A computational investigation of the electrocardiogram with healthy and diseased human ventricles

Louie Cardone-Noott

Linacre College

University of Oxford

A thesis submitted for the degree of

Doctor of Philosophy

Trinity Term 2016

Louie Cardone-Noott
Linacre College

Doctor of Philosophy
Trinity Term 2016

A computational investigation of the electrocardiogram with healthy and diseased human ventricles

Abstract

Cardiovascular diseases are the leading cause of death worldwide, and are estimated to kill over 17 million people each year, about 31 % of all deaths. In the clinic, the first diagnostic procedure for a suspected cardiac abnormality is often acquisition of an electrocardiogram (ECG), which measures the electrical potential of the heart at the body surface. Understanding the mechanisms underlying generation of the ECG waveforms is crucial for optimal clinical benefit. Computer simulations possess several strengths as a tool to gain this understanding, particularly in terms of human-specificity, flexibility, repeatability, and ethics.

The ventricles make up the majority of the cardiac volume and are therefore responsible for the majority of ECG waveforms. Ventricular disorders are the most life-threatening, because the ventricles are responsible for pumping blood to the body. Due to their size it has only recently become possible to perform biophysically detailed simulations of the ventricles and torso using supercomputers.

In this thesis, multiscale, mathematical models of the ventricles and torso using the Chaste software library are simulated on high performance computing systems. A description is included of the performance enhancements made in Chaste to improve resource efficiency and accelerate job turnaround, particularly in data storage and the auxiliary tasks of post-processing and data conversion. A novel model of ventricular activation is presented and parametrized using multi-modal human data, and successfully used to simulate normal and pathological QRS complexes. Similarly, repolarization gradients are imposed based on the literature and result in a variety of T waves. Finally, the developed human whole-ventricular and torso models are utilized to gain new insights into possible ionic mechanisms underlying the clinical manifestations of the early repolarization syndrome.

Overall, this thesis presents a novel framework for simulation of the human ECG using high performance computers, with possible applications in basic science and computational medicine.

Acknowledgements

Many people played a part in this thesis, in more or less obvious ways.

Thanks go to: my supervisors, Blanca Rodriguez, Alfonso Bueno, Ana Mincholé, and Kevin Burrage; my short-project supervisor Vicente Grau; the Chaste team, especially Jonathan Cooper, Gary Mirams, and Joe Pitt-Francis (double thanks for also being my departmental advisor); my colleagues and friends in the Computational Biology group; current and former students and staff of the Systems Biology Doctoral Training Centre, especially David Gavaghan and Elspeth Garman; members of Linacre College past and present, especially those who served with me as a Common Room officer; my College advisor Klaus Ebmeier; the administrators and staff at Linacre College and in the Department of Computer Science; and last but not least to my parents, family, and to Maddie.

Contents

Glossary

AB Apico-basal.
AP Action potential.
APD Action potential duration.
ARI Activation-recovery interval.
AT Activation time.

BSP Body-surface potential.

CV Conduction velocity.

ECG Electrocardiography/electrocardiogram.
ECGI Electrocardiographic imaging.
ER Early repolarization.
ERP Early repolarization pattern.
ERS Early repolarization syndrome.

HDF Hierarchical Data Format.
HOM Homogeneous.
HPC High performance computing.

IV Interventricular.

LAFB Left anterior fascicular block.
LBBB Left bundle branch block.
LPFB Left posterior fascicular block.
LV Left ventricle.

MPI Message passing interface.

ODE Ordinary differential equation.

PDE Partial differential equation.

RBBB Right bundle branch block.
RT Recovery time.
RV Right ventricle.

SA Sinoatrial node.

TM Transmural.

VTK Visualization Toolkit.

Chapter 1

Introduction

1.1 Motivation

Cardiovascular diseases are the leading cause of death worldwide, and are estimated to kill 17.3 million people each year, or about 31 % of all deaths (GBD Collaborators 2015). A common diagnostic procedure for cardiac pathologies in the clinic is electrocardiography (ECG[1]), which involves recording the electrical potential at multiple positions on the body surface. Since the heart is bioelectrically active and the body somewhat electrically conductive, information about the activity of the heart can be extracted from analysis of the potentials on the body surface. Continuous ECG monitoring is ubiquitous in many hospitals: in a recent analysis of 17 hospitals in the USA, Canada, and Hong Kong, 94 % of patients in cardiac units (4383/4678) were undergoing such monitoring (Funk et al. 2013). Understanding ECG generation in healthy and diseased conditions is therefore essential for good clinical outcomes and optimal use of this technology.

Ventricular disorders are the most life-threatening, because the ventricles are responsible for pumping blood to the body. Each heartbeat the ventricles contract

[1]The abbreviation 'ECG' is used for the procedure (electrocardiography) as well as the result (an electrocardiogram).

(activate) and relax (repolarize). Ventricular activity is reflected in the ECG in two features called the QRS complex and the T wave, as well as the intervening period called the ST segment. The QRS complex represents activation of the muscle, and the T wave its recovery; the evolution of the latter depends somewhat on the former. Abnormalities are manifested in the ECG in different ways: conduction disorders typically impact the QRS complex, whereas infarction often alters the ST segment.

Computational studies are well-suited to investigating human ECG for several reasons. Firstly, virtual models may be constructed based on human experimental data, avoiding the uncertainty in translating from animal models. Secondly, simulations have fewer ethical considerations than invasive experiments, for example those during surgery. Thirdly, virtual studies can be designed to methodically isolate and investigate small components of the system, which is impossible to do in animal or patient models.

Following the seminal work of Noble (1962) in modelling the Purkinje action potential, the first molecular-biological model of mammalian ventricular myocytes (heart cells) was published fifty years ago (Krause, Antoni, and Fleckenstein 1966). Since then there has been a proliferation of models for many species and cell types (Noble, Garny, and Noble 2012). Likewise there have been several approaches to formulate mathematical models of cardiac tissue. The model used in this thesis is called the *bidomain model*, which is based on the cable equation. The major advantage of the bidomain equations is their close link to the underlying physics; they 'represent the most general description of electrical behavior in the cardiac syncytium' (Trayanova 2011). Their main disadvantage, however, is the computational burden in solving them; it was not practical to solve them at the scale of the human heart until this decade (Bernabeu, Wallman, and Rodriguez 2010; Potse et al. 2006; Washio, Okada, and Hisada 2010), shortly before the commencement

of my DPhil. Improvements in performance and extension of capabilities are still needed before the field of computational medicine can achieve clinical impact, and this constitutes the first aim of this thesis.

Many ECG simulation studies have investigated recovery of the ventricles (for example, to evaluate the effect of drugs on the ST segment) without a thorough treatment of the preceding activation. This might partly be due to difficulties in modelling the activation system. Despite the anatomy being understood for over one hundred years (Silverman, Grove, and Upshaw 2006), complete three-dimensional human activation system data with which to construct simulations are still lacking[2]. In their absence, several strategies have been suggested, each requiring their own assumptions. The second aim of this thesis is to contribute to this endeavour by developing methods for reproducing the activation sequence of the ventricles. The approach should be closely based on a variety of human data, and be as simple and general as possible (for convenience and reusability) while still capturing a range of normal and abnormal behaviours.

The mechanisms behind the T wave have long been a source of interest for researchers, not least because T waves are usually concordant with the QRS complex (if the QRS is positive then so is the T wave, and *vice versa*). This implies that, at least in some regions, the heart muscle does *not* recover in the same way that it activates (i.e. the earliest regions to activate are not necessarily the earliest to recover) but the physiological explanation for this is unknown. Although experimental data exist from animal and human experiments, the differing methodologies and experimental conditions result in an inconsistent picture. Furthermore, as above for the activation problem, the data that are available are typically sampled from only a small number of locations in the ventricles. Recent simulation studies have therefore attempted to identify the most likely explanation by evaluating a variety

[2]There have been some promising results published recently by Forder, Hwang, and Blackband (2015) from imaging the conduction system in rabbits.

of possibilities. My third aim is an extension of this work, using a state-of-the-art modelling framework, data exclusively from human experiments, and a realistic activation sequence.

As well as simulating normal activation and repolarization, the models will be used to simulate pathological states. The activation simulations will be adjusted to represent conduction abnormalities in the activation system. Normally, the ventricles are activated in a short amount of time by a network of fast-conducting branches. If a branch of the network becomes damaged or fails to develop, the ventricles can only activate dyssynchronously, leading to characteristic waverforms in the ECG depending on the location of the block.

The repolarization models will be used as a starting point for an investigation of early repolarization pattern (ERP). This is an ECG sign involving notches or slurs in the QRS complex, the cause and mechanisms of which are unknown. There are two main hypotheses, one of which will be considered in this thesis.

1.2 Aims

1. *Improve simulation performance.* Substantial efforts have been made in recent years to improve the performance of the software used in this thesis (Chaste), in particular the numerics and parallel scaling. Several routines remain unoptimized, however, specifically relating to data storage and the solution of cardiac ODEs. My first aim is to reduce the computer time required in these functions to improve the overall efficiency in running large bidomain simulations in Chaste.

2. *Develop models of ventricular activation.* My second aim is to develop techniques for simulating ventricular activation in agreement with human data from a range of experimental modalities. The approach should be relatively

general considering the inconsistency in the data and to avoid extensive tuning and patient specificity. Normal activation sequences that generate realistic waveforms in all leads, as well as models of conduction abnormalities, will be considered.

3. *Evaluate generation of the T wave during repolarization.* Since the repolarization sequence depends somewhat on the activation sequence, my third aim builds on the preceding one in investigating the factors underlying the T wave. The effect on the T wave of combinations of gradients in action potential duration mediated by regionally heterogeneous ion channel density will be examined. The quality of simulated T waves will be determined by comparison to a set of ECGs from healthy controls.

4. *Test repolarization hypothesis in early repolarization syndrome (ERS).* The leading hypotheses for the cause of the early repolarization pattern (ERP) are (1) localized, irregular/delayed activation and (2) irregular repolarization due to localized pathological electrophysiology. A model of the latter will be informed by the activation and repolarization models described in the preceding two aims, with the addition of a region of altered electrophysiology.

1.3 Publications

L. Cardone-Noott, A. Mincholé, B. Rodriguez, and A. Bueno-Orovio (in press). "Strategies of Data Layout and Cache Writing for Input-Output Optimization in High Performance Scientific Computing: Applications to the Forward Electrocardiographic Problem". In: *IEEE Transactions on Parallel And Distributed Systems*

L. Cardone-Noott, A. Bueno-Orovio, A. Mincholé, N. Zemzemi, and B. Rodriguez (2016). "Human ventricular activation sequence and the simulation of the elec-

trocardiographic QRS complex and its variability in healthy and intraventricular block conditions". In: *Europace* 18.suppl 4, pp. iv4–iv15

B. Rodriguez, A. Carusi, N. Abi-Gerges, R. Ariga, O. Britton, G. Bub, A. Bueno-Orovio, R. A. B. Burton, V. Carapella, L. Cardone-Noott, M. J. Daniels, M. R. Davies, S. Dutta, A. Ghetti, V. Grau, S. Harmer, I. Kopljar, P. Lambiase, H. R. Lu, A. Lyon, A. Minchole, A. Muszkiewicz, J. Oster, M. Paci, E. Passini, S. Severi, P. Taggart, A. Tinker, J.-P. Valentin, A. Varro, M. Wallman, and X. Zhou (2015). "Human-based approaches to pharmacology and cardiology: an interdisciplinary and intersectorial workshop". In: *Europace*, euv320

L. Cardone-Noott, A. Bueno-Orovio, A. Mincholé, K. Burrage, M. Wallman, N. Zemzemi, E. Dall'Armellina, and B. Rodriguez (2014). "A Computational Investigation into the Effect of Infarction on Clinical Human Electrophysiology Biomarkers". In: *Computing in Cardiology*. Vol. 41, pp. 673–676

1.4 Thesis outline

This thesis is organized as follows.

Chapter 2 presents background information including: the anatomy, physiology, and function of the heart; the cardiac action potential; and history and fundamentals of the ECG.

Chapter 3 contains the methods and materials used in the study. This includes a summary of the simulation framework, an overview of cardiac cell models and details of the specific model used, an introduction to the bidomain equations and finite element method, a description of the tetrahedral mesh used, a description of the algorithm used to generate the tissue structure, a section on measured tissue conductivities and the values used in the study, and a summary of the Chaste software library used for simulation.

Chapter 4 describes performance improvements in Chaste on high-performance computing platforms. The time required in several auxiliary (not the main solver) tasks was reduced, including mesh reading and writing, accessing the main results file for a simulation, and performing post-processing and data conversion for visualization. Also presented are details of enabling routine use of a specialized, high-performance, differential-equation solver, and an adaptive time-stepper for long simulations. Several benchmarks are used to illustrate the efficiency improvements.

