 6.1 and 6.2 complete the picture by illustrating the impact on the errors and uncertainty in the parameter estimates resulting from the three strategies. The results in Fig. 6.7 indicate that the location of measurements have a significant impact on the uncertainty inherent in the parameter estimates. In fact, for high levels of measurement noise, the figure shows that the standard deviation based measure might sometimes outperform the reduction in posterior width (corresponding to the *greedy* measurement strategy). This can also be seen in both Tables 6.1 and 6.2, most clearly on the cv_n parameter, but it should be noted that the marginalization procedure might obscure some of the reduction in uncertainty of the full posterior PDF. An important consequence of this is the possibility of obtaining an *a priori* estimate of how to place the electrophysiological measurements. Given a geometrical description of the tissue, along with a parameterization such as the one presented in this work, a map of S across the tissue can be constructed prior to obtaining data, and the calculated map can thus be used to guide the data acquisition.

The fact that the standard deviation in activation time (S) works well for optimizing measurement placement also highlights the relationship between the way the model is parameterized and where to place measurements. In other words, if we change the parameterization of the model, we also likely change the optimal pattern of measurements. The results in Section 6.3.1 seem to imply that the information yield gets higher far away from the site of initial activation, and downstream of conduction heterogeneities. However, the picture is further complicated by the fact that the concept of downstream itself depends on the parameters and thus has to be defined probabilistically. Thus, there is no general solution to the problem of how to best obtain data during an electroanatomical mapping procedure. Rather, in order to maximize the information content of the measurements, the results presented here suggest that the parameterization should dictate the data acquisition, a fact which will be further discussed in the next chapter.

7

Conclusions and Future Directions

We conclude this dissertation by reviewing and discussing the achievements and results described in the previous chapters. We begin this chapter with Section 7.1, in which we review the original motivation for this dissertation and summarise the work done to achieve these aims. Section 7.2 lists the key contributions presented. Section 7.3 discusses potential future improvements to the inference framework, as well as potential applications of the developed methods. Finally, in Section 7.4, we end the dissertation with some concluding remarks.

7.1 Summary

Cardiovascular diseases (CVDs) are the number one cause of death globally. More people die annually from CVDs than from any other cause. Although the CVD category encompasses a heterogeneous range of conditions, many of these diseases are associated with abnormalities in the cardiac electrical propagation. Despite the fact that the fundamental mechanisms underlying many conduction abnormalities remain unclear, cardiac infarctions, leading to the formation of scar tissue, have been implicated as a contributing factor. There currently exists a number of clinical interventions that rely on modifying the electrical propagation through direct intervention in order to improve cardiac function. Notable examples of this are cardiac

resynchronization therapy (CRT) and ablation therapy.

Computational modelling and simulation of the heart have emerged as increasingly important tools for cardiac research, with the potential to aid in optimizing existing therapies, as well as to provide new insights into the disease mechanisms. A crucial part of a computational cardiac model is the definition of tissue conduction properties, which have a profound impact on the cardiac activation sequence predicted by simulations. The conduction properties may vary spatially throughout the anatomical model to represent different regions of tissue, and may also vary depending on direction, reflecting the complex spatial orientation of the cells that make up the cardiac tissue.

Thus, a necessary step in constructing a personalized cardiac model is to obtain measured data that can be used as a basis for inference about the underlying electrophysiological properties of the cardiac tissue. In a clinical, *in vivo*, setting, the available data are generally sparse. Spatial data can be obtained non-invasively using imaging modalities such as MRI, CT or ultrasound, but reliable electrophysiological (EP) information still needs to be obtained invasively using catheterization. This limits the amount and quality of data that can be obtained, because of constrains in both time and accessibility. Clinical EP contact measurements typically consist of point measurements obtained exclusively from the endocardial surface. Since the data sets are often small, and associated with various degrees of noise, extracted parameter values will be associated with a potentially large degree of uncertainty. This poses two challenges: Firstly, it is important to be able to quantify the uncertainty in extracted parameters, to prevent estimation errors from being translated to errors in therapy. Secondly, under clinical time constraints, it becomes important to maximize the information content of the obtained data by placing the measurements in an optimal way. It has been the main goal of this dissertation to contribute to addressing these two challenges.

In Chapter 2 we gave a description of the structure and function of the heart, with special focus on the electrophysiological properties. Further, we described how conduction blocks caused

by abnormalities such as scars can impair the normal function of the heart, leading to disease conditions such as arrhythmias and heart failure. Later in the chapter, we also reviewed the main electroanatomical mapping systems for collecting clinical electrophysiological data, with special focus on the accuracy with respect to the electrical activation time. This was followed by a historical overview of the main efforts within the field of cardiac modelling during the last 60 years, before finally giving an account of electrophysiological parameter estimation.

In Chapter 3 we treated the main computational methods used in the work presented in this dissertation. We started by describing the monodomain and bidomain models, representing the current state of the art of detailed modelling of the cardiac electrophysiology, along with an overview of the main numerical methods used to solve the resulting equations. Next, we gave an account of the eikonal equations, representing the main phenomenological model for describing the current propagation in cardiac tissue. After this, we discussed the fast marching method for solving the eikonal equations. Finally, we gave a detailed account of the inverse problem encountered during parameter estimation, setting the scene for the final part of the chapter, describing the Bayesian inference methodology used throughout this dissertation.

Next, in Chapter 4, we presented a novel graph based alternative to the fast marching method. As a validation of the graph based method, the chapter includes a comprehensive comparison between results obtained using the proposed method and corresponding results obtained using FEM solutions of the monodomain equation and fast marching solutions of the eikonal equations. Additionally, the computational performance of the graph based method was studied, and was shown to compare favourably to the the performance of the two other methods.

In Chapter 5 we developed a methodology for Bayesian inference, driven by the graph based method developed in Chapter 4. The Bayesian inference procedure was tested in a simulation study performed in an image-derived human geometry, incorporating conduction anisotropy and a conduction block in the form of a transmural scar. In addition to demonstrating the usefulness of the proposed approach, the results showed that multimodal posterior distributions

are common, even for the relatively low-dimensional problem treated in the study. This strongly indicates that careful analysis of the Bayesian posterior PDFs is necessary.

Building on the results presented in Chapter 5, in Chapter 6 we presented a study of how to optimally position endocardial point measurements during clinical electroanatomical mapping in order to maximize their information content. Here we compared three different strategies for placing the measurements: one naive, random strategy, distributing the measurements randomly across the entire cardiac surface; one greedy strategy, where after each measurement the posterior PDF resulting from each possible choice of next measurement point was pre-computed, and the one leading to the largest decrease in uncertainty was chosen; and one heuristic strategy, where the standard deviation in activation time across the entire parameter space was used as an index to decide the order and placement of the measurement points. It was shown that the greedy strategy, utilizing *a posteriori* information was usually superior. However, the heuristic strategy, only relying on *a priori* information, was always superior to the random strategy, and sometimes also outperformed the greedy strategy, especially in settings where measurement data were associated with a high degree of uncertainty.

Finally, in this chapter we list the key contributions presented in this dissertation. Further, we discuss future refinements, extensions and possible applications of the developed methods, before ending the dissertation with some concluding remarks.

7.2 Key Contributions

The first major contribution within the scope of this dissertation is the graph based method for activation times approximation presented in Chapter 4. The method facilitates very fast evaluation of the activation times in different points in the cardiac tissue, in particular for single or small numbers of points. This makes the method especially convenient to use for matching simulated data and experimental data in the context of electroanatomical mapping studies,

where a relatively small amount of point measurements are obtained in a sequential manner. The graph based method makes it possible to simulate the outcome of a given pacing protocol essentially in real-time on a current desktop computer with accuracy comparable to that of the fast marching method, making it an attractive alternative for parameter estimation applications, where activation times must be calculated for a large number of different parameter values.

The second major contribution of this dissertation is the development of a Bayesian inference procedure, incorporating the graph based method for activation time approximation from Chapter 4, to produce probabilistic parameter estimates based on anatomical information from imaging and electroanatomical mapping data. Using this method, probability distributions of parameters governing the electrophysiological conduction can be estimated with feasible computational costs, giving the investigator a comprehensive description of the uncertainty associated with a particular parameter estimate. The uncertainty in estimates is a crucial component in any real-world clinical application, where the cost of mistakes can be very high.

The third major contribution presented in this dissertation is the development of an index for optimizing the placement of electrophysiological measurements on the cardiac surface. Using this index, which can be computed *a priori* given knowledge of the cardiac model and its parameterization, the surface of the accessible tissue can be mapped according to its spatial information content, facilitating an optimized distribution of the limited number of measurements that are possible to obtain within clinical time constraints. These results also highlight the important relationship between the model, the parameterization and the spatial distribution of information, with different models and parameterizations giving rise to different optimized distributions of measurements.

7.3 Future Directions

Finally, we suggest a number of areas of possible further research. We begin by considering aspects of the presented methods and their validation that could be improved or developed. We then discuss some of potential applications of the methods, with the potential to further improve the amount and type of information that can be extracted from clinical electrophysiological data.

7.3.1 Extensions of the Parameter Estimation Framework

In Chapters 5 and 6 we have demonstrated that the developed method for probabilistic parameter estimation is able to accurately compute posterior PDFs of the parameter values of interest, given correct error distributions and a sufficiently flexible parameterization of the conduction model. The choice of geometries used in the studies presented in the studies in Chapters 4, 5 and 6 have inevitably affected the outcomes to some degree. As demonstrated in the two last chapters, the conduction velocity in the normal direction is possible to estimated using only endocardial measurements. Although it seems reasonable that some information can be gained from this type of data, it is perhaps surprising that the parameter can be estimated with the demonstrated accuracy. This might be an artifact of the algorithmically generated fibre directions. To investigate this further, it would be interesting to study the relation between fibre orientation of the model and the amount of information that can be extracted from measurements about each parameter. Another related line of expanding the inference study is to incorporate the epicardium into the surface accessible for measurements. Experiments of this type could help to quantify how much extra information can be gained by incorporating epicardial measurements, and thus provide a basis for decisions about how to design an EP study to maximise the information yield with respect to a given model.

An important feature of the Bayesian inference framework, as described in Chapters 3 and 5, is the ability to formally incorporate previous knowledge into the inference procedure. This is

done by modifying the prior distribution in for example Eq. (3.15) to incorporate any *a priori* information we might have on the probability of different parameter values. In this work, we have taken a conservative approach in that we have chosen the prior to be completely flat, reflecting a total lack of preconceptions regarding the parameter values (except for assumptions on the range of possible values). However, given the increasing amount of clinical data accumulated during the last decade, along with the numerous efforts to structure this knowledge in the form of atlases and data bases [186, 187, 188], we are approaching a situation where reliable priors can be constructed from measured data. An informative and correct prior can greatly improve the accuracy in new estimates, even when working with very limited data sets. However, great care must be taken when constructing and choosing these priors, to assure that the case for which the new estimation is performed really is part of the population from which the prior was constructed. A mismatch at this stage could easily lead to a significant bias, possibly resulting in erroneous conclusions. A large set of clinical data from electroanatomical mapping systems, such as from EnSite NavX or CARTO, would also serve to elucidate what error distributions are likely to occur in a clinical scenario, and what parameterizations of the model are adequate to capture the structural and functional heterogeneities associated with various disease conditions. Using this information to adapt error distributions and the prior, the developed inference method could be further validated.

The studies presented in Chapters 5 and 6 concern the estimation of four parameters; conduction velocities in the transverse and normal directions, a proportionality factor between infarcted and non-infarcted tissue, and a parameter regulating the extent of the scar border zone. Of course, given the plethora of other possible parameterizations of any cardiac model, an obvious extension to the presented studies would be to investigate the impact of these. An increased detail in the parameterization of the model corresponds to an increase in dimensionality of the parameter estimation problem. For higher dimensional problems, such as those treated in for example [125], memory and computation time requirements will make the grid based evaluation imprac-

tical, and alternative ways of computing the posteriors, such as MCMC methods [25] must be employed. Although this presents an additional challenge in terms of the computational effort, the great majority of results presented here are independent of how the posterior is computed and thus equally valid when other methods are employed.

Another interesting line of investigation is to improve the measurement placement further by successively incorporating information from newly obtained measurements into the index to interactively improve it at the measurement collection stage. The heuristic index S, as presented in Chapter 6, relies on analysing the full range of possible activation times for each node, prior to having seen any data. However, as measurements are obtained, some of these possible activation times may start to appear very unlikely, given our prior knowledge about measurement accuracy. This, in turn, may restrict the amount of variation in activation time that is likely for some nodes of the computational mesh, and thereby modify S. Another way of viewing this is that since S for any node depends on the expected value of the activation time, and since the expected value is affected by the measurements, it is reasonable that S is affected as well. In a sense, this reasoning is already partly incorporated in the proposed strategy for measurement placement; the reason that we do not simply continue to place measurements in the the point with the largest S value is that we get significantly more certain of the activation time at that point after having seen the result of the measurement. This reduces the information about the parameters that remains to be extracted from that particular point. This is easily realized when considering a situation where we are completely certain about the activation time in a point. In such cases, additional measurements in that point will tell us nothing we do not already know, and thus cannot possibly affect our state of knowledge about the model parameters.

Yet another option for extending the parameter estimation framework presented in this dissertation is the incorporation of additional electrophysiological data modalities, such a body surface potential measurements, into the parameter estimation process. To achieve this, the graph based model might be extended to incorporate a phenomenological description of the action potential,

together with for example a double layer surface source model to compute the body surface potential, similar to what was done in [171]. This approach requires the additional characterization of the uncertainty in the computed body surface signal, but in principle this could be achieved using simultaneous invasive and body potential measurements.

7.3.2 Refinement of Models for Anatomy and Errors

The application of the methods presented in this dissertation to the clinic presents some additional challenges. In the work presented in this dissertation, the dimensionality of the parameter estimation problem has been intentionally kept low in order to facilitate a detailed analysis of the posterior distributions associated with the different parameters, using a grid based evaluation scheme.

Before deployment in a clinical scenario, further refinement of the model parameterization will likely be required. Generally, the optimal parameterization of a given model always represents a balance between flexibility, computational effort and access to data, all of which are very likely to vary between different clinical cases. One option for refining the parameterization is to divide the left ventricle into a number of fixed segments, such as the standardized AHA segments defined in [189], and to associate each segment with a separate local parameter or set of parameters for conduction velocity. These parameters could be interpolated between the centres of the defined regions, preventing unrealistically rapid variations in parameter values. This approach might also be employed to represent heterogeneous conduction properties of the endocardial tissue, resulting from for example purkinje conduction, by creating a separate set of segments and associated parameters comprising only the endocardial tissue. Further, the scar BZ might be described as a gradient with a separate parameter controlling the gradient shape. For added realism, the entire scar might be subdivided into regions, each with its own set of conduction parameters. A more in-depth study of optimal parameterization in the light of

clinical data on inter-individual and intra-individual variability would be highly valuable in this context.

As described in Section 5.2.2, the errors resulting from inaccuracies in modelling and measurement techniques have been approximated as normally distributed and unbiased. The measurement errors used in this study should be viewed as an accumulated description of errors from several sources, including errors from signal analysis of electrograms, errors in location of the electrodes and errors in registration of electroanatomical data to CT/MRI. Out of these, the two positional error sources are likely to have the largest impact. An interesting line of investigation is thus that of estimating the uncertainty incorporated in the measurements. As described in Section 2.5, when co-registering geometries derived from electroanatomical mapping systems with e.g. MRI-derived geometries, the nodes of the electroanatomical map are projected onto the image-derived mesh. Following the registration, a measure of how far each node has been moved is generally available to the user, giving some indication of the error which might be utilized in estimating for example the measurement related standard deviation of Eq. (5.3).

The errors in modelling electrical conduction are largely unknown, although a few quantitative studies have been performed for the bidomain model [6]. This type of error will furthermore differ between different models and parameterizations. Because of this, the description of the errors in modelling used in this dissertation should be viewed as a first approximation, useful for studying the behaviour of the inference procedure in a setting where the error distribution is approximately known. In principle, a more accurate description of the errors in modelling could obviously be derived from a large number of comparisons between computer models and clinically obtained data, but this was outside the limits of the current study.

7.3.3 Applications of the Developed Methods

The main application for the methods developed and tested in this dissertation lies within the field of parameter estimation. However, within this scope there is a great number of ways to apply it. In a clinical context, certain features of a model might be readily available from imaging modalities. These include the size and location of infarcts, available from late enhancement MRI [10], and the main myocyte orientation, available from diffusion tensor MRI [190]. However, in some situations these modalities might not be available or associated with significant uncertainty, raising the need to obtain information about these features by other means. Here, the developed parameter estimation method could be utilized. As mentioned in the discussion of Chapter 4, the graph based activation time approximations give a strong response to variations in preferential conduction direction. An interesting application would thus be to obtain endocardial and epicardial measurements located on opposite sides of a ventricular wall and use these to estimate the average orientation of the cells in the volume between the measurements. The size and location of scars could be found in a corresponding way, either by dividing the myocardium into regions and estimating parameters for each region, similar to that in [121, 125] for regions of the endocardial surface, or by having a number of moveable and resizable regions, corresponding to scar tissue, and including their size and position among the parameters to be estimated.