Chapter 5 is concerned with the activation sequence of the ventricles, specifically a model for initiating a beat without explicit inclusion of the activation-system branches and fibres. Multi-modal human data from the literature are used to guide model construction. The relationships between model parameters and ECG features are investigated, in normal conditions and with simulated conduction abnormalities.

Chapter 6 presents modelling and simulations of ventricular repolarization following normal activation as shown in the preceding chapter. Regional heterogeneities in cell properties are used to modify the repolarization times, and the quality of the resultant ECG features ascertained by comparison to a dataset from healthy control patients. The abnormalities generated by some configurations are shown, as well as one combination of heterogeneities that produce a normal ECG.

Chapter 7 draws on the models of the preceding chapters to investigate a proposed mechanism in an ECG sign known as early repolarization pattern. Two modifications to the cell properties are invoked in a localized region of the ventricles, and their impact on the ECG quantified.

Chapter 2

The human heart and the body surface electrogram

The purpose of this chapter is to provide the physiological and clinical foundations for the rest of this thesis. Section 2.1 contains descriptions of the physiology of the human heart, the activation system, cardiomyocytes, the action potential, and regional differences in the ventricles. Section 2.2 provides an introduction to electrocardiography.

2.1 The human heart

2.1.1 Physiology

The heart is a variable-speed biomechanical pump that regulates the movement of blood throughout the body in order to transport biological resources and waste products. The human heart, like all mammalian hearts, consists of four chambers (fig. 2.1). These chambers are efficiently filled and emptied through an organized series of contraction and relaxation of the walls.

The heart is made up mainly of *myocardium*, bordered by thin internal and ex-

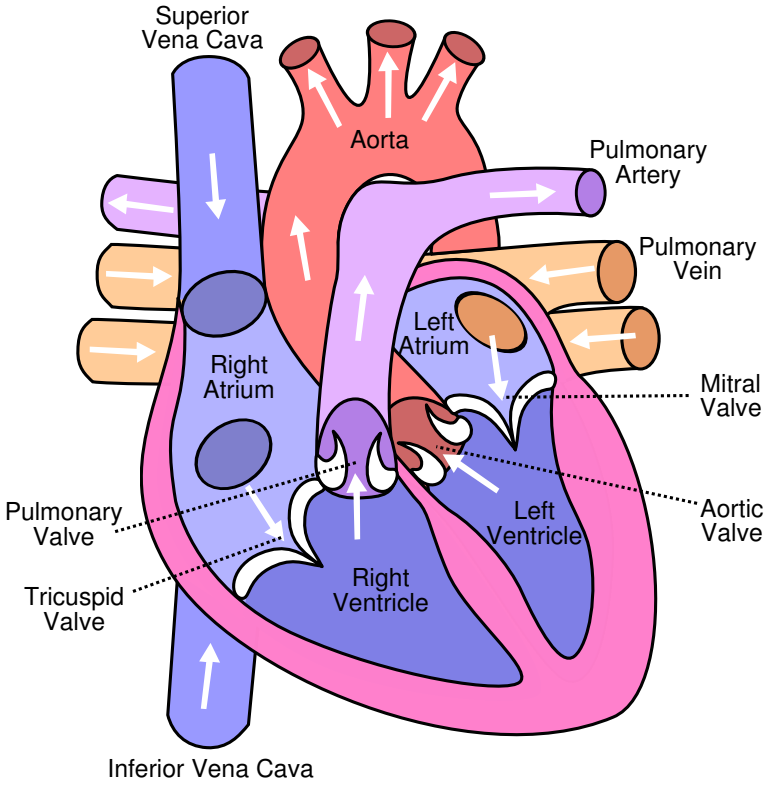

Figure 2.1: A cartoon of the human heart, with direction of blood flow indicated by arrows. Deoxygenated blood enters through the vena cava and goes to the lungs through the pulmonary artery. Oxygenated blood returns from the lungs through the pulmonary vein, and goes to the body through the aorta. A number of valves act to maintain this direction of flow. "Diagram of the human heart" by Eric Pierce is licensed under CC BY-SA 3.0. `https://commons.wikimedia.org/wiki/File:Diagram_of_the_human_heart_(cropped).svg`

ternal layers called *endocardium* and *epicardium*, respectively. Cardiac myocytes (muscle cells) make up the majority of the myocardium by mass, the remainder (about 70 % by number) being vascular or connective tissue. Myocytes are coupled to one another at *intercalated discs* which have low electrical resistance, hence they form an electrically continuous structure.

In healthy hearts, the pacing and initiation of a heartbeat occurs in the *sinoatrial node*, which transmits a bio-electrical voltage signal called an *action potential* (AP) to the muscular heart walls via a network of specialized cells comprising various nodes, bundles, and branches, described below (fig. 2.2). The atria contract first, during which time the signal from the sinoatrial node is delayed at the *atrioventric-*

ular node, located between the atria and ventricles. It then passes along the *bundle of His* to the left and right *bundle branches*, which further split into several *fascicles* that travel through the septum to the endocardia of the left and right ventricles. Finally, a network of sub-endocardial *Purkinje fibres* covering much of the endocardial surfaces leads to relatively synchronous activation of the myocardium via a large number of Purkinje-ventricular junctions. For a myocyte, the sudden change in voltage from the activation system (or its neighbouring myocytes) causes it to undergo its own action potential, altering the concentration of ions within and ultimately causing it to contract and thus contribute to the heart's function.

2.1.2 The cardiomyocyte action potential

Cardiomyocytes are electrically excitable due to a variety of voltage-sensitive pores in the membrane. The *resting membrane potential* (measured inside relative to the outside) is negative, and is maintained by *ion transporters* (pumps and exchangers) as well as the inward rectifier current (I_{K1}). When the cell is sufficiently depolarized (i.e. the cell interior made less negative compared to the exterior) the myocyte exhibits an AP. The AP is due to a complex series of openings and closings of *voltage-gated ion channels* resulting in a movement of ions in both directions across the cell membrane and thus a characteristic change in transmembrane potential (fig. 2.3). Since it is a vitally important process to the organism the sequence of events which produce an AP are very robust, such that the AP is an all-or-nothing event and may often be recovered even with disruption to multiple ion channels.

Like all ion channels found throughout the body's cells, those found in cardiomycocytes are membrane proteins composed of one or more types of subunit that permit the selective movement of ions across the cell membrane. This movement of ions is a movement of charge and is therefore called an *ionic current*. The most important ionic currents in the present context are sodium, potassium, cal-

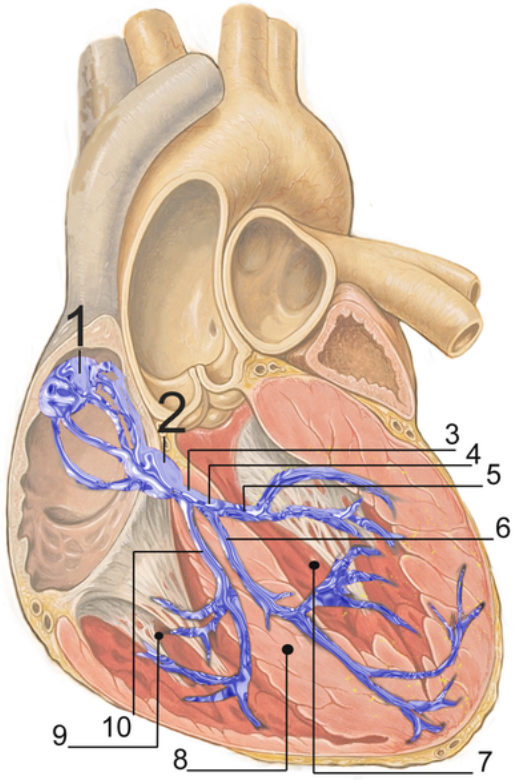

Figure 2.2: A diagram of the conduction system in the human heart. Normal heart rhythm is initiated by the sinoatrial node (1) which activates the atria. The signal is delayed at the atrioventricular node (2) before passing along the bundle of His (3) to the left and right bundle branches (4 and 10). The remainder of the network consists of fascicles (5 and 6) which travel through the septum (8) to the endocardium of the left (7) and right ventricles. A network of Purkinje fibres then cover much of the endocardial surface (9). "Electrical conduction system of the heart" by Patrick J. Lynch and C. Carl Jaffe is licensed under CC BY 2.5. `https://commons.wikimedia.org/wiki/File:RLS_12blauLeg.png`

cium, and chloride currents. The sodium current is an inward current (sodium cations enter the cell) that opens briefly to generate the fast initial depolarization of an AP. Several potassium currents operate in the opposite direction during the AP, variously to maintain the resting membrane potential or regulate action potential duration (APD). Two main inward calcium currents are found in the heart: the larger of which (L-type, I_{CaL}) is important in maintaining the plateau phase and causing contraction; the other (T-type) is more important in the SA node. Finally, chloride currents are repolarizing and contribute mainly to volume regulation.

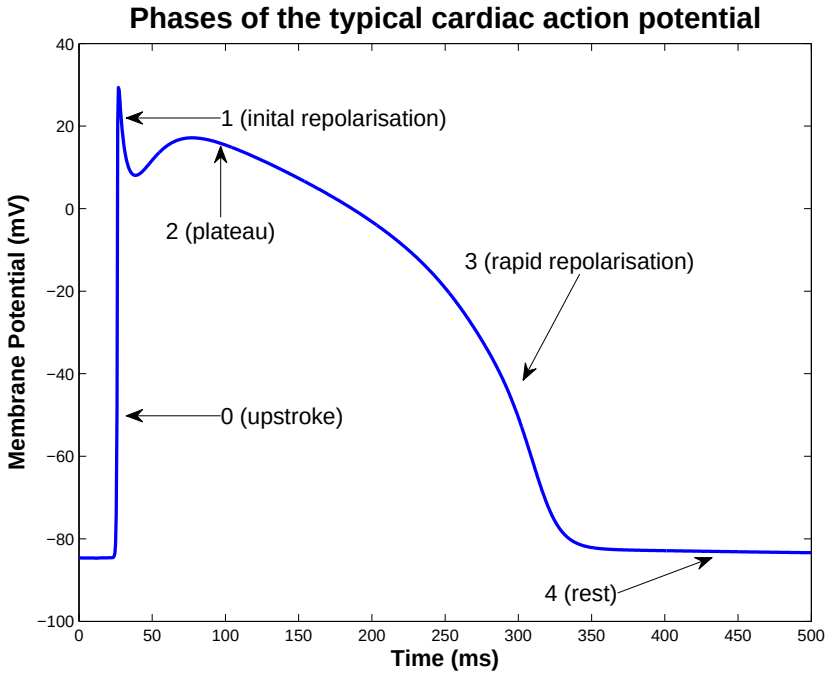

Figure 2.3: A schematic of the phases of the human cardiac action potential. The notch immediately after the spike (between phase 1 and 2) is characteristic of a subepicardial cell. Figure courtesy of Doug Bruce.

2.1.3 Tissue microstructure

Cardiomyocytes are approximately cylindrical muscle cells of length 100 µm and diameter 25 µm. Pores in the membrane called gap junctions connect neighbouring cells together, allowing ions to travel directly between their sarcoplasms. The majority of gap junctions are found at the 'ends' of the cells at features called intercalated disks, with a smaller density found on the 'sides'. This difference in coupling leads to preferential ion movement along the long-axis of neighbouring cells, which leads us to term this the 'fibre direction'.

Furthermore, the fibres are arranged in 'sheets', with impaired ion transport between sheets (Greenbaum et al. 1981; LeGrice et al. 1995; Sands et al. 2006; Young et al. 1998). This configuration means that at every point in the tissue the local conductivity may be described as a set of orthogonal vectors, with the largest member describing conductivity along the fibre direction, the intermediate one

conductivity within the sheet, and the smallest one conductivity between sheets. It has recently been suggested that fully orthotropic conductivity is important in reproducing accurate conduction dynamics (Caldwell et al. 2009).

2.1.4 Electrophysiological heterogeneities

Mammalian hearts are electrophysiologically inhomogeneous (Boukens et al. 2009; Boukens and Christoffels 2012). Experiments in the late 1980s led to the discovery that myocytes near the epicardial (epi) surface are not identical to those near the endocardial (endo) surface, but instead have different ion channel densities, conductivities, and rate behaviours. This is known as transmural heterogeneity. In the 1990s, researchers identified a third 'cell type' in the deep tissue of the ventricles, and named it the M-cell (Antzelevitch et al. 1991; Sicouri and Antzelevitch 1991) as it was isolated from the midmyocardium. As shown by Liu and Antzelevitch (1995, fig. 2), the majority of samples from the midmyocardium behave similarly to epicardial and endocardial cells, but a minority display very extended APDs.

Data on the relative abundance of each cell type has been widely cited from Drouin et al. (1995) as approximately 10 % epicardial and 30 % M-cell by mass in the left ventricle (the remainder being endocardial, or 'transitional' between M and endocardial). More recent evidence, however, suggests that M-cells may not make up a contiguous layer between Endo and Epi cells, but rather form 'islands' (Akar et al. 2002; Glukhov et al. 2010). Whether the regions of long APD attributed to M-cells in wedge preparations are responsible for any transmural heterogeneity in intact hearts is controversial (Anyukhovsky et al. 1999; Janse et al. 2012).

There have been many studies of electrophysiological heterogeneity in large mammalian hearts. Typically they report differences in: APD or recovery time (RT) (see Opthof et al. (2016)) or ion channel expression (see Boukens et al. (2009)); transmurally, between apical and basal regions, or between the ventricles. Below

is a brief review of each of these.

Transmural

The APD of endocardial tissue is longer than epicardial in human (Näbauer et al. 1996; Szabó et al. 2005), canine (Burgess et al. 1972; Janse et al. 2005; Obreztchikova et al. 2006; Sekiya et al. 1984; Szabó et al. 2005), and guinea pig (Bryant et al. 1998). Other evidence in guinea pig suggests little variation (Watanabe, Rautaharju, and McDonald 1985), and an *in vivo* human study found no variation in APD transmurally (Taggart et al. 2001).

Different electrophysiology in samples taken from the endocardium and epicardium has been reported. The transient outward current (I_{to}) has been measured with higher density in epicardial cells in human (Li et al. 1998; Näbauer et al. 1996; Wettwer et al. 1994), canine (Litovsky and Antzelevitch 1988), and feline (Furukawa et al. 1990). Second, the slow delayed rectifier potassium current (I_{Ks}) has been measured to be larger in the epicardium in canine (Liu and Antzelevitch 1995; Zygmunt et al. 2001). Third, the sodium-calcium exchanger current (I_{NaCa}) has been measured as largest in the epicardium in canine (Zygmunt, Goodrow, and Antzelevitch 2000). Fourth, the rapid delayed rectifier potassium current (I_{Kr}) was measured to be larger in the epicardium than endocardium in canine (Obreztchikova et al. 2006) and ferret (Brahmajothi et al. 1997).

Recent experimental data suggest there may be more currents involved. For example, the 2011 O'Hara et al. human ventricular model goes further by altering the following depending on cell type: potassium channel currents (I_{Kr}, I_{K1}, I_{Kb}), late sodium (I_{NaL}) and calcium (I_{CaL}) channel currents, the sodium-potassium pump, and calcium transport (SERCA and RyR).

Interventricular

APD is longer in the left than the right ventricle in human (Bueno-Orovio et al. 2012; Ramanathan et al. 2006), canine (Volders et al. 1999; Vos et al. 1998), and rat (Watanabe et al. 1983). Results from guinea pig show longer epicardial APD in the right ventricle, but longer endocardial in the left (Watanabe, Rautaharju, and McDonald 1985), while computer simulations show a longer APD in the right ventricle than the left (Samie et al. 2001).

I_{to} is larger in the right than left ventricle in canine when measured in the epicardium (Di Diego, Sun, and Antzelevitch 1996) or midmyocardium (Volders et al. 1999). However, values reported in human show no significant interventricular difference, not allowing for differing methodologies (Li et al. 1998; Wettwer et al. 1994).

I_{Ks} is also larger in the right ventricle than the left in canine (Volders et al. 1999).

Apico-basal

Studies on apico-basal gradients are contradictory. Studies have found longer APD at the apex than the base in canine (Bauer et al. 2002; Burgess et al. 1972; Sekiya et al. 1984) and rabbit (Cheng et al. 1999). Others have found longer APD in the base than the apex in human and canine (Szentadrassy et al. 2005), rat (Watanabe et al. 1983) and guinea pig (Watanabe, Rautaharju, and McDonald 1985). Another reports longer repolarisation time (activation time + activation recovery interval) at the apex than the base (Janse et al. 2005).

In terms of currents, I_{Ks} is larger in the apex than the base in canine (Bauer et al. 2002; Szentadrassy et al. 2005). In rabbit however, I_{Ks} is larger in the base, but I_{Kr} larger in the apex, while total potassium current is larger in the base than the apex (Cheng et al. 1999).

2.2 Electrocardiography

2.2.1 Principles of ECG

The standard 12 lead ECG is a set of twelve time-varying voltages constructed from electric potentials sampled at ten electrodes on the body. Four electrodes are placed on the limbs (either on the wrists/ankles or near the shoulders/hips) and six on the chest surface. The latter are called the precordial electrodes and should be positioned as in table 2.1, illustrated in cartoon form in fig. 2.5. The right leg electrode serves as a reference electrode (ground).

A normal ECG from a healthy volunteer is shown in fig. 2.4. The upper section shows the 12 leads in the conventional layout with limb leads in the first column, augmented limb leads in the second column, and precordial leads in the third and fourth columns. The lower section shows two leads (II and V5 in this case) for a longer duration. The conventional grid is used: one small square represents 40 ms horizontally and 0.1 mV vertically; and is 1 mm square.

Electrode	Position
V1	Fourth intercostal space at the right sternal edge
V2	Fourth intercostal space at the left sternal edge
V3	Midway between V2 and V4
V4	Fifth intercostal space in the mid-clavicular line
V5	Left anterior axillary line at same horizontal level as V4
V6	Left mid-axillary line at same horizontal level as V4 and V5

Table 2.1: Precordial electrode placement guidelines (Society for Cardiological Science and Technology 2010).

Each of the leads will be defined below. In the following, the potential at the left arm and right arm electrodes are written LA and RA, respectively. Similarly, LL and RL refer to the left and right leg electrodes, respectively.

Einthoven, Fahr, and de Waart (1950) proposed the first three leads in 1908. Leads I, II, and III each result from difference between two of the limb signals, and

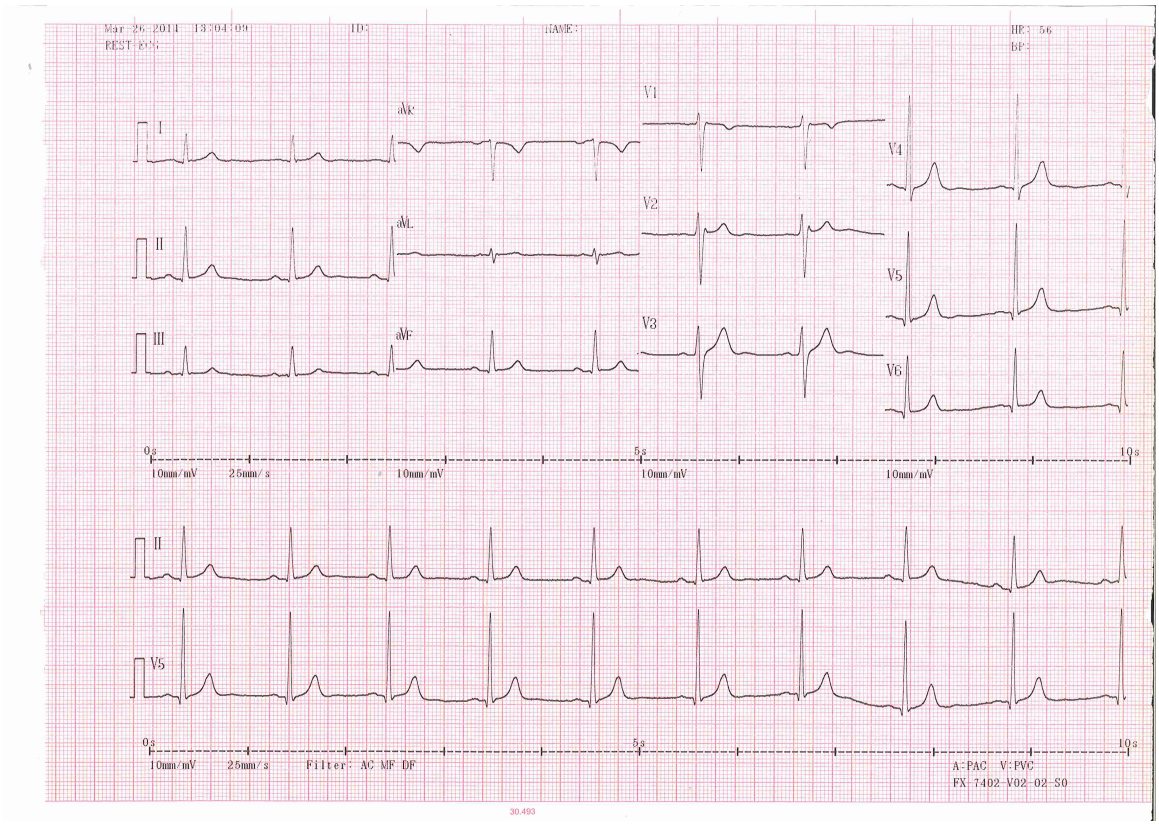

Figure 2.4: An ECG from a healthy 24-year-old male volunteer. The upper half shows two beats for each of the 12 leads arranged in a 3 by 4 grid. The lower half shows 9 beats for leads II and V5 only. The small grid spacing is 1 mm, the horizontal scale is 25 mm s^{-1} (making one small square equal to 40 ms), and the vertical scale is 10 mm mV^{-1} (making one small square equal to 0.1 mV).

are therefore called bipolar leads:

$$\mathrm{I} = \mathrm{LA} - \mathrm{RA},$$

$$\mathrm{II} = \mathrm{LL} - \mathrm{RA},$$

$$\mathrm{III} = \mathrm{LL} - \mathrm{LA}.$$

Kirchhoff's second law states that the sum of voltages in a closed loop is zero:

$$(\mathrm{LA} - \mathrm{RA}) + (\mathrm{RA} - \mathrm{LL}) + (\mathrm{LL} - \mathrm{LA}) = 0 \tag{2.1}$$

$$\implies \mathrm{I} - \mathrm{II} + \mathrm{III} = 0 \tag{2.2}$$

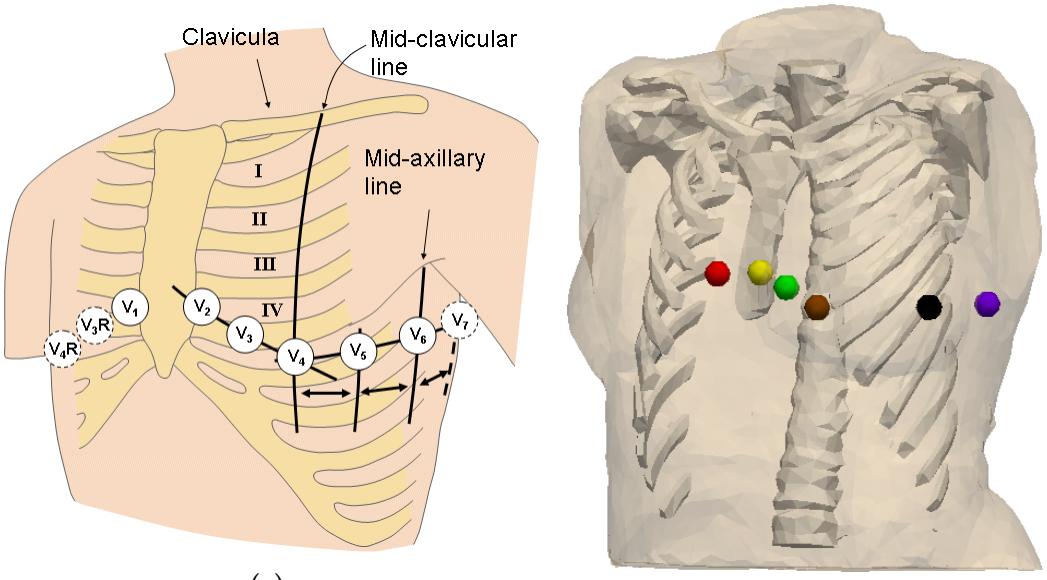

(a) (b)

Figure 2.5: (a) Illustration of the locations of the precordial electrodes. The solid circles indicate the standard V1 to V6 electrodes, the dashed circles indicate optional extensions. Reproduced from Plonsey and Malmivuo (1995, fig. 15.8) under licence (`http://www.bem.fi/book/`). (b) The V1–V6 electrode positions on the torso mesh (section 3.3.2), positioned relative to the intercostal spaces (table 2.1).

so we see that only two of the three leads are independent. Next, Goldberger (1942) suggested the three 'augmented' limb leads

$$\begin{aligned}
\mathrm{aVR} &= \mathrm{RA} - \tfrac{1}{2}\left(\mathrm{LA} + \mathrm{LL}\right) = \tfrac{3}{2}\left(\mathrm{RA} - \mathrm{V}_W\right), \\
\mathrm{aVL} &= \mathrm{LA} - \tfrac{1}{2}\left(\mathrm{RA} + \mathrm{LL}\right) = \tfrac{3}{2}\left(\mathrm{LA} - \mathrm{V}_W\right), \\
\mathrm{aVF} &= \mathrm{LL} - \tfrac{1}{2}\left(\mathrm{RA} + \mathrm{LA}\right) = \tfrac{3}{2}\left(\mathrm{LL} - \mathrm{V}_W\right),
\end{aligned}$$

where V_W is a reference potential called the Wilson central terminal (Wilson et al. 1934), which is the average of the (three non-earthed) limb leads

$$\mathrm{V}_W = \frac{\mathrm{RA} + \mathrm{LA} + \mathrm{LL}}{3}.$$

Finally, in 1944 Wilson et al. proposed the six unipolar precordial leads. These are the differences between potentials recorded at each of the six chest electrodes and V_W, which acts as a reference approximating the potential at infinity. The augmented and precordial leads are the difference between a recording electrode and a combination of other electrodes and are therefore called unipolar leads. The orientations of these lead vectors are shown in fig. 2.6.