Bayesian inference is also closely related to the process of model selection, which could be used to select an adequate parameterization, either on a case-by-case basis or as general best practice for a class of disease conditions. It is established that the Bayesian evidence term, $P(\mathcal{D})$, is well suited to be used as an index for model selection [14, 15], since it gives information about how likely a given data set is to observe, given a specific parameterization or model. Although the computational challenge is significant, efforts in this direction also have the potential to shed some objective light on more general modelling questions, such as how to adapt the level of model complexity to the problem at hand. A full study of every conceivable model seems well

beyond the computational resources available at the time of writing, but the model selection approach can still feasibly be used to study sub-groups of models, providing answers to questions such as how to parameterize a conduction model or what cell model to choose, given access to a limited set of data modalities. This step also requires additional computational machinery to be developed in order to numerically compute the integrals involved in the evidence term.

In Chapter 6 we saw that the choice of model and strategy for data collection are intimately related. The distribution of measurements resulting from the strategy developed in the chapter (as shown in Figs. 6.4 and 6.5) relies on the particular form of the model. Thus, a re-parameterization of the model is likely to lead to a new optimized distribution of measurements. This suggests an interesting alternative: we can turn the problem around by asking for what model and parameterization the collected data leads to the lowest uncertainty about the true parameter values. One way of doing this is using adaptive models and parameterizations. Steps towards this have been taken in [157], where sub-divisions of the computational domain were made in response to new measurement data. Adapting the model to the collected data might also have a positive impact on the shape of the posterior PDF, possibly working towards avoiding multimodal distributions.

7.4 Concluding Remarks

In this dissertation, we have developed a novel phenomenological model of cardiac conduction with minimal computational requirements. Further, the developed model has enabled us to use Bayesian inference to address the challenge of probabilistic parameter estimation from electroanatomical mapping data in a simulation study. Finally, the resulting probabilistic parameter estimates have been used to study how the uncertainty in estimated parameters is affected by the number and position of the invasive measurements, and how their position can be chosen in order to reduce the uncertainty. It is hoped that the techniques, models and studies presented in

this dissertation will provide a valuable contribution to the field of computational biology and to cardiac electrophysiology research.

Bibliography

[1] José Leal, Ramón Luengo-Fernández, Alastair Gray, Sophie Petersen, and Mike Rayner. Economic burden of cardiovascular diseases in the enlarged european union. *European Heart Journal*, February 2006.

[2] J. Keener and J. Sneyd. *Mathematical Physiology*. Springer, New York, 2nd edition, 1998.

[3] Leslie Tung. *A bi-domain model for describing ischemic myocardial d-c potentials*. PhD thesis, Dept. of Electrical Engineering and Computer Science, MIT, Cambridge, Mass, 1978.

[4] Peter Lee, Christian Bollensdorff, T. Alexander Quinn, Joseph P. Wuskell, Leslie M. Loew, and Peter Kohl. Single-sensor system for spatially resolved, continuous, and multiparametric optical mapping of cardiac tissue. *Heart Rhythm*, 8(9):1482–1491, September 2011.

[5] R. Bordas, K. Gillow, Q. Lou, I.R. Efimov, D. Gavaghan, P. Kohl, V. Grau, and B. Rodriguez. Rabbit-specific ventricular model of cardiac electrophysiological function including specialized conduction system. *Progress in Biophysics and Molecular Biology*, 107(1):90–100, October 2011.

[6] A L Muzikant and C S Henriquez. Validation of three-dimensional conduction models using experimental mapping: are we getting closer? *Progress in biophysics and molecular biology*, 69(2-3):205–223, 1998. PMID: 9785939.

[7] Seong-Min Hwang, Tae Yun Kim, and Kyoung J. Lee. Complex-periodic spiral waves in

confluent cardiac cell cultures induced by localized inhomogeneities. *Proceedings of the National Academy of Sciences of the United States of America*, 102(29):10363–10368, July 2005.

[8] Jichao Zhao, Mark L. Trew, Ian J. Legrice, Bruce H. Smaill, and Andrew J. Pullan. A tissue-specific model of reentry in the right atrial appendage. *Journal of Cardiovascular Electrophysiology*, 20(6):675–684, 2009.

[9] Gernot Plank, Rebecca A. B. Burton, Patrick Hales, Martin Bishop, Tahir Mansoori, Miguel Bernabeu, Alan Garny, Anton J. Prassl, Christian Bollensdorff, Fleur Mason, Fahd Mahmood, Blanca Rodriguez, Vicente Grau, Jürgen E. Schneider, David Gavaghan, and Peter Kohl. Generation of histo-anatomically representative models of the individual heart: tools and application. *Philosophical transactions. Series A, Mathematical, physical, and engineering sciences*, 367(1896):2257–2292, June 2009. PMID: 19414455 PMCID: PMC2881535.

[10] H. Ashikaga, T. Sasano, J. Dong, M. M. Zviman, R. Evers, B. Hopenfeld, V. Castro, R. H. Helm, T. Dickfeld, S. Nazarian, J. K. Donahue, R. D. Berger, H. Calkins, M. R. Abraham, E. Marban, A. C. Lardo, E. R. McVeigh, and H. R. Halperin. Magnetic resonance based anatomical analysis of scar-related ventricular tachycardia: Implications for catheter ablation. *Circulation Research*, 101(9):939–947, 2007.

[11] W R Janowitz, A S Agatston, and Jr Viamonte, M. Comparison of serial quantitative evaluation of calcified coronary artery plaque by ultrafast computed tomography in persons with and without obstructive coronary artery disease. *The American journal of cardiology*, 68(1):1–6, July 1991. PMID: 2058541.

[12] Chattanong Yodwut, Lynn Weinert, Berthold Klas, Roberto M Lang, and Victor Mor-Avi. Effects of frame rate on three-dimensional speckle-tracking-based measurements of myocardial deformation. *Journal of the American Society of Echocardiography*, page In

press, July 2012. PMID: 22766029.

[13] D. J. Abrams, M. J. Earley, S. C. Sporton, P. M. Kistler, M. A. Gatzoulis, M. J. Mullen, J. A. Till, S. Cullen, F. Walker, M. D. Lowe, J. E. Deanfield, and R. J. Schilling. Comparison of noncontact and electroanatomic mapping to identify scar and arrhythmia late after the fontan procedure. *Circulation*, 115(13):1738–1746, March 2007.

[14] D. S. Sivia and J Skilling. *Data Analysis: A Bayesian Tutorial*. Oxford University Press, Oxford, 2nd edition, 2006.

[15] Edwin T. Jaynes. *Probability Theory: The Logic of Science*. Cambridge University Press, Cambridge, 2nd edition, 2003.

[16] Mikael Wallman, Nicolas P Smith, and Blanca Rodriguez. A comparative study of graph-based, eikonal, and monodomain simulations for the estimation of cardiac activation times. *IEEE transactions on bio-medical engineering*, 59(6):1739–1748, 2012.

[17] M. Wallman, N. Smith, and B. Rodriguez. Estimation of activation times in cardiac tissue using graph based methods. *Functional Imaging and Modeling of the Heart*, page 71–79, 2011.

[18] M. Wallman, N P Smith, and B. Rodriguez. Computational methods to reduce uncertainty in the estimation of cardiac conduction properties from electroanatomical recordings. *Submitted to Medical Image Analysis*, In review, 2013.

[19] M von Ludinghausen. Clinical anatomy of cardiac veins, vv cardiacae. *Surgical Radiologic Anatomy*, 9:159–168, 1987.

[20] Stephen N. Oesterle, Nicolaus Reifart, Eugen Hauptmann, Motoya Hayase, and Alan C. Yeung. Percutaneous in situ coronary venous arterialization report of the first human catheter-based coronary artery bypass. *Circulation*, 103(21):2539–2543, May 2001. PMID: 11382720.

[21] André G. Kléber. Gap junctions and conduction of cardiac excitation. *Heart Rhythm*, 8(12):1981–1984, December 2011.

[22] Shigeto Kanno and Jeffrey E. Saffitz. The role of myocardial gap junctions in electrical conduction and arrhythmogenesis. *Cardiovascular Pathology*, 10(4):169–177, July 2001.

[23] J. E. Saffitz, H. L. Kanter, K. G. Green, T. K. Tolley, and E. C. Beyer. Tissue-specific determinants of anisotropic conduction velocity in canine atrial and ventricular myocardium. *Circulation Research*, 74(6):1065–1070, January 1994.

[24] R. H. Hoyt, M. L. Cohen, and J. E. Saffitz. Distribution and three-dimensional structure of intercellular junctions in canine myocardium. *Circulation Research*, 64(3):563–574, January 1989.