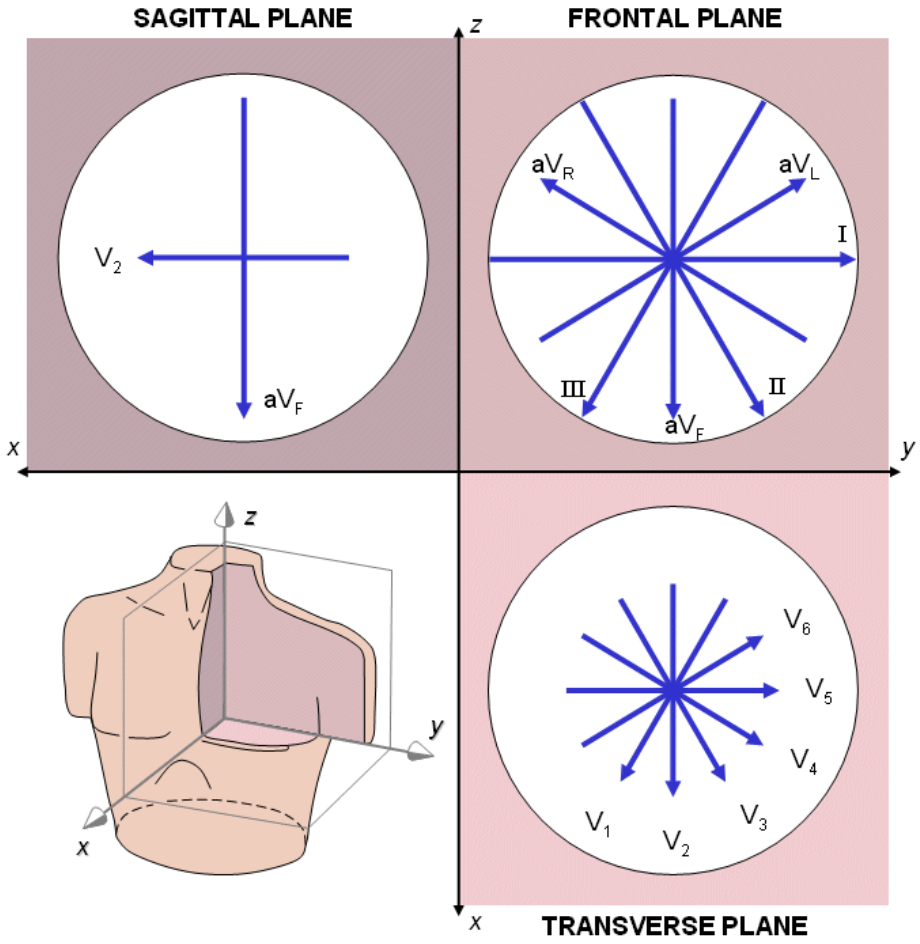

Figure 2.6: An illustration of the twelve lead vectors for a homogeneous conductor. Reproduced from Plonsey and Malmivuo (1995, fig. 15.9) under licence (`http://www.bem.fi/book/`).

Before describing any specific ECG features it is worth noting how deflections are formed in the ECG signals. In the context of cardiac tissue (which has a negative resting potential), given a 'positive' and a 'negative' electrode:

- A wave of **depolarization** moving **towards the positive** electrode and away from the negative electrode causes a **positive** deflection.
- A wave of **depolarization** moving **towards the negative** electrode and away from the positive electrode causes a **negative** deflection.
- A wave of **repolarization** moving **towards the positive** electrode and away from the negative electrode causes a **negative** deflection.
- A wave of **repolarization** moving **towards the negative** electrode and away from the positive electrode causes a **positive** deflection.

In terms of the limb leads: in lead I for example, the positive electrode would be LA and the negative electrode RA. In terms of the augmented leads: in lead aVR for example, the positive electrode would be RA and the negative electrode the average of LA and LL. In terms of the precordial leads: in V1 for example, the positive electrode would be the V1 electrode and the negative electrode would be V_W. Note that in addition, a near-zero or small and biphasic feature is formed when net activity is perpendicular to the vector between electrodes.

2.2.2 Interpreting an ECG

A normal heartbeat consists of depolarization and repolarization waves in the atria and ventricles, resulting in deflections or *waveforms* in most of the leads. A schematic of the usual waves is shown in fig. 2.7, but note that the amplitudes and polarities vary from lead to lead. First, the P wave is due to atrial depolarization. Second, the QRS complex is due to ventricular depolarization (activation). Third, the T wave is due to ventricular repolarization (recovery). Fourth, the U wave is only sometimes seen, possibly due to delayed papillary muscle repolarization.

The QRS complex contains a combination of Q, R, and S waves, but not necessarily one of each (zero or more than one are possible). A Q wave is any downward

deflection following a P wave. An R wave is an upward deflection following a P or Q wave. An S wave is a downward deflection following an R wave.

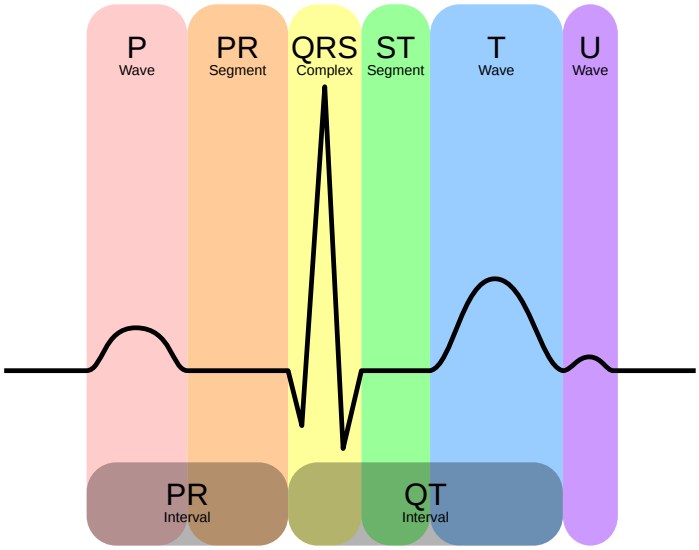

Figure 2.7: A cartoon of an ECG lead recording with the waves highlighted separately. "Normal ECG/EKG complex with labels" by Hank van Helvete and Hazmat2 is licensed under CC BY-SA 3.0. `https://commons.wikimedia.org/wiki/File:EKG_Complex_en.svg`

A complete guide to ECG interpretation is unnecessary here, and for the purposes of this chapter it will suffice to highlight only a few of its aspects. As this thesis is concerned only with the ventricles we can disregard the P wave and focus instead on the QRS complexes and T waves.

QRS complex

As stated above, the QRS complex contains information about ventricular activation. It is normally 60 ms to 120 ms in duration in healthy conditions, although durations over 100 ms may be abnormal.

The gradient of electric potential throughout the heart and body is a vector field. At a given moment during activation most of these vectors in the ventricles cancel out because of the geometry and the activation of distant regions of the ventricles in a short space of time. The sum of these vectors is called the *mean*

instantaneous QRS vector, and the sum of all instantaneous vectors throughout time is the *mean QRS vector*.

Einthoven used his bipolar limb leads and a model called the *Einthoven triangle* to estimate the QRS axis (fig. 2.8). In this model the edges of an equilateral triangle represent the vectors between the non-earthed electrodes, and the centre of the triangle the position of the heart. By plotting the net deflection of each lead onto the corresponding edge, the mean QRS vector is approximately oriented from the centre to the intersection between the three perpendiculars (fig. 2.9). The perpendiculars intersect at a single point due to the linear dependency of the three leads (eq. (2.2)). By convention, an axis of 0° is aligned with lead I and the angle increases in a clockwise direction (fig. 2.10). An axis of between −30° and 110° is considered normal according to population studies (Katz 2006, p. 454). An example of calculating the QRS axis angle using Einthoven's triangle is shown in fig. 2.11 for a control subject. Note the small amount of error in the intersection point.

Figure 2.8: Cartoon of the Einthoven triangle. The vertices represent the three non-earthed electrodes, and the edges the limb leads, with the heart in the centre. Reproduced from Katz (2006, fig. 15-27) (© 2006 Lippincott Williams & Wilkins).

A simple way to estimate the QRS axis, to guide interpretation of the results in chapter 5, is to examine leads I and aVF, because in Einthoven's triangle these are orthogonal and bound the 0° to 90° quadrant. Below are some illustrative examples of lead magnitudes and QRS axes.

Permission to make the image originally presented here freely available via ORA was not granted by Wolters Kluwer Health, Inc.

Figure 2.9: An example of using the Einthoven triangle to estimate the QRS axis. The origin of the QRS axis lies in the centre of the triangle, and the origin of each 'lead vector' at the midpoint of the respective triangle edge. A vector is drawn on each edge of the triangle with (relative) magnitude equal to the net QRS amplitude of the respective lead (the difference in the number positive and negative squares, written above each lead). The intersection of the perpendiculars of each lead vector defines the end of the QRS vector. In this example from a normal ECG, the QRS axis is approximately 70°. Reproduced from Katz (2006, fig. 15-28) (© 2006 Lippincott Williams & Wilkins).

Permission to make the image originally presented here freely available via ORA was not granted by Wolters Kluwer Health, Inc.

Figure 2.10: A schematic of QRS axis angles. The normal range is from −30° to 110°, where 0° is aligned with lead I and the angle increases in a clockwise direction. Reproduced from Katz (2006, fig. 15-30) (© 2006 Lippincott Williams & Wilkins).

1. Leads I and aVF are both positive: the QRS axis is in the lower-left quadrant, i.e. 0° to 90° (normal).

2. Lead I is negative and aVF is positive: the QRS axis is in the lower-right quadrant, i.e. 90° to 180°. If the tallest R wave is in lead aVF the QRS axis is <105° (normal). If the tallest R wave is in lead III or aVR the axis is >110° (right-deviated).

3. Lead I is positive and aVF is negative: the QRS axis is in the upper-left quadrant, i.e. 0° to −90°. If the largest R wave is in lead I the axis is >−15° (normal). If the largest R wave is in aVL and the net polarity of lead II is negative the axis is <−30° (left-deviated).

4. Leads I and aVF are negative: the QRS axis is −90° to −180° (extreme axis).

In the **precordial leads**, lead V1 is usually negatively deflected (QS type), leads V4 to V6 are positively deflected, and in between the R wave increases (called *R wave progression*).