[25] Bryan J. Caldwell, Mark L. Trew, Gregory B. Sands, Darren A. Hooks, Ian J. LeGrice, and Bruce H. Smaill. Three distinct directions of intramural activation reveal nonuniform side-to-side electrical coupling of ventricular MyocytesCLINICAL PERSPECTIVE. *Circulation: Arrhythmia and Electrophysiology*, 2(4):433–440, January 2009.

[26] J. P Keener. An eikonal-curvature equation for action potential propagation in myocardium. *Journal of Mathematical Biology*, 29(7):629–651, 1991.

[27] Michael P. Czubryt. Common threads in cardiac fibrosis, infarct scar formation, and wound healing. *Fibrogenesis & Tissue Repair*, 5(1):19, November 2012. PMID: 23114500.

[28] I. E. Willems, M. G. Havenith, J. G. De Mey, and M. J. Daemen. The alpha-smooth muscle actin-positive cells in healing human myocardial scars. *The American Journal of Pathology*, 145(4):868–875, October 1994. PMID: 7943177 PMCID: PMC1887334.

[29] Sally L. Rutherford, Mark L. Trew, Gregory B. Sands, Ian J. LeGrice, and Bruce H. Smaill. High-resolution 3-dimensional reconstruction of the infarct border ZoneNovelty

and significance impact of structural remodeling on electrical activation. *Circulation Research*, 111(3):301–311, July 2012.

[30] Michele Miragoli, Nicolò Salvarani, and Stephan Rohr. Myofibroblasts induce ectopic activity in cardiac tissue. *Circulation Research*, 101(8):755–758, December 2007.

[31] J Pu and P A Boyden. Alterations of na+ currents in myocytes from epicardial border zone of the infarcted heart. a possible ionic mechanism for reduced excitability and postrepolarization refractoriness. *Circulation research*, 81(1):110–119, July 1997. PMID: 9201034.

[32] Wen Dun, Shigeo Baba, Takuya Yagi, and Penelope A. Boyden. Dynamic remodeling of k+ and ca2+ currents in cells that survived in the epicardial border zone of canine healed infarcted heart. *American Journal of Physiology - Heart and Circulatory Physiology*, 287(3):H1046–H1054, January 2004.

[33] J A Yao, M Jiang, J S Fan, Y Y Zhou, and G N Tseng. Heterogeneous changes in k currents in rat ventricles three days after myocardial infarction. *Cardiovascular research*, 44(1):132–145, October 1999. PMID: 10615397.

[34] P. C Ursell, P. I Gardner, A. Albala, J. J Fenoglio Jr, and A. L Wit. Structural and electrophysiological changes in the epicardial border zone of canine myocardial infarcts during infarct healing. *Circulation Research*, 56(3):436–451, 1985.

[35] H. Kottkamp, B. Hugl, B. Krauss, U. Wetzel, A. Fleck, G. Schuler, and G. Hindricks. Electromagnetic versus fluoroscopic mapping of the inferior isthmus for ablation of typical atrial flutter: a prospective randomized study. *Circulation*, 102(17):2082–2086, 2000.

[36] M. J Earley, R. Showkathali, M. Alzetani, P. M Kistler, D. Gupta, D. J Abrams, J. A Horrocks, S. J Harris, S. C Sporton, and R. J Schilling. Radiofrequency ablation of arrhythmias guided by non-fluoroscopic catheter location: a prospective randomized trial. *European Heart Journal*, 27(10):1223–1229, 2006.

[37] R. J Schilling, N. S Peters, and D. W Davies. Simultaneous endocardial mapping in the human left ventricle using a noncontact catheter: comparison of contact and reconstructed electrograms during sinus rhythm. *Circulation*, 98(9):887–898, 1998.

[38] S. Y. Ho and P. Nihoyannopoulos. Anatomy, echocardiography, and normal right ventricular dimensions. *Heart*, 92(suppl 1):i2–i13, April 2006.

[39] C. C Gornick, S. W Adler, B. Pederson, J. Hauck, J. Budd, and J. Schweitzer. Validation of a new noncontact catheter system for electroanatomic mapping of left ventricular endocardium. *Circulation*, 99(6):829–835, 1999.

[40] C M Chiang, J M Jenkins, and L A DiCarlo. Digital signal processing chip implementation for detection and analysis of intracardiac electrograms. *Pacing and clinical electrophysiology: PACE*, 17(8):1373–1379, August 1994. PMID: 7971398.

[41] Shlomo A. Ben-Haim, Daniel Osadchy, Israel Schuster, Lior Gepstein, Gal Hayam, and Mark E. Josephson. Nonfluoroscopic, in vivo navigation and mapping technology. *Nature Medicine*, 2(12):1393–1395, 1996.

[42] Lior Gepstein, Gal Hayam, and Shlomo A. Ben-Haim. A novel method for nonfluoroscopic catheter-based electroanatomical mapping of the heart : In vitro and in vivo accuracy results. *Circulation*, 95(6):1611–1622, 1997.

[43] 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*, 96:2016–2021, 1997.

[44] F. H.M Wittkampf, E. F.D Wever, R. Derksen, A. A.M Wilde, H. Ramanna, R. N.W Hauer, and E. O Robles de Medina. LocaLisa: new technique for real-time 3-dimensional localization of regular intracardiac electrodes. *Circulation*, 99(10):1312–1317, 1999.

[45] A. G Brooks, L. Wilson, P. Kuklik, M. K Stiles, B. John, H. Dimitri, D. H Lau, R. L

Roberts-Thomson, C. X Wong, G. D Young, et al. Image integration using NavX fusion: initial experience and validation. *Heart Rhythm*, 5(4):526–535, 2008.

[46] P. M Kistler, K. Rajappan, S. Harris, M. J Earley, L. Richmond, S. C Sporton, and R. J Schilling. The impact of image integration on catheter ablation of atrial fibrillation using electroanatomic mapping: a prospective randomized study. *European Heart Journal*, 29(24):3029–3036, 2008.

[47] Laura Richmond, Kim Rajappan, Eric Voth, Vamsee Rangavajhala, Mark J. Earley, Glyn Thomas, Stuart Harris, Simon C. Sporton, and Richard J. Schilling. Validation of computed tomography image integration into the EnSite NavX mapping system to perform catheter ablation of atrial fibrillation. *Journal of Cardiovascular Electrophysiology*, 19(8):821–827, 2008.

[48] Laurens F. Tops, Jeroen J. Bax, Katja Zeppenfeld, Monique R.M. Jongbloed, Hildo J. Lamb, Ernst E. van der Wall, and Martin J. Schalij. Fusion of multislice computed tomography imaging with three-dimensional electroanatomic mapping to guide radiofrequency catheter ablation procedures. *Heart Rhythm*, 2(10):1076–1081, October 2005.

[49] Jun Dong, Darshan Dalal, Daniel Scherr, Aamir Cheema, Saman Nazarian, Kenneth Bilchick, Ibrahim Almasry, Alan Cheng, Charles A. Henrikson, David Spragg, Joseph E. Marine, Ronald D. Berger, and Hugh Calkins. Impact of heart rhythm status on registration accuracy of the left atrium for catheter ablation of atrial fibrillation. *Journal of Cardiovascular Electrophysiology*, 18(12):1269–1276, 2007.

[50] 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*, 33(2):161–169, March 2012.

PMID: 22119854 PMCID: PMC3305875.

[51] Hua Zhong, Joan M. Lacomis, and David Schwartzman. On the accuracy of CartoMerge for guiding posterior left atrial ablation in man. *Heart Rhythm*, 4(5):595–602, May 2007.

[52] Timo Mäkelä, Student Member, Patrick Clarysse, Outi Sipilä, Nicoleta Pauna, Quoc Cuong Pham, Toivo Katila, Isabelle E. Magnin, and La Long Axis. *A Review of Cardiac Image Registration Methods*. 2002.

[53] A. L. Hodgkin and A. F. Huxley. A quantitative description of membrane current and its application to conduction and excitation in nerve. *The Journal of Physiology*, 117(4):500–544, August 1952. PMID: 12991237 PMCID: PMC1392413.

[54] D. Noble. Cardiac action and pacemaker potentials based on the hodgkin-huxley equations. *, Published online: 05 November 1960; | doi:10.1038/188495b0*, 188(4749):495–497, November 1960.

[55] D. Noble. A modification of the Hodgkin—Huxley equations applicable to purkinje fibre action and pacemaker potentials. *The Journal of Physiology*, 160(2):317–352, February 1962. PMID: 14480151 PMCID: PMC1359535.

[56] R E McAllister, D Noble, and R W Tsien. Reconstruction of the electrical activity of cardiac purkinje fibres. *The Journal of physiology*, 251(1):1–59, September 1975. PMID: 1185607.

[57] G. W. Beeler and H. Reuter. Reconstruction of the action potential of ventricular myocardial fibres. *The Journal of Physiology*, 268(1):177–210, June 1977. PMID: 874889 PMCID: PMC1283659.

[58] D DiFrancesco and D Noble. A model of cardiac electrical activity incorporating ionic pumps and concentration changes. *Philosophical transactions of the Royal Society*

of London. Series B, Biological sciences, 307(1133):353–398, January 1985. PMID: 2578676.

[59] C H Luo and Y Rudy. A model of the ventricular cardiac action potential. depolarization, repolarization, and their interaction. *Circulation research*, 68(6):1501–1526, June 1991. PMID: 1709839.

[60] C. H. Luo and Y. Rudy. A dynamic model of the cardiac ventricular action potential. II. afterdepolarizations, triggered activity, and potentiation. *Circulation Research*, 74(6):1097–1113, January 1994.

[61] C. H. Luo and Y. Rudy. A dynamic model of the cardiac ventricular action potential. i. simulations of ionic currents and concentration changes. *Circulation Research*, 74(6):1071–1096, January 1994.

[62] J Zeng and Y Rudy. Early afterdepolarizations in cardiac myocytes: mechanism and rate dependence. *Biophysical Journal*, 68(3):949–964, March 1995. PMID: 7538806 PMCID: PMC1281819.

[63] P C Viswanathan, R M Shaw, and Y Rudy. Effects of IKr and IKs heterogeneity on action potential duration and its rate dependence: a simulation study. *Circulation*, 99(18):2466–2474, May 1999. PMID: 10318671.

[64] G M Faber and Y Rudy. Action potential and contractility changes in [na(+)](i) overloaded cardiac myocytes: a simulation study. *Biophysical journal*, 78(5):2392–2404, May 2000. PMID: 10777735.

[65] Thomas O'Hara, László Virág, András Varró, and Yoram Rudy. Simulation of the undiseased human cardiac ventricular action potential: model formulation and experimental validation. *PLoS computational biology*, 7(5):e1002061, May 2011. PMID: 21637795.

[66] Pan Li and Yoram Rudy. A model of canine purkinje cell electrophysiology and ca(2+)

cycling: rate dependence, triggered activity, and comparison to ventricular myocytes. *Circulation research*, 109(1):71–79, June 2011. PMID: 21566216.

[67] Leo Priebe and Dirk J. Beuckelmann. Simulation study of cellular electric properties in heart failure. *Circulation Research*, 82(11):1206–1223, June 1998.

[68] K. H. W. J. ten Tusscher, D. Noble, P. J. Noble, and A. V. Panfilov. A model for human ventricular tissue. *American Journal of Physiology - Heart and Circulatory Physiology*, 286(4):H1573–H1589, January 2004.

[69] K. H. W. J. ten Tusscher and A. V. Panfilov. Alternans and spiral breakup in a human ventricular tissue model. *American Journal of Physiology - Heart and Circulatory Physiology*, 291(3):H1088–H1100, January 2006.

[70] Eleonora Grandi, Francesco S. Pasqualini, and Donald M. Bers. A novel computational model of the human ventricular action potential and ca transient. *Journal of molecular and cellular cardiology*, 48(1):112, January 2010. PMID: 19835882 PMCID: PMC2813400.

[71] A. Mahajan, Y. Shiferaw, D. Sato, A. Baher, R. Olcese, L. H Xie, M. J Yang, P. S Chen, J. G Restrepo, A. Karma, A. Garfinkel, Z. Qu, and J. N. Weiss. A rabbit ventricular action potential model replicating cardiac dynamics at rapid heart rates. *Biophysical journal*, 94(2):392–410, 2008.

[72] Denis Noble, Alan Garny, and Penelope J. Noble. How the hodgkin-huxley equations inspired the cardiac physiome project. *The Journal of Physiology*, April 2012.

[73] Richard FitzHugh. Impulses and physiological states in theoretical models of nerve membrane. *Biophysical Journal*, 1(6):445–466, July 1961. PMID: 19431309 PMCID: PMC1366333.

[74] J. Nagumo, S. Arimoto, and S. Yoshizawa. An active pulse transmission line simulating

nerve axon. *Proceedings of the IRE*, 50(10):2061 –2070, October 1962.

[75] Rubin R. Aliev and Alexander V. Panfilov. A simple two-variable model of cardiac excitation. *Chaos, Solitons & Fractals*, 7(3):293–301, March 1996.

[76] Flavio Fenton and Alain Karma. Vortex dynamics in three-dimensional continuous myocardium with fiber rotation: Filament instability and fibrillation. *Chaos (Woodbury, N.Y.)*, 8(1):20–47, March 1998. PMID: 12779708.

[77] Colleen C Mitchell and David G Schaeffer. A two-current model for the dynamics of cardiac membrane. *Bulletin of mathematical biology*, 65(5):767–793, September 2003. PMID: 12909250.

[78] A Bueno-Orovio, E Cherry, and F Fenton. Minimal model for human ventricular action potentials in tissue. *Journal of Theoretical Biology*, 253(3):544–560, August 2008.

[79] Alfonso Bueno-Orovio, Ben M. Hanson, Jaswinder S. Gill, Peter Taggart, and Blanca Rodriguez. In vivo human left-to-right ventricular differences in rate adaptation transiently increase pro-arrhythmic risk following rate acceleration. *PLoS ONE*, 7(12):e52234, 2012.

[80] R. Plonsey and R. C Barr. Mathematical modeling of electrical activity of the heart. *Journal of electrocardiology*, 20(3):219–226, 1987.

[81] C S Henriquez, N Trayanova, and R Plonsey. Potential and current distributions in a cylindrical bundle of cardiac tissue. *Biophysical Journal*, 53(6):907–918, June 1988. PMID: 3395660 PMCID: PMC1330272.

[82] Natalia Trayanova and Craig S. Henriquez. Modification of a cylindrical bidomain model for cardiac tissue. *Mathematical Biosciences*, 104(1):59–72, April 1991.

[83] C S Henriquez, A L Muzikant, and C K Smoak. Anisotropy, fiber curvature, and bath loading effects on activation in thin and thick cardiac tissue preparations: simulations

in a three-dimensional bidomain model. *Journal of cardiovascular electrophysiology*, 7(5):424–444, May 1996. PMID: 8722588.

[84] N. Trayanova. Discrete versus syncytial tissue behavior in a model of cardiac stimulation. i. mathematical formulation. *Biomedical Engineering, IEEE Transactions on*, 43(12):1129 –1140, December 1996.

[85] A V Panfilov and J P Keener. Generation of reentry in anisotropic myocardium. *Journal of cardiovascular electrophysiology*, 4(4):412–421, August 1993. PMID: 8269309.

[86] J P Keener and A V Panfilov. A biophysical model for defibrillation of cardiac tissue. *Biophysical Journal*, 71(3):1335–1345, September 1996. PMID: 8874007 PMCID: PMC1233600.

[87] P. M. Nielsen, I. J. Le Grice, B. H. Smaill, and P. J. Hunter. Mathematical model of geometry and fibrous structure of the heart. *American Journal of Physiology - Heart and Circulatory Physiology*, 260(4):H1365–H1378, January 1991.

[88] Frederick J. Vetter and Andrew D. McCulloch. Three-dimensional analysis of regional cardiac function: a model of rabbit ventricular anatomy. *Progress in Biophysics and Molecular Biology*, 69(2–3):157–183, March 1998.

[89] Natalia Trayanova, James Eason, and Felipe Aguel. Computer simulations of cardiac defibrillation: a look inside the heart. *Computing and Visualization in Science*, 4(4):259–270, July 2002.

[90] Blanca Rodríguez and Natalia Trayanova. Upper limit of vulnerability in a defibrillation model of the rabbit ventricles. *Journal of electrocardiology*, 36 Suppl:51–56, 2003. PMID: 14716592.

[91] Blanca Rodríguez, Brock M. Tice, James C. Eason, Felipe Aguel, and Natalia Trayanova.

Cardiac vulnerability to electric shocks during phase 1A of acute global ischemia. *Heart Rhythm*, 1(6):695–703, December 2004.

[92] B. Rodriguez, Li Li, James C. Eason, Igor R. Efimov, and Natalia A. Trayanova. Differences between left and right ventricular chamber geometry affect cardiac vulnerability to electric shocks. *Circulation Research*, 97(2):168–175, June 2005.

[93] Mark Potse, Bruno Dube, Jacques Richer, Alain Vinet, and Ramesh M. Gulrajani. A comparison of monodomain and bidomain reaction-diffusion models for action potential propagation in the human heart. *IEEE Transactions on Biomedical Engineering*, 53(12):2425–2435, 2006.

[94] K. H. W. J. ten Tusscher, R. Hren, and A. V. Panfilov. Organization of ventricular fibrillation in the human heart. *Circulation Research*, 100(12):87–101, 2007.