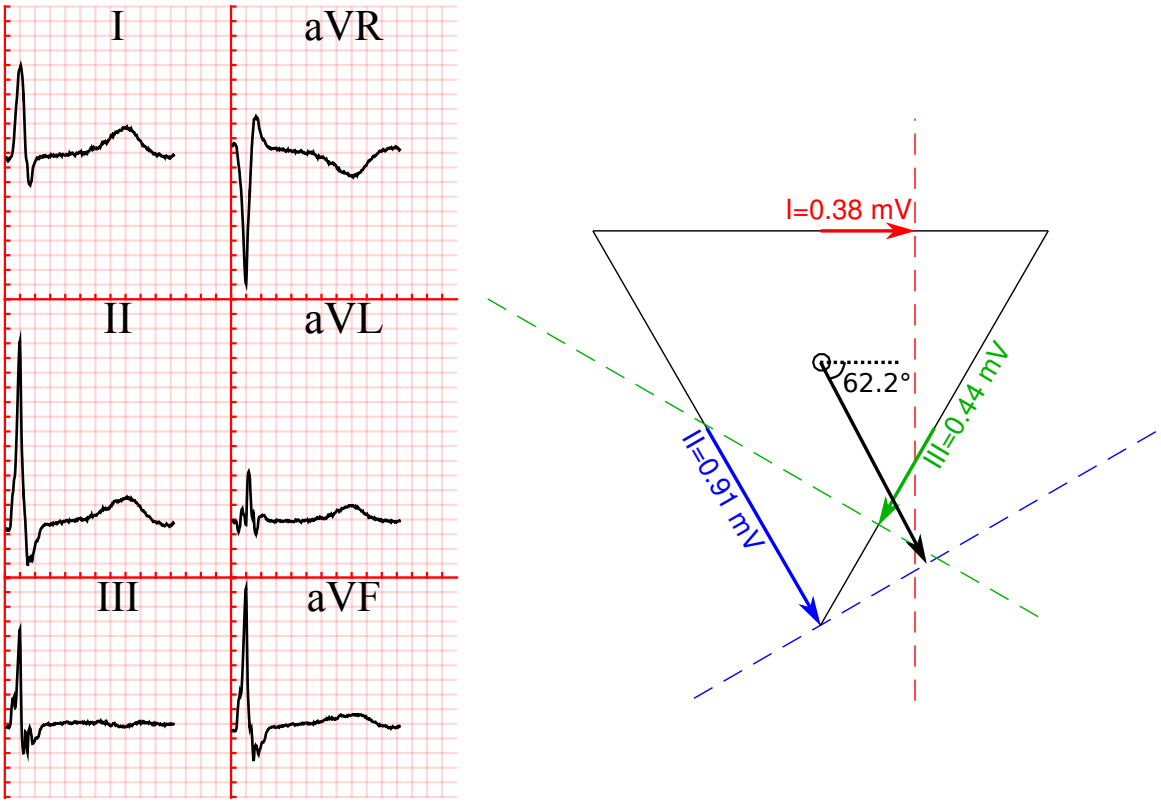

Figure 2.11: An example of calculating the QRS axis using Einthoven's triangle. Left: the limb leads and augmented limb leads for a control patient (subject #246, Bousseljot, Kreiseler, and Schnabel (2009)). Right: the net deflection of each of the leads I, II, and III are plotted on the corresponding sides of an equilateral triangle. The approximate intersection of the perpendiculars triangulates the QRS axis angle, in this case 62°.

T wave

The T wave is a sign of repolarization of the myocardium, which occurs asynchronously throughout the ventricles. Unlike the QRS complex, which is narrow and tall due to the rapid activation of the ventricles, the T wave is smooth and broad. The T wave can be described by its shape (symmetry, skew), amplitude, slopes, and other metrics.

An interesting property of the T wave in healthy patients is its concordance with the QRS complex, i.e. a positive QRS is usually followed by a positive T wave and *vice versa*. Recall that a wave of activation followed by a wave of repolarization

produce opposite deflections (section 2.2.1). There are many possible explanations for this, including: (1) long action potential duration in Purkinje fibre (Katz 2006, p. 455); (2) a mid-wall layer of cells called *M-cells* (Antzelevitch 2001) with long APD, though their existence in intact hearts is controversial (Janse et al. 2012; Taggart et al. 2001); (3) regional APD differences due e.g. to differences in ion channel expression (Boukens et al. 2009); (4) shortening of APD along activation path due to inhomogeneity in the tissue microstructure (Bueno-Orovio et al. 2014; Hanson et al. 2009).

In an ECG from a healthy person the T waves are normally positively deflected in all leads except aVR, apart from V1 which may be positive, negative, or flat (table 6.3). A negative T wave in lead III or aVL is rare in control patients (table 6.3). Inversion in leads V1–V3 is a normal variant in children and occasionally persists into adulthood (Hayden et al. 2002). Biphasic T waves are usually indicative of ischaemia or hypokalaemia (Mattu 2003).

Chapter 3

Multiscale modelling of the heart: from ion channel to body surface

Simulating the electrical fields in the heart and torso during a heartbeat requires mathematical and computational techniques and materials from multiple disciplines. The problem is intrinsically multi-scale, spanning spatial scales from the human body to cardiac cells and ion channels (fig. 3.1). In this chapter I will describe the mathematical model of the ventricle-torso system used in this thesis, the three-dimensional domain, the simulation parameters, and the software used for simulation.

3.1 Models of cardiomyocytes

Mathematical analysis of physiological function was revolutionized sixty-five years ago in a series of papers by Hodgkin and Huxley (1952). The authors used measurements of ion channels in a squid axon to formulate differential equations describing the opening and closing of the channels. Remarkably, the mathematical formulations they used would later be identified as reflections of underlying

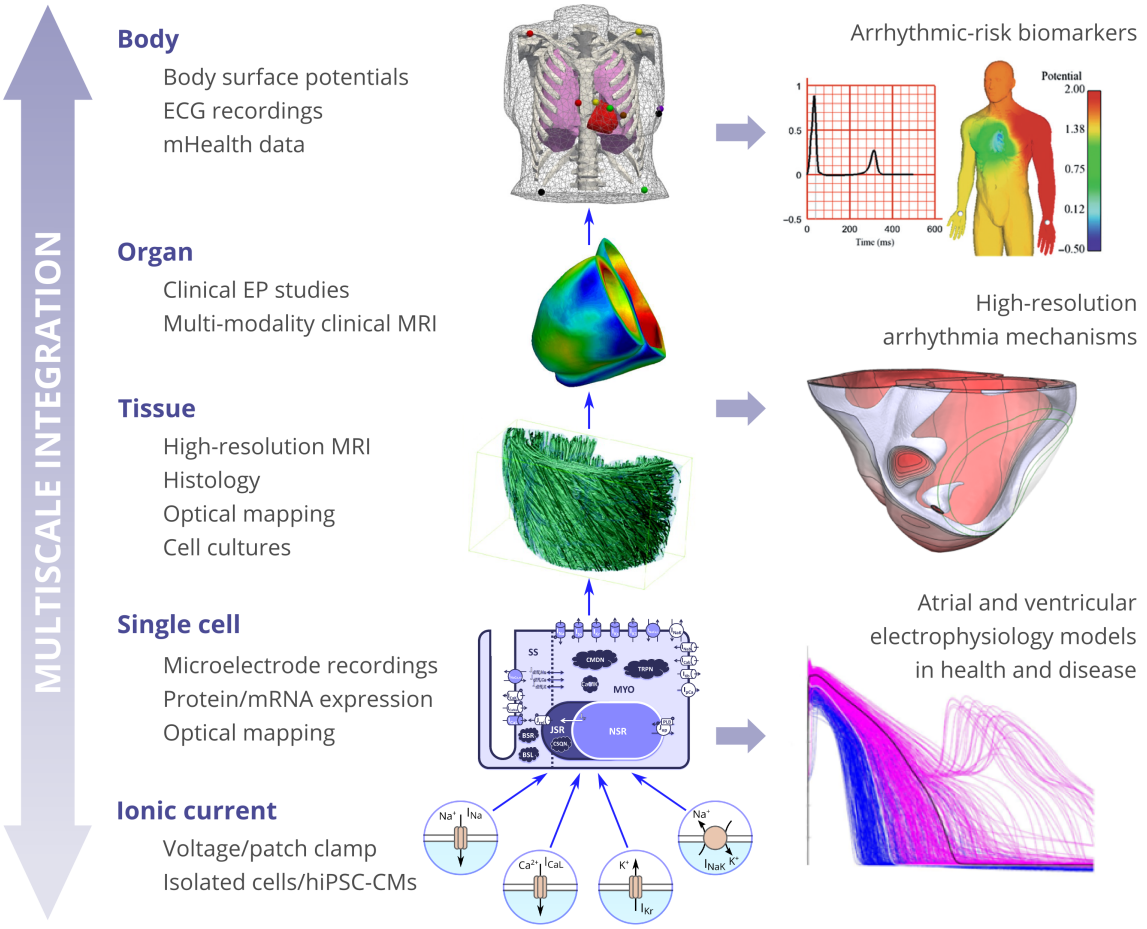

Figure 3.1: Cardiac computational biology experiment are intrinsically multi-scale, bridging spatial scales from the human body to the subcellular movement of ions. Reproduced from Rodriguez et al. (2015) with permission from Oxford University Press.

molecular configurations, leading to a molecular-biological interpretation. These forms are widely used in myocyte models. This section contains a brief summary of the different types of cardiac models, including selected examples. For a thorough review see Noble, Garny, and Noble (2012).

The simplest models contain just two variables describing excitation and recovery. The Fitzhugh-Nagumo model published independently in the early 1960s (FitzHugh 1961; Nagumo, Arimoto, and Yoshizawa 1962) denotes these as u and v respectively, and is of the form $\frac{\partial u}{\partial t} = f(u, v)$, $\frac{\partial v}{\partial t} = g(u, v)$, where f is cubic in u and linear in v, and g is linear in both u and v. The main advantage of this model is

its simplicity, as analysis of excitation and recovery is possible using a phase plane with u plotted against v. The main disadvantages however are certain unphysiological characteristics (e.g. absence of refractory period at fast pacing lengths), and a lack of detail at the molecular level.

Biophysically detailed models are based on experimentally derived data of the voltage and time dependence of different ion channel currents. The vast majority follow the formulation set out by Hodgkin and Huxley (1952). That is, for an ion species S, the current I_S is the product of a maximal conductance $\overline{G}_S$, some gating variables $x, y, \ldots, z$ raised to powers $l, m, \ldots, n$ respectively, and the 'driving force' (the difference between the transmembrane potential V and the Nernst potential E_S):

$$C_m \frac{\mathrm{d}V}{\mathrm{d}t} = -\sum_S I_S \tag{3.1}$$

$$I_S = \overline{G}_S x^l y^m \ldots z^n (V - E_S) \tag{3.2}$$

where C_m is the membrane capacitance. Their original model contained three currents: sodium (I_{Na}), potassium (I_K), and a leakage current (I_L). The gating variables take the range $[0, 1]$, and may represent activation, inactivation, or deactivation of the pore, where a value of 1 indicates fully open (or *not* inactivated, or *not* deactivated) and a value of 0 indicates fully closed (or fully inactivated, or fully deactivated). They are governed by ordinary differential equations (ODEs) of the form

$$\frac{\mathrm{d}x}{\mathrm{d}t} = \frac{x_\infty(V) - x}{\tau_x(V)} \tag{3.3}$$

where x_∞ is the steady state value of the gating variable, and τ_x is a time constant. Both of these are voltage dependent and must be fitted from data.

The first cardiac model of this type was published in 1962 by Noble and de-

scribed Purkinje cells by reformulating the potassium current and refitting the parameters. The first widely used model of mammalian ventricular myocytes was the Beeler-Reuter model (Beeler and Reuter 1977), which included four currents (attributed to sodium, calcium, and two to potassium) and allowed the intracellular calcium concentration to vary. Luo and Rudy (1991) extended this model to include six currents, and again refitted the parameters to new experimental data from guinea pig.

The second generation models include detailed descriptions of intracellular ion concentrations and/or separate subspaces within the cell. Conservation is achieved by including differential equations of the form

$$\frac{\mathrm{d}[S]}{\mathrm{d}t} = \sum_{T} I_T(V) \tag{3.4}$$

where T currents are implicated in determining the concentration of the species S. For example, in the case of potassium, T might enumerate the multitude of potassium currents (I_{Ks}, I_{Kr}, I_{Kb}, I_{K1}, $I_{\mathrm{K,ATP}}$, $\ldots$), the sodium-potassium pump ($I_{\mathrm{Na/K}}$), etc.

The first of these models was again a Purkinje cell model, by DiFrancesco and Noble (1985), the first model to contain pumps and exchangers. The first human ventricular models (two of them) were published in 2004, by ten Tusscher et al. (2004) and Iyer, Mazhari, and Winslow (2004), and several more have been developed since. Some widely used human ventricular models which are especially worthy of note are:

- The ten Tusscher and Panfilov (2006) model (an extension of ten Tusscher et al. (2004))
- The Grandi, Pasqualini, and Bers (2010) model (a refit of the Shannon et al. (2004) rabbit model with human data)

- The Carro et al. (2011) model (a refit of the Grandi, Pasqualini, and Bers (2010) model to study arrhythmias)
- The O'Hara et al. (2011) model

The latter is noteworthy as it was developed using very thorough and consistent human data, making it perhaps the most human-specific model to date. This is in contrast to most models, which are developed using data from multiple species and a variety of experimental procedures.

In this thesis the model of ten Tusscher and Panfilov (2006) was used, as it provides a compromise between speed and detail. The model is well-established for use in tissue simulations of activation and repolarization (Okada et al. 2013; Villongco et al. 2014). A cartoon of the transmembrane currents and sarcoplasmic reticulum is shown in fig. 3.2. The model is an extension of the ten Tusscher et al. (2004) model, with 'a more extensive description of intracellular calcium handling, incorporating a subsarcolemmal space, describing CICR with a Markov-state model for the ryanodine receptor, including both fast and slow voltage inactivation for the L-type calcium current, and applying some minor changes to parameter values and to slow delayed rectifier time dynamics'.