[95] R. H. Keldermann, K. H. W. J. ten Tusscher, M. P. Nash, R. Hren, P. Taggart, and A. V. Panfilov. Effect of heterogeneous APD restitution on VF organization in a model of the human ventricles. *American Journal of Physiology - Heart and Circulatory Physiology*, 294(2):H764–H774, January 2008.

[96] R.H. Clayton and A.V. Panfilov. A guide to modelling cardiac electrical activity in anatomically detailed ventricles. *Progress in Biophysics and Molecular Biology*, 96(1–3):19–43, January 2008.

[97] N. A. Trayanova. Whole-heart modeling: Applications to cardiac electrophysiology and electromechanics. *Circulation Research*, 108(1):113–128, 2011.

[98] Roy Kerckhoffs, Andrew McCulloch, Jeffrey Omens, and Lawrence Mulligan. Effect of pacing site and infarct location on regional mechanics and global hemodynamics in a model based study of heart failure. In Frank Sachse and Gunnar Seemann, editors, *Functional Imaging and Modeling of the Heart*, volume 4466 of *Lecture Notes in Computer Science*, pages 350–360. Springer Berlin / Heidelberg, 2007.

[99] H.J. Arevalo, P.A. Helm, and N.A. Trayanova. Development of a model of the infarcted canine heart that predicts arrhythmia generation from specific cardiac geometry and scar distribution. In *Computers in Cardiology, 2008*, pages 497 –500, September 2008.

[100] R. C.P Kerckhoffs, A. D McCulloch, J. H Omens, and L. J Mulligan. Effects of biventricular pacing and scar size in a computational model of the failing heart with left bundle branch block. *Medical Image Analysis*, 13(2):362–369, 2009.

[101] Fijoy Vadakkumpadan, Hermenegild Arevalo, Anton J. Prassl, Junjie Chen, Ferdinand Kickinger, Peter Kohl, Gernot Plank, and Natalia Trayanova. Image-based models of cardiac structure in health and disease. *Wiley interdisciplinary reviews. Systems biology and medicine*, 2(4):489–506, July 2010. PMID: 20582162 PMCID: PMC2889712.

[102] N Wiener and A Rosenblueth. The mathematical formulation of the problem of conduction of impulses in a network of connected excitable elements, specifically in cardiac muscle. *Archivos del Instituto de Cardiología de México*, 16(3):205–265, July 1946. PMID: 20245817.

[103] Gordon K Moe, Werner C Rheinboldt, and J.A Abildskov. A computer model of atrial fibrillation. *American Heart Journal*, 67(2):200–220, February 1964.

[104] Mitsuharu Okajima, Tetsuro Fujino, Toshiji Kobayashi, and Kazuo Yamada. Computer simulation of the propagation process in excitation of the ventricles. *Circulation Research*, 23(2):203–211, January 1968.

[105] W J Eifler and R Plonsey. A cellular model for the simulation of activation in the ventricular myocardium. *Journal of electrocardiology*, 8(2):117–128, 1975. PMID: 1151192.

[106] W. J. Eifler, E. Macchi, H. J. Ritsema van Eck, B. M. Horacek, and P. M. Rautaharju. Mechanism of generation of body surface electrocardiographic p-waves in normal, middle, and lower sinus rhythms. *Circulation Research*, 48(2):168–182, January 1981.

[107] P. Colli-Franzone, L. Guerri, and S. Rovida. Wavefront propagation in an activation model of the anisotropic cardiac tissue: asymptotic analysis and numerical simulations. *Journal of Mathematical Biology*, 28(2):121–176, 1990.

[108] Toyomi Sano, Noriyuki Takayama, and Takio Shimamoto. Directional difference of conduction velocity in the cardiac ventricular syncytium studied by microelectrodes. *Circulation Research*, 7(2):262–267, January 1959.

[109] M. H. Draper and M. Mya-Tu. A comparison of the conduction velocity in cardiac tissues of various mammals. *Experimental Physiology*, 44(1):91–109, January 1959.

[110] L. Clerc. Directional differences of impulse spread in trabecular muscle from mammalian heart. *The Journal of Physiology*, 255(2):335–346, 1976.

[111] D. E. Roberts and A. M. Scher. Effect of tissue anisotropy on extracellular potential fields in canine myocardium in situ. *Circulation Research*, 50(3):342–351, January 1982.

[112] D. E. Roberts, L. T. Hersh, and A. M. Scher. Influence of cardiac fiber orientation on wavefront voltage, conduction velocity, and tissue resistivity in the dog. *Circulation Research*, 44(5):701–712, May 1979.

[113] D.A. Hooks and M.L. Trew. Construction and validation of a plunge electrode array for three-dimensional determination of conductivity in the heart. *Biomedical Engineering, IEEE Transactions on*, 55(2):626 –635, February 2008.

[114] F. M. Weber, A. Luik, C. Schilling, G. Seemann, M. W. Krueger, C. Lorenz, C. Schmitt, and O. Dossel. Conduction velocity restitution of the human Atrium—An efficient measurement protocol for clinical electrophysiological studies. *IEEE Transactions on Biomedical Engineering*, 58(9):2648–2655, September 2011.

[115] Richard J Schilling, Nicholas S Peters, Jeffrey Goldberger, Alan H Kadish, and D.Wyn Davies. Characterization of the anatomy and conduction velocities of the human right

atrial flutter circuit determined by noncontact mapping. *Journal of the American College of Cardiology*, 38(2):385–393, August 2001.

[116] Stanley Rush. Methods of measuring the resistivities of anisotropic conducting media in situ. *Journal of Research of the National Bureau of Standards - C. Engineering and Instrumentation*, 66C(3):217–222, 1962.

[117] Robert Plonsey and Roger Barr. The four-electrode resistivity technique as applied to cardiac muscle. *Biomedical Engineering, IEEE Transactions on*, BME-29(7):541 –546, July 1982.

[118] P. Le Guyader, P. Savard, and F. Trelles. Measurement of myocardial conductivities with a four-electrode technique in the frequency domain. In *Engineering in Medicine and Biology Society, 1997. Proceedings of the 19th Annual International Conference of the IEEE*, volume 6, pages 2448 –2449 vol.6, November 1997.

[119] Pierre Le Guyader, Francisco Trelles, and Pierre Savard. Extracellular measurement of anisotropic bidomain myocardial conductivities. i. theoretical analysis. *Annals of Biomedical Engineering*, 29(10):862–877, 2001.

[120] P. Steendijk, G. Mur, E.T. Van Der Velde, and J. Baan. The four-electrode resistivity technique in anisotropic media: theoretical analysis and application on myocardial tissue in vivo. *Biomedical Engineering, IEEE Transactions on*, 40(11):1138 –1148, November 1993.

[121] V. Moreau-Villeger, H. Delingette, M. Sermesant, H. Ashikaga, E. McVeigh, and N. Ayache. Building maps of local apparent conductivity of the epicardium with a 2-d electrophysiological model of the heart. *Biomedical Engineering, IEEE Transactions on*, 53(8):1457 –1466, August 2006.

[122] 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*, 16(1):201–215, January 2012.

[123] Jatin Relan, Phani Chinchapatnam, Maxime Sermesant, Kawal Rhode, Matt Ginks, Hervé Delingette, C. Aldo Rinaldi, Reza Razavi, and Nicholas Ayache. Coupled personalization of cardiac electrophysiology models for prediction of ischaemic ventricular tachycardia. *Interface Focus*, 1(3):396–407, June 2011.

[124] Oscar Camara, Ali Pashaei, Rafael Sebastian, and Alejandro F. Frangi. Personalization of fast conduction purkinje system in eikonal-based electrophysiological models with optical mapping data. In Oscar Camara, Mihaela Pop, Kawal Rhode, Maxime Sermesant, Nic Smith, and Alistair Young, editors, *Statistical Atlases and Computational Models of the Heart*, number 6364 in Lecture Notes in Computer Science, pages 281–290. Springer Berlin Heidelberg, January 2010.

[125] E. Konukoglu, J. Relan, U. Cilingir, B. H Menze, P. Chinchapatnam, A. Jadidi, H. Cochet, M. Hocini, H. Delingette, N. Ayache, et al. Efficient probabilistic model personalization integrating uncertainty on data and parameters: application to eikonal-diffusion models in cardiac electrophysiology. *Progress in Biophysics and Molecular Biology*, 2011.

[126] Jazmin Aguado-Sierra, Adarsh Krishnamurthy, Christopher Villongco, Joyce Chuang, Elliot Howard, Matthew J. Gonzales, Jeff Omens, David E. Krummen, Sanjiv Narayan, Roy C.P. Kerckhoffs, and Andrew D. McCulloch. Patient-specific modeling of dyssynchronous heart failure: A case study. *Progress in Biophysics and Molecular Biology*, 107(1):147–155, October 2011.