3.2 Models of cardiac tissue

There are two approaches in tissue modelling: discrete and continuum based. Discrete models such as cellular automaton are relatively easy to implement, and sophisticated models have certain strengths, for example in modelling tissue architecture (e.g. explicitly including gap junctions). For the large-scale (whole-organ) simulations presented here however they are computationally unfeasible and a continuous approximation must be used. At this scale, the propagation of an action potential through the tissue is sufficiently smooth to use a continuous descrip-

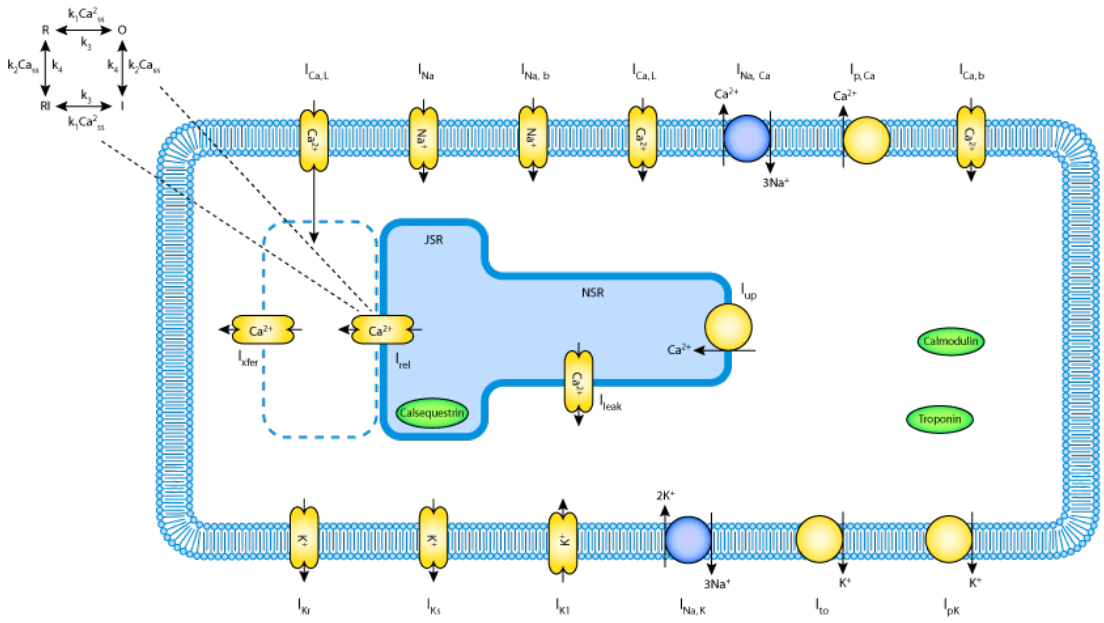

Figure 3.2: Cartoon of the ten Tusscher and Panfilov (2006) model showing the transmembrane currents and subcellular spaces. Figure licensed under CC BY 3.0 (`https://models.physiomeproject.org/exposure/a7179d94365ff0c9c0e6eb7c6a787d3d`).

tion of the cells. In this section the conventional notation is used of a subscript i and a subscript e for intracellular and extracellular quantities, respectively.

3.2.1 Bidomain equations

The bidomain equations were first expressed independently by Tung (1978) and Miller and Geselowitz (1978). The physical interpretation is of a region Ω containing two overlapping electric fields ϕ_e and ϕ_i, representing the electrical potentials in the extracellular and intracellular domains, respectively. The intracellular and extracellular quantities necessary for the cell model (such as ionic concentrations) are also defined at all points in space. The cell model couples all the other fields together, as the difference between the electric fields is the transmembrane potential, which is coupled to the activity of the ion channels, and therefore the ionic

concentrations.

Using the parabolic-elliptic form of the equations, the bidomain problem is: find V and ϕ_e satisfying

$$\chi \left(C_m \frac{\partial V}{\partial t} + I_{\text{ion}}(\vec{u}, V) \right) - \nabla \cdot \left(\sigma_i \vec{\nabla} \left(V + \phi_e \right) \right) = -I_i^{(\text{vol})}, \tag{3.5}$$

$$\nabla \cdot \left(\sigma_i \vec{\nabla} V + (\sigma_i + \sigma_e) \vec{\nabla} \phi_e \right) = I_{\text{total}}^{(\text{vol})}, \tag{3.6}$$

$$\frac{\partial \vec{u}}{\partial t} = \vec{f}(\vec{u}, V), \tag{3.7}$$

where: σ_i and σ_e are conductivity tensors of the intracellular and extracellular domains, respectively; χ the surface-area-to-volume ratio of a cardiac cell; C is the membrane capacitance (per area); V is the transmembrane potential ($V \equiv \phi_i - \phi_e$); I_{ion} is the total ionic transmembrane current due to the activity of pores in the membrane (depending on state variables $\vec{u}$, which themselves are governed by a system of differential equations $\vec{f}$); and the $I^{(\text{vol})}$ are stimulus currents (with units of current/volume) where $I_{\text{total}}^{(\text{vol})} \equiv I_i^{(\text{vol})} + I_e^{(\text{vol})}$.

I_i terms are interpreted as current injected into the intracellular space, and I_e terms as current into the extracellular space (usually reserved for simulating shocks e.g. in defibrillation). As this thesis only uses intracellular stimuli, stimulus currents are applied such that $I_e^{(\text{vol})} = -I_i^{(\text{vol})}$ (i.e. $I_{\text{total}}^{(\text{vol})} = 0$) to ensure conservation of current (Pathmanathan et al. 2010).

General boundary conditions are

$$\vec{n} \cdot \left(\sigma_i \vec{\nabla} (V + \phi_e) \right) = I_i^{(\text{surf})}, \tag{3.8}$$

$$\vec{n} \cdot \left(\sigma_e \vec{\nabla} \phi_e \right) = I_e^{(\text{surf})}, \tag{3.9}$$

where $\vec{n}$ is the outward-facing unit normal and the $I^{(\text{surf})}$ are stimulus currents (with units of current/area). No surface stimuli are used in this thesis, i.e. $I_i^{(\text{surf})} =$

$I_e^{(\text{surf})} = 0$.

3.2.2 Bidomain equations with a perfusing bath

In this thesis the ventricle-torso system is modelled using the previously reported formulation for cardiac tissue in a conductive bath (Pathmanathan et al. 2010), by identifying the torso tissue volume as a bath. An additional bath region Ω_b is introduced with conductivity σ_b (fig. 3.3). The transmembrane voltage V is not defined in this bath region, but ϕ_e is defined throughout the entire domain. Therefore, in addition to equations (3.5)–(3.7) we also have

$$\nabla \cdot (\sigma_b \nabla \phi_e) = 0 \quad \text{in } \Omega_b \tag{3.10}$$

with boundary conditions (3.8) on the tissue boundary and (3.9) on the bath boundary, as well as an interface boundary condition

$$\vec{n} \cdot (\sigma_e \nabla \phi_e) \big|_{\partial\Omega^{\text{tiss}}} + \vec{n} \cdot (\sigma_b \nabla \phi_e) \big|_{\partial\Omega^{\text{bath}}} = 0 \qquad \text{on } \partial\Omega \tag{3.11}$$

where $\big|_{\partial\Omega^{\text{tiss}}}$ denotes the evaluation on $\partial\Omega$ as from the tissue side and $\big|_{\partial\Omega^{\text{bath}}}$ the bath side.

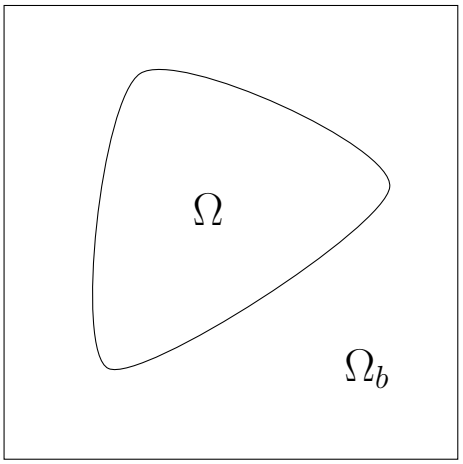

Figure 3.3: Domains in a model of cardiac tissue Ω within a conductive bath Ω_b.

Many ECG simulation studies simplify the problem by uncoupling the heart and torso, allowing separate solution of the bidomain equations in the heart and the Laplace equation in the torso (Lines et al. 2003). For superior accuracy and solver convergence I instead solve the bidomain equations for the entire domain simultaneously (Southern et al. 2009; Zemzemi et al. 2012).

3.2.3 Cardiac microstructure

Tissue fibre structure was generated using a rule-based method to replicate the findings of Streeter et al. (1969) (fig. 3.4). Essentially this method generates fibres that run circumferentially around the ventricles. Near the mid-wall of the ventricles they lie within the short axis plane. If we define 90° and −90° to point towards the base and apex, respectively, then at the inner and outer surfaces the fibres are angled relative to the mid-wall fibres by 60° and −60°, respectively. This angle varies continuously between the surfaces. This idealization has been validated experimentally (Jouk et al. 2007) and in simulation (Bishop et al. 2009).

In Chaste, the Streeter fibre generation is performed as a pre-processing step on the mesh, resulting in a 3×3 tensor at each element that describes the (cartesian) components of each of the three vectors in the fibre, sheet, and inter-sheet directions. The magnitude of each of these vectors represents the ventricular tissue conductivities (see below).

3.2.4 Tissue conductivities

Ventricular conductivities

Measurements have been reported in the literature of the conductivities and action potential conduction velocities (CV) of cardiac tissue in the fibre direction and transversely (Clerc 1976; Roberts, Hersh, and Scher 1979; Roberts and Scher 1982).

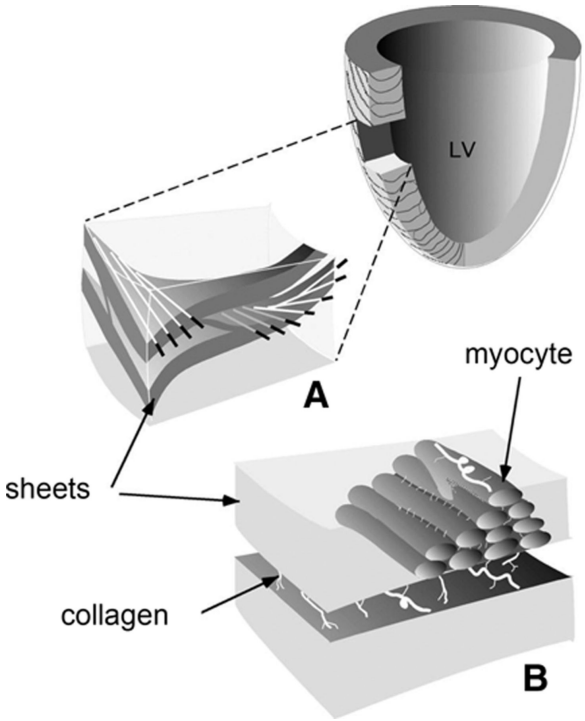

Figure 3.4: Cartoon of myocyte fibre orientations in the ventricles. (A) Fibres run circumferentially in the left ventricle (LV) at an angle that varies transmurally. (B) The fibres form sheets separated by collagen. Reproduced from Li et al. (2014) with permission from Wolters Kluwer Health, Inc.

Since the discovery of heterogeneity in the non-fibre direction, the ratio of conductivity in the fibre, sheet, and inter-sheet directions in the ventricles in human has been reported as approximately 4:2:1 (Caldwell et al. 2009; Hooks et al. 2007).

Simply using the experimentally derived conductivities as parameters in a simulation does not necessarily lead to physiological CV, because the CV is a result of the complete modelling approach, including the mesh, cell model, and numerical techniques used. The approach taken in this thesis, therefore, was to select conductivities that simultaneously generated the desired CVs in simulation given the cell model, mesh, etc., as well as conforming to experimentally measured conductivity ratios.

The relationship between cardiac tissue conductivities and conduction velocity is shown in fig. 3.5. Data for this figure were calculated from a 5 cm 1-D fibre model

with 0.4 mm node spacing (corresponding to the mean edge length of the mesh presented in section 3.3.2) and the endocardial variant of the 2006 ten Tusscher and Panfilov model. A stimulus was applied at one end of the cable and the CV was calculated from the middle 3 cm to mitigate edge effects.