[127] Mihaela Pop, Maxime Sermesant, Garry Liu, Jatin Relan, Tommaso Mansi, Alan Soong, Jean-Marc Peyrat, Michael V. Truong, Paul Fefer, Elliot R. McVeigh, Herve Delingette,

Alexander J. Dick, Nicholas Ayache, and Graham A. Wright. Construction of 3D MR image-based computer models of pathologic hearts, augmented with histology and optical fluorescence imaging to characterize action potential propagation. *Medical Image Analysis*, 16(2):505–523, February 2012.

[128] Sándor M. Szilágyi, László Szilágyi, and Zoltán Benyó. A patient specific electromechanical model of the heart. *Computer Methods and Programs in Biomedicine*, 101(2):183–200, February 2011.

[129] Linwei Wang, Ken C. L. Wong, Heye Zhang, Huafeng Liu, and Pengcheng Shi. A statistical physiological-model-constrained framework for computational imaging of subject-specific volumetric cardiac electrophysiology using optical imaging and MRI data. In Oscar Camara, Mihaela Pop, Kawal Rhode, Maxime Sermesant, Nic Smith, and Alistair Young, editors, *Statistical Atlases and Computational Models of the Heart*, number 6364 in Lecture Notes in Computer Science, pages 261–269. Springer Berlin Heidelberg, January 2010.

[130] Charulatha Ramanathan, Raja N Ghanem, Ping Jia, Kyungmoo Ryu, and Yoram Rudy. Noninvasive electrocardiographic imaging for cardiac electrophysiology and arrhythmia. *Nature medicine*, 10(4):422–428, April 2004. PMID: 15034569 PMCID: PMC1950745.

[131] Xin Zhang, Indiresha Ramachandra, Zhongming Liu, Basharat Muneer, Steven M. Pogwizd, and Bin He. Noninvasive three-dimensional electrocardiographic imaging of ventricular activation sequence. *American Journal of Physiology - Heart and Circulatory Physiology*, 289(6):H2724–H2732, January 2005.

[132] Phillip S. Cuculich, Yong Wang, Bruce D. Lindsay, Mitchell N. Faddis, Richard B. Schuessler, Ralph J. Damiano, Li Li, and Yoram Rudy. Noninvasive characterization of epicardial activation in humans with diverse atrial fibrillation patterns. *Circulation*, 122(14):1364–1372, May 2010. PMID: 20855661.

[133] Linwei Wang, Heye Zhang, K.C.L. Wong, Huafeng Liu, and Pengcheng Shi. Physiological-model-constrained noninvasive reconstruction of volumetric myocardial transmembrane potentials. *IEEE Transactions on Biomedical Engineering*, 57(2):296–315, 2010.

[134] Linwei Wang, Ken C L Wong, Heye Zhang, Huafeng Liu, and Pengcheng Shi. Noninvasive computational imaging of cardiac electrophysiology for 3-d infarct. *IEEE transactions on bio-medical engineering*, 58(4):1033–1043, April 2011. PMID: 21156386.

[135] W. T. Miller and D. B. Geselowitz. Simulation studies of the electrocardiogram. II. ischemia and infarction. *Circulation Research*, 43(2):315–323, January 1978.

[136] C. M. Lloyd, J. L. Lawson, P. J Hunter, and P. F. Nielsen. The CellML model repository. *Bioinformatics*, 24(18):2122–2123, 2008.

[137] Martin J Bishop and Gernot Plank. Representing cardiac bidomain bath-loading effects by an augmented monodomain approach: application to complex ventricular models. *IEEE transactions on bio-medical engineering*, 58(4):1066–1075, April 2011. PMID: 21292591.

[138] Alfonso Bueno-Orovio, Víctor M. Pérez-García, and Flavio H. Fenton. Spectral methods for partial differential equations in irregular domains: The spectral smoothed boundary method. *SIAM J. Sci. Comput.*, 28(3):886–900, March 2006.

[139] P. Chinchapatnam, K. Rhode, M. Ginks, P. Nair, R. Razavi, S. Arridge, and M. Sermesant. Voxel based adaptive meshless method for cardiac electrophysiology simulation. In *Functional Modeling and Imaging of the Heart, Proceedings*, pages 182–190, 2009.

[140] G Fischer, B Tilg, R Modre, G J Huiskamp, J Fetzer, W Rucker, and P Wach. A bidomain model based BEM-FEM coupling formulation for anisotropic cardiac tissue. *Annals of biomedical engineering*, 28(10):1229–1243, 2000. PMID: 11144984.

[141] Andrew J. Pullan, Leo K. Cheng, and Martin L. Buist. *Mathematically Modelling the Electrical Activity of the Heart: From Cell to Body Surface and Back Again*. World Scientific, 2005.

[142] Mark L. Trew, Bruce H. Smaill, David P. Bullivant, Peter J. Hunter, and Andrew J. Pullan. A generalized finite difference method for modeling cardiac electrical activation on arbitrary, irregular computational meshes. *Mathematical Biosciences*, 198(2):169–189, December 2005.

[143] R.H. Clayton, O. Bernus, E.M. Cherry, H. Dierckx, F.H. Fenton, L. Mirabella, A.V. Panfilov, F.B. Sachse, G. Seemann, and H. Zhang. Models of cardiac tissue electrophysiology: Progress, challenges and open questions. *Progress in Biophysics and Molecular Biology*, 104(1–3):22–48, January 2011.

[144] Yves Coudière and Charles Pierre. Stability and convergence of a finite volume method for two systems of reaction-diffusion equations in electro-cardiology. *Nonlinear Analysis: Real World Applications*, 7(4):916–935, September 2006.

[145] D M Harrild and C S Henriquez. A finite volume model of cardiac propagation. *Annals of biomedical engineering*, 25(2):315–334, April 1997. PMID: 9084837.

[146] Mark Trew, Ian Le Grice, Bruce Smaill, and Andrew Pullan. A finite volume method for modeling discontinuous electrical activation in cardiac tissue. *Annals of biomedical engineering*, 33(5):590–602, May 2005. PMID: 15981860.

[147] Martin J Bishop, Gernot Plank, Rebecca A B Burton, Jürgen E Schneider, David J Gavaghan, Vicente Grau, and Peter Kohl. Development of an anatomically detailed MRI-derived rabbit ventricular model and assessment of its impact on simulations of electrophysiological function. *American Journal of Physiology. Heart and Circulatory Physiology*, 298(2):H699–718, February 2010. PMID: 19933417.

[148] John N. Tsitsilkis. Efficient algorithms for globally optimal trajectories. *IEEE Transactions on Automatic Control*, 40(9):1528–1538, September 1995.

[149] J. A Sethian. A fast marching level set method for monotonically advancing fronts. *Proceedings of the National Academy of Sciences of the United States of America*, 93(4):1591–1595, 1996.

[150] J. A Sethian and A. Vladimirsky. Fast methods for the eikonal and related Hamilton–Jacobi equations on unstructured meshes. *Proceedings of the National Academy of Sciences of the United States of America*, 97(11):5699–5703, 2000.

[151] J. A. Sethian and A. Vladimirsky. Ordered upwind methods for static hamilton-jacobi equations: theory and algorithms. *SIAM Journal on Numerical Analysis*, 41(1):325–363, 2004.

[152] M. Sermesant, Y. Coudiere, V. Moreau-Villéger, K. S. Rhode, D. L. G. Hill, and R. S. Razavi. A fast-marching approach to cardiac electrophysiology simulation for XMR interventional imaging. In *Medical Image Computing and Computer-Assisted Intervention - MICCAI 2005*, pages 607–615, 2005.

[153] Christiaan Huygens. *Traité de la Lumière*. Pieter van der Aa, Leiden, 1690.

[154] E. W. Dijkstra. A note on two problems in connexion with graphs. *Numerishe Mathematik*, 1:269–271, 1959.

[155] Ender Konukoglu, Maxime Sermesant, Olivier Clatz, Jean-Marc Peyrat, Hervé Delingette, and Nicholas Ayache. A recursive anisotropic fast marching approach to reaction diffusion equation: Application to tumor growth modeling. In *Information processing in medical imaging*, page 687–699, 2007.

[156] E. Pernod, M. Sermesant, E. Konukoglu, J. Relan, H. Delingette, and N. Ayache. A multi-

front eikonal model of cardiac electrophysiology for interactive simulation of radio-frequency ablation. *Computers {\&} Graphics*, 35:431–440, 2011.

[157] P. Chinchapatnam, K.S. Rhode, M. Ginks, C.A. Rinaldi, P. Lambiase, R. Razavi, S. Arridge, and M. Sermesant. Model-based imaging of cardiac apparent conductivity and local conduction velocity for diagnosis and planning of therapy. *Medical Imaging, IEEE Transactions on*, 27(11):1631 –1642, November 2008.