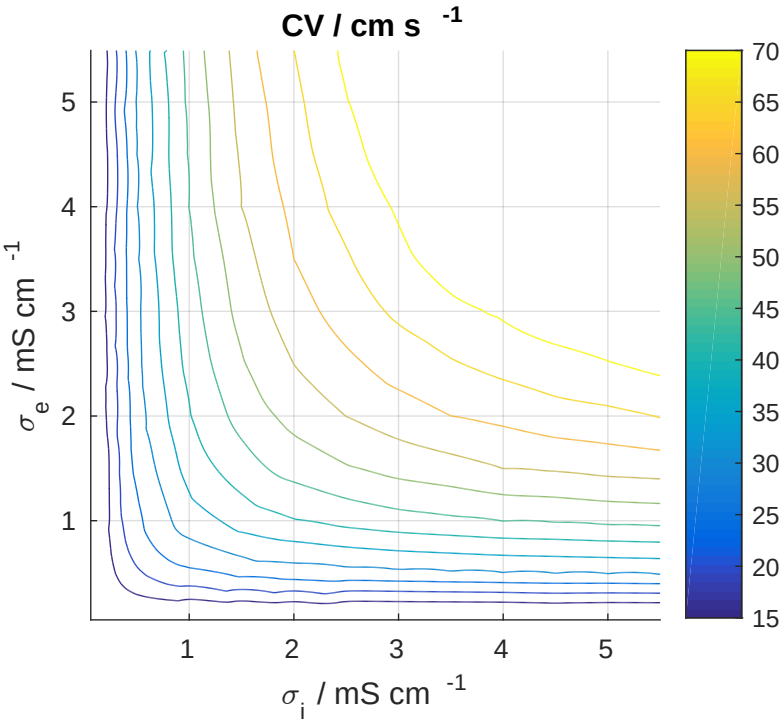

Figure 3.5: Estimated conduction velocity (CV) in a 1-D fibre depending on intracellular and extracellular conductivities under similar conditions to the 3-D simulations (see text). Contours are shown from 15 cm s^{-1} to 70 cm s^{-1} in steps of 5 cm s^{-1}.

Unless otherwise noted, the tissue conductivities used in the simulations presented in this thesis are listed in table 3.1. They were determined as follows. The three orthotropic intracellular conductivities were 1.5, 0.45, and 0.225 mS/cm in the fibre, sheet, and inter-sheet directions. Axisymmetric extracellular conductivities were calculated from these intracellular values scaled by measured resistivity ratios (Clerc 1976; Sebastian et al. 2008). In the fibre direction, 1.5 mS/cm was increased by a factor of 3.64 (the ratio of the extracellular fibre and intracellular fibre conductivities) to give 5.46 mS/cm. The transverse conductivity was set relative to

this, by scaling 5.46 mS/cm by 2.69^{-1} to give 2.03 mS/cm. In a 1D fibre model these conductivities resulted in conduction velocities of 57.0, 25.9, 16.2 cm/s, which are close to the 67, 30, and 17 cm/s wave speeds reported in pig (Caldwell et al. 2009).

	σ_i / mS cm^{-1}	σ_e / mS cm^{-1}	CV / cm s^{-1}
Fibre	1.5	5.46	57.0
Sheet	0.45	2.03	25.9
Inter-sheet	0.225	2.03	16.2

Table 3.1: Ventricular conductivity parameters used in this thesis unless otherwise noted. σ_i and σ_e are intracellular and extracellular conductivities respectively. Speed refers to the activation wave speed in a 1-D cable simulation.

Extra-cardiac conductivities

The conductivities of the various tissues in the chest have been measured (Foster and Schwan 1989; Gabriel, Gabriel, and Corthout 1996; Gabriel, Lau, and Gabriel 1996; Rush, Abildskov, and Mcfee 1963). We assigned isotropic conductivities of 0.389 mS cm^{-1} and 0.2 mS cm^{-1} to the lung and bone elements in the mesh (see section 3.3.2) respectively (Gabriel, Lau, and Gabriel 1996). The remaining torso elements were given an average trunk measurement of 2.16 mS cm^{-1} (Rush, Abildskov, and Mcfee 1963).

3.3 Numerical solution of the electrocardiographic problem

3.3.1 The finite element method

The bidomain equations are usually discretized and solved in space with the finite difference, finite element, or finite volume methods (Clayton et al. 2011), which are numerical techniques for finding approximate solutions to differential equations.

In this thesis I use the finite element method (FEM), which essentially involves dividing the continuous domain into small elements and approximating the solution at each of these elements using a set of basis functions. The FEM proceeds by first finding a weak form of the solution by multiplying through by a test function, integrating, and then expressing it as a system of linear equations which can be efficiently solved by a specialized solver. See Bernabeu et al. (2014) and Pathmanathan et al. (2010) for further mathematical and implementation details.

An advantage of the FEM over others (including the finite difference method) for cardiac modelling is that it allows us to easily use unstructured (irregular) meshes to describe the domain, which is useful as the problem geometry is typically generated from medical images. These images are segmented to isolate the myocardium, and processed to generate a collection of elements (nodes and edges) which may be refined or coarsened to achieve a given level of anatomic detail and solver accuracy.

3.3.2 Anatomical ventricle-torso tetrahedral mesh

The biventricular heart mesh used in this thesis was generated from a multi-slice computer tomography (CT) scan, as described in a previous study (Sebastian et al. 2008). Briefly, segmentation was performed automatically using an atlas of scans from healthy and diseased patients. The ventricles were isolated by truncating with a plane at the base below the valves. The mesh was generated with an edge length of 0.4 mm, and contains 2.51 million nodes and 14.2 million tetrahedral elements. This edge length is sufficiently small as Chaste has been shown to generate accurate and consistent conduction velocities for meshes with resolutions of 0.5 mm, 0.2 mm, and 0.1 mm (Niederer et al. 2011).

The biventricular mesh was embedded in a torso geometry generated from CT images of a 43-year-old woman. These CT DICOM images were segmented using

the medical imaging software Osirix[1] into lung, bone, and the rest. These surfaces were meshed using INRIA meshing Software MMG3D[2]. The heart and torso meshes were translated and rotated into anatomically realistic concordance (Nejib Zemzemi, INRIA, France). The combined mesh has a total of about 3.25 million nodes and about 19.4 million tetrahedra. A visualization of the resultant heart-torso mesh is shown in fig. 3.6.

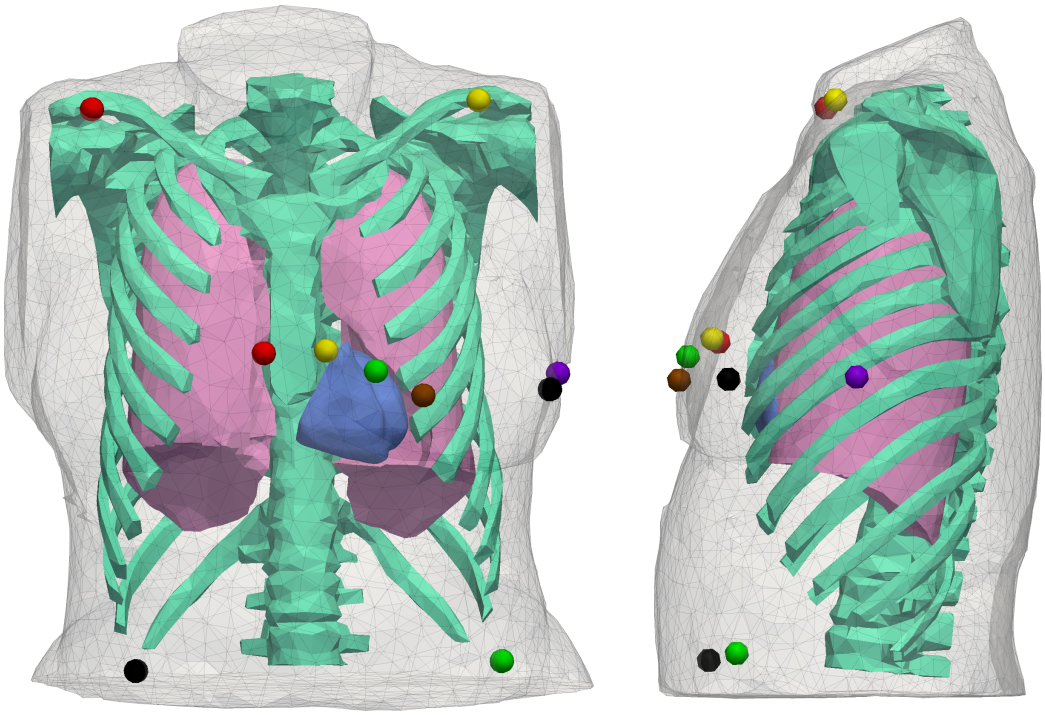

Figure 3.6: Visualizations of the combined ventricle-torso mesh (anterior and lateral views). The coloured spheres indicate the location of the recording electrodes using standard colour-coding (Society for Cardiological Science and Technology 2010).

3.3.3 Generating a simulated ECG

By recording the extracellular potentials satisfying the bidomain equations at nodes corresponding to virtual electrodes we can calculate a simulated ECG using the definitions in section 2.2.1. The final step before doing so is filtering the signals,

[1] http://www.osirix-viewer.com
[2] http://www.math.u-bordeaux1.fr/~dobrzyns/logiciels/mmg3d.php

as is done by commercial ECG machines in the clinic to reduce noise and baseline drift in the leads. Guidelines suggest using a bandpass filter of 0.05 Hz to 100 Hz (Society for Cardiological Science and Technology 2010).

The high pass filter was not required because the simulations include neither baseline drift nor motion artefacts like breathing. All the simulated ECGs in this thesis have therefore been calculated from ϕ_e time series filtered with a fourth-order low-pass Butterworth filter with cutoff frequency 100 Hz (−3 dB). See fig. 3.7 for a magnitude plot. The MATLAB function `filtfilt` was used to achieve zero-phase filtering.

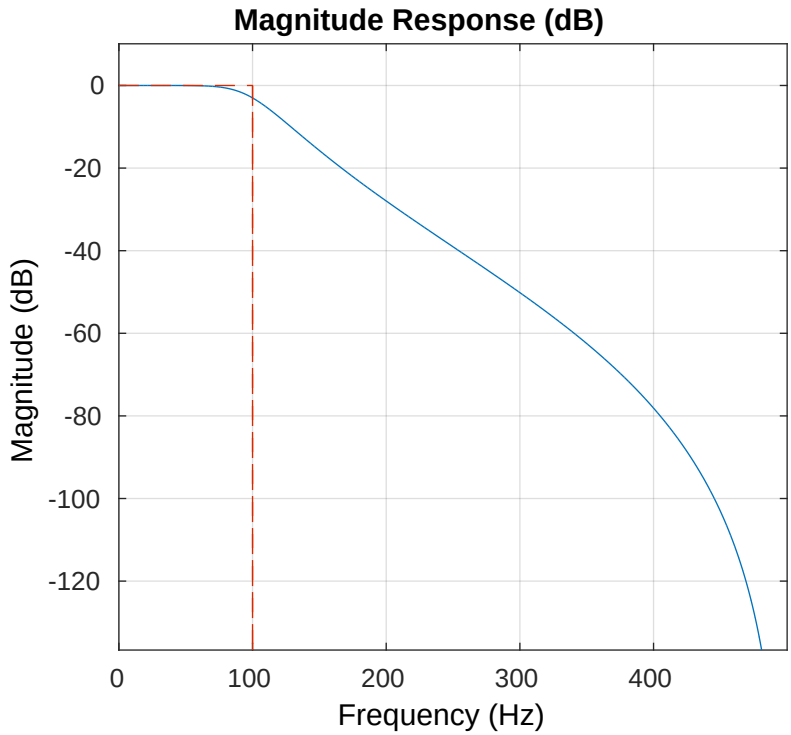

Figure 3.7: Magnitude plot of the filter applied to the ϕ_e signals recorded at virtual electrodes on the body surface during simulations before calculating the ECG leads.

3.3.4 Chaste

All computation was performed using the open source software package Chaste (Cancer, Heart, and Soft Tissue Environment) (Mirams et al. 2013; Pitt-Francis et al. 2009), a flexible set of software libraries for solving a wide range of problems in computational biology. Here we use it to: (1) solve the bidomain equations on the unstructured mesh described above, (2) produce visualizations (voltage distribution, action potential duration (APD) maps, and activation time maps), and (3) output extracellular potential signals at the virtual electrodes. Chaste is free and open source software that integrates many components useful for cardiac tissue simulations, as summarized in table 3.2.