[158] Jacques Hadamard. Sur les problèmes aux dérivés partielles et leur signification physique. *Princeton University Bulletin*, 13:49–52, 1902.

[159] G. Larry Bretthorst. *Bayesian Spectrum Analysis and Parameter Estimation*. Springer, 1 edition, November 1988.

[160] H. Jeffreys. An invariant form for the prior probability in estimation problems. *Proceedings of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 186(1007):453–461, September 1946.

[161] Joakim Sundnes, Bjørn Fredrik Nielsen, Kent Andre Mardal, Xing Cai, Glenn Terje Lines, and Aslak Tveito. On the computational complexity of the bidomain and the monodomain models of electrophysiology. *Annals of Biomedical Engineering*, 34(7):1088–1097, 2006.

[162] E. J. Vigmond, R. Weber dos Santos, A. J. Prassl, M. Deo, and G. Plank. Solvers for the cardiac bidomain equations. *Progress in Biophysics and Molecular Biology*, 96(1-3):3–18, 2008.

[163] P. Pathmanathan, M. O Bernabeu, R. Bordas, J. Cooper, A. Garny, J. M Pitt-Francis, J. P Whiteley, and D. J Gavaghan. A numerical guide to the solution of the bidomain equations of cardiac electrophysiology. *Progress in Biophysics and Molecular Biology*, 102(2-3):136–155, 2010.

[164] Peter Hart, Nils J. Nilsson, and Bertram Raphael. A formal basis for the heuristic determination of minimum cost paths. *IEEE Transactions on Systems Science and Cybernetics*, 4(2):100–107, 1968.

[165] J. A. Sethian and A. Vladimirsky. Ordered upwind methods for static Hamilton–Jacobi equations. *Proceedings of the National Academy of Sciences*, 98(20):11069–11074, September 2001.

[166] J. Kim and J. P Hespanha. Discrete approximations to continuous shortest-path: Application to minimum-risk path planning for groups of UAVs. In *Proceedings of the 42nd IEEE Conference on Decision and Control*, volume 2, page 1734–1740, 2003.

[167] S. Patil and B. Ravi. Voxel-based representation, display and thickness analysis of intricate shapes. *Proceedings of the Ninth International Conference on Computer Aided Design and Computer Graphics*, pages 415–422, 2005.

[168] D. D Streeter, H. M Spotnitz, D. P Patel, J. Ross, and E. H Sonnenblick. Fiber orientation in the canine left ventricle during diastole and systole. *Circulation Research*, 24(3):339–347, 1969.

[169] J. Pitt-Francis, P. Pathmanathan, M. O Bernabeu, R. Bordas, J. Cooper, A. G Fletcher, G. R Mirams, P. Murray, J. M Osborne, A. Walter, et al. Chaste: a test-driven approach to software development for biological modelling. *Computer Physics Communications*, 180(12):2452–2471, 2009.

[170] J. M. Hare. Cardiac-resynchronization therapy for heart failure. *New England Journal of Medicine*, 346(24):1902 – 1904, 2002.

[171] Peter M. van Dam, Thom F. Oostendorp, and Adriaan Oosterom. Application of the fastest route algorithm in the interactive simulation of the effect of local ischemia on the ECG. *Medical & Biological Engineering & Computing*, 47(1):11–20, September 2008.

[172] A. Pashaei, R. Sebastian, V. Zimmerman, B. H. Bijnens, and A. F. Frangi. A meshless approach for fast estimation of electrical activation time in the ventricular wall. In *Computers in Cardiology, 2009*, pages 209–212, 2010.

[173] Hakan Oral, Aman Chugh, Eric Good, Sundar Sankaran, Stephen S. Reich, Petar Igic, Darryl Elmouchi, David Tschopp, Thomas Crawford, Sujoya Dey, Alan Wimmer, Kristina Lemola, Krit Jongnarangsin, Frank Bogun, Frank Pelosi, and Fred Morady. A tailored approach to catheter ablation of paroxysmal atrial fibrillation. *Circulation*, 113(15):1824–1831, April 2006.

[174] J. M. Morgan and V. Delgado. Lead positioning for cardiac resynchronization therapy: techniques and priorities. *Europace*, 11(Supplement 5):v22–v28, October 2009.

[175] P. V. Bayly, B. H. KenKnight, J. M. Rogers, E. E. Johnson, R. E. Ideker, and W. M. Smith. Spatial organization, predictability, and determinism in ventricular fibrillation. *Chaos: An Interdisciplinary Journal of Nonlinear Science*, 8(1):103–115, March 1998.

[176] Annamaria Carusi, Kevin Burrage, and Blanca Rodríguez. Bridging experiments, models and simulations: an integrative approach to validation in computational cardiac electrophysiology. *American Journal of Physiology - Heart and Circulatory Physiology*, 303(2):H144–H155, July 2012.

[177] M.O. Bernabeu, M. Wallman, and B. Rodriguez. Shock-induced arrhythmogenesis in a human model with patient-specific anatomy author. *Heart Rhythm*, 2010, 2010.

[178] M.O. Bernabeu, M.J. Bishop, J. Pitt-Francis, D.J. Gavaghan, V. Grau, and B. Rodriguez. High performance computer simulations for the study of biological function in 3D heart models incorporating fibre orientation and realistic geometry at para-cellular resolution. In *Computers in Cardiology, 2008*, pages 721–724, 2008.

[179] Shiro Nakahara, Roderick Tung, Rafael J. Ramirez, Jean Gima, Isaac Wiener, Aman Mahajan, Noel G. Boyle, and Kalyanam Shivkumar. Distribution of late potentials within

infarct scars assessed by ultra high-density mapping. *Heart Rhythm*, 7(12):1817–1824, 2010. In Press.

[180] Mark Girolami and Ben Calderhead. Riemann manifold langevin and hamiltonian monte carlo methods. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 73(2):123–214, 2011.

[181] Magnus Carlsson, Hakan Arheden, Charles B. Higgins, and Maythem Saeed. Magnetic resonance imaging as a potential gold standard for infarct quantification. *Journal of Electrocardiology*, 41:614–620, 2008.

[182] A. Tarantola. Popper, bayes and the inverse problem. *Nature Physics*, 2(8):492–494, 2006.

[183] Albert Tarantola. *Inverse problem theory and methods for model parameter estimation*. Society for Industrial and Applied Mathematics, Philadelphia, PA, 2005.

[184] T Cormen, C Leiserson, and R Rivest. *Introduction to Algorithms*. The MIT Press, 1990.

[185] Sir Ronald Aylmer Fisher. *Statistical methods and scientific inference*. Hafner Press, June 1973.

[186] Alistair A. Young and Alejandro F. Frangi. Computational cardiac atlases: from patient to population and back. *Experimental Physiology*, 94(5):578–596, January 2009.

[187] Carissa G. Fonseca, Michael Backhaus, David A. Bluemke, Randall D. Britten, Jae Do Chung, Brett R. Cowan, Ivo D. Dinov, J. Paul Finn, Peter J. Hunter, Alan H. Kadish, Daniel C. Lee, Joao A. C. Lima, Pau Medrano−Gracia, Kalyanam Shivkumar, Avan Suinesiaputra, Wenchao Tao, and Alistair A. Young. The cardiac atlas project—an imaging database for computational modeling and statistical atlases of the heart. *Bioinformatics*, 27(16):2288–2295, August 2011.

[188] Nic Smith, Adelaide de Vecchi, Matthew McCormick, David Nordsletten, Oscar Ca-

mara, Alejandro F. Frangi, Hervé Delingette, Maxime Sermesant, Jatin Relan, Nicholas Ayache, Martin W. Krueger, Walther H. W. Schulze, Rod Hose, Israel Valverde, Philipp Beerbaum, Cristina Staicu, Maria Siebes, Jos Spaan, Peter Hunter, Juergen Weese, Helko Lehmann, Dominique Chapelle, and Reza Rezavi. euHeart: personalized and integrated cardiac care using patient-specific cardiovascular modelling. *Interface Focus*, 1(3):349–364, June 2011. PMID: 22670205 PMCID: PMC3262448.

[189] Manuel D. Cerqueira, Neil J. Weissman, Vasken Dilsizian, Alice K. Jacobs, Sanjiv Kaul, Warren K. Laskey, Dudley J. Pennell, John A. Rumberger, Thomas Ryan, and Mario S. Verani. Standardized myocardial segmentation and nomenclature for tomographic imaging of the heart a statement for healthcare professionals from the cardiac imaging committee of the council on clinical cardiology of the american heart association. *Circulation*, 105(4):539–542, January 2002.

[190] L. Zhukov and A.H. Barr. Heart-muscle fiber reconstruction from diffusion tensor MRI. pages 597–602. IEEE, 2003.