Monodomain and bidomain equation solvers and preconditioners using PETSc (Portable, Extensible Toolkit for Scientific Computation) (Balay et al. 1997, 2014) *Krylov space (KSP) solvers:* PETSc: Conjugate gradient (CG), generalized minimal residual method (GMRES), SYMMLQ Custom: hybrid CG-Chebyshev solver (Bernabeu et al. 2014) *Preconditioners:* PETSc: Block Jacobi, analytic multigrid (HYPRE and ML), SPAI Cutom: BD and LDU (Bernabeu and Kay 2011)
ODE solvers for cell models Euler method, Backward Euler method, Classical Runge-Kutta method (RK4), Runge-Kutta-Fehlberg method, Heun's method, Rush-Larsen schemes (Sundnes et al. 2009), CVODE (Hindmarsh et al. 2005)
Visualization generation VTK, Meshalyzer, CMGUI
Miscellaneous Parallel mesh partitioning using parMETIS HDF5 (Hierarchical Data Format) output Streeter et al. (1969) cardiac tissue anisotropy model PyCml to automatically convert CellML models to Chaste code (Cooper, Spiteri, and Mirams 2015)

Table 3.2: Summary of Chaste software library features and external software package support.

A simulation typically proceeds as shown schematically in fig. 3.8. The first step is to load the mesh, which must be provided in the Chaste format (based on the TetGen[3] format) consisting of several files such as a node list (`.node`), an element list (`.ele`), fibre definitions (`.axi` or `.ortho` for axisymmetric or orthotropic fibres, respectively), and lists of LV/RV or endo/epi surface nodes. A tissue model is then constructed based on the nodes in the mesh, including cardiac and other regions. Within the main solve loop (grey box) there is a hierarchy of time step sizes: the cell models are usually solved repeatedly before solving the PDEs, and the PDEs are usually solved repeatedly before writing to the HDF5 file (the main results file, see section 4.4). For example, the ratio of ODE time step to PDE time step to HDF5 file 'print' time step might typically be of the order of 1:10:1000. Once the simulation is complete the HDF5 file is closed and optionally used as an input for one or more postprocessing tasks (e.g. for calculating upstroke times or conduction velocities) and/or data conversion tasks (e.g. for producing visualization files).

3.3.5 Solver methods and parameters

In this thesis, simulations were performed with Chaste using the SYMMLQ solver and BD (block diagonal) preconditioner (table 3.2). Cell models were solved with CVODE, with relative and absolute tolerances set to 1×10^{-5} and 1×10^{-7}, respectively. Visualization files were generated using the 'parallel VTK' writer for rendering in Paraview. Other features were used according to the defaults (parMETIS, PyCml, etc.).

A convergence analysis was performed to determine the choice of PDE time step. Details of two leads from a typical bidomain simulation with four different PDE time steps (0.0125 ms to 0.1 ms) are shown in fig. 3.9. The signal generated using the second-smallest time step (0.025 ms) was within 1 % of the signal using

[3] `http://wias-berlin.de/software/tetgen/`

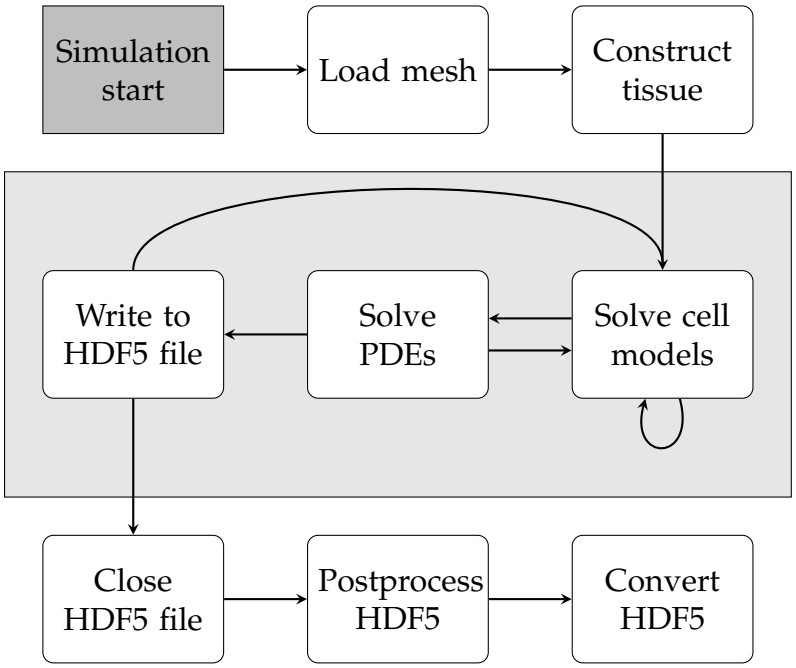

Figure 3.8: A schematic of the stages of a typical cardiac simulation in Chaste. The first row shows the initialization tasks, the second row (grey box) the 'main solve loop', and the third row the finalization tasks.

the smallest time step (0.0125 ms). Also, under certain conditions the solver was unstable with a step size of 0.1 ms. Based on these findings and considering the relative performance, a PDE time step of 0.025 ms was used for the simulations in this thesis.

Cell model ODEs were solved using CVODE, which uses variable-order, variable-step multistep methods. The main practical advantages of this solver over traditional fixed time-step solvers such as backward Euler are accuracy and speed. Solving cardiac action potential models demands a small time step around the upstroke, which is only a brief part of the action potential. The rest of the action potential is relatively slowly varying, allowing a longer time step to be used. This separation of temporal scales in the action potential makes CVODE several times faster than backward Euler as it can dynamically adapt the time step and numerical order to the problem stiffness. Furthermore, the solver stays within maximum relative and absolute errors prescribed by the user, obviating time-consuming con-

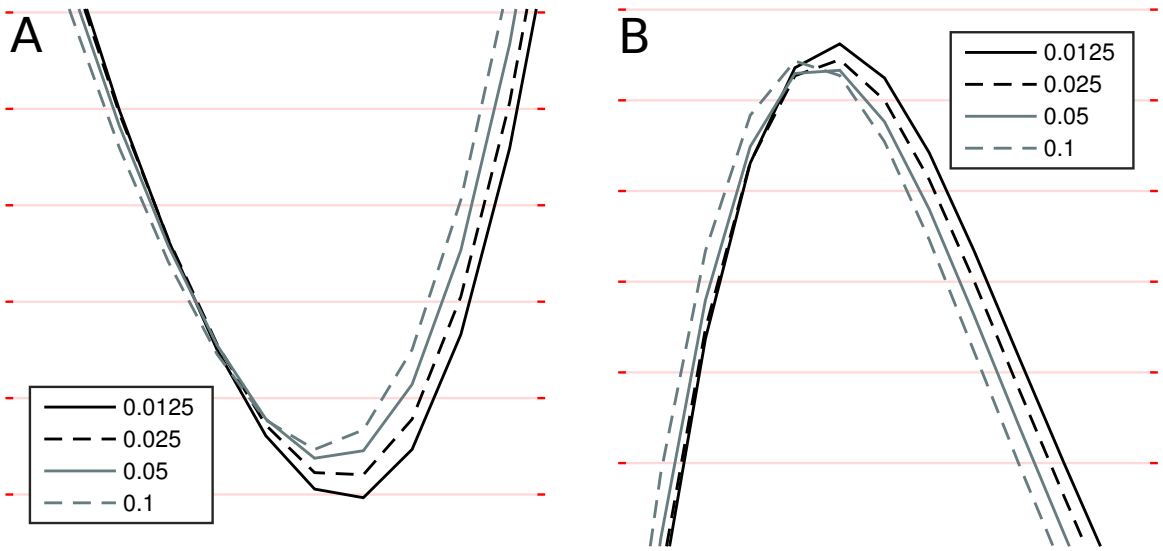

Figure 3.9: Detail of two leads, (A) V1 and (B) I, for four sizes of PDE time step ranging from 0.0125 ms to 0.1 ms. Each plot shows about 12 ms of data sampled every 1 ms, at the S and R wave peaks (36 ms and 29 ms), respectively. In (A) the grid has 20 µV spacing, in (B) it is 5 µV.

vergence testing.

3.3.6 HPC performance summary

On Archer, following the performance improvements described in chapter 4, simulations were typically performed on 8 compute nodes, i.e. 192 Intel Xeon "Ivy Bridge" cores. With these compute resources a simulation of a complete heart beat (of the order of 500 ms) required about one hour of wall time.

Chapter 4

High performance optimization of the electrocardiographic problem

4.1 Introduction

Chaste's parallel scaling on high performance computing (HPC) systems was substantially improved shortly before commencement of my DPhil, mainly through PETSc Krylov subspace (KSP) method solver options (Bernabeu et al. 2014). In my preliminary testing, however, it was clear that several areas required improvement in order to use Chaste for routine simulation of the human heart and torso. In this chapter, I describe the main improvements to Chaste resulting from my work. Several optimizations were implemented with assistance from the other Chaste developers: mesh reading and writing, cell model ODE integration, and input/output (I/O) performance with the HDF5 library. Finally, a speed-up specific to my simulations was realized from the use of adaptive time step size.

Early results were generated on the supercomputer HECToR, a Cray XE6 system with 90 112 CPU cores that was located at the University of Edinburgh. It was superseded in 2013 by Archer, a 118 080-core Cray XC30 system at the same facility.

4.2 Binary mesh reading

The first step in any Chaste simulation is to load (or generate) a mesh to describe the problem domain. Despite substantial improvements made by Bernabeu et al. (2014) in mesh loading on a former supercomputer, HECToR, particularly the implementation of binary (BIN) mesh files and the node-connectivity list (`.ncl`) file, binary mesh reading was still unsatisfactory towards the end of my first year (July 2013). When we measured the time to load a 2.51 million node ventricular mesh with 512 processes on HECToR, reading a binary version of the mesh took over five times *longer* than reading the regular text (ASCII) version, despite the binary mode ostensibly being more optimized (table 4.1).

	InMesh/s	Init/s	Total/s
ASCII	460	170	520
BIN (before)	2755	35	2821
BIN (after)	80	38	121

Table 4.1: Mesh loading (InMesh) and initialization (Init) times required to load a 2.51 million node mesh in text (ASCII) and binary (BIN) formats, the latter before and after improvements from this thesis. Times are given in seconds. BIN load time was reduced from over five times that of ASCII to over four times *less* than ASCII.

The main cause of slow BIN mesh reading was identified as a random reordering mode during partitioning (Bernabeu et al. 2014). The idea of this earlier optimization was to randomly shuffle the order in which each process would read the element file in order to reduce the likelihood of many processes overwhelming a given storage device (in other words, to reduce *disk contention*). In practice this meant that reading each and every element involved a call to `seekg`[1] with a random offset relative to the start of the file, and on HECToR this operation was extremely slow.

[1]The files are opened as `std::istreams` (input streams). The method `std::istream::seekg` 'sets the position of the next character to be extracted from the input stream'.

The implemented solution was as follows:

- Removed the randomization code completely, thus removing any BIN-specific behaviour from the partitioning method.
- Made the mesh reader check if it is already in the correct position before doing a `seekg` for the node, element, face, and ncl files. In the case of reading an entire file contiguously (as for the element file when partitioning) this avoids calling `seekg` entirely.
- Replaced `seekg(offset, beg)` with `seekg(offset, cur)` for `offset` greater than the current position. In other words, if the access is monotonic (but not necessarily contiguous), reposition from the current location (`cur`) rather than calculating from the start of the file (`beg`).

These changes resulted in a speed-up of over 23 times with the BIN mesh, making BIN over four times faster than ASCII (table 4.1). On a current supercomputer, Archer, the mesh load time is only about 50 s for the larger (3.25 million node) mesh, due to superior hardware.

4.3 VTK mesh writing

Several output formats (table 3.2) are available to visualize a simulation in Chaste but, as the name suggests, the 'Parallel VTK' one is the most appropriate for HPC. An MPI issue was identified in the VTK mesh writer when writing large meshes to disk.

In the VTK mesh writer a deadlock was discovered under certain MPI implementations when writing VTK files in parallel. Each process writes its partition (chunk) to disk independently, which has better scaling than the sequential writer which concentrates all data onto a single process before writing. In order to write

consistent chunks, each process must communicate with its neighbours to request the value of variables at its halo nodes. The deadlock would occur when exchanging halo node values if the number of halos was sufficiently large.

The specific problem was a call to `MPI_Send` followed by one to `MPI_Recv`. Each process would attempt to: 1. send to the rank above followed by a receive from the rank below; 2. send to two ranks above followed by receive from two ranks below; and so on until every process had communicated with every other. For a small payload, MPI buffers the data and allows the program to continue, but I discovered that for a sufficiently large payload the send would become *blocking*, i.e. unable to continue until the matching receive was complete. If the receive was on a process that was also stuck in a blocking send then a deadlock would result. The solution was to replace the two separate operations with one `MPI_Sendrecv`, which is immune to such cyclic dependencies.

4.4 HDF5 performance

The primary means of storing Chaste simulation results on disk is in an HDF5 file. The HDF (Hierarchical Data Format) is designed for large and/or complex multi-dimensional datasets, is portable between systems, and offers high-performance for parallel computing. The `Hdf5DataWriter` class in Chaste was introduced in 2008 to replace a text-based file writer, but compared to other Chaste components it remained relatively overlooked in terms of optimization efforts until I began investigating in 2013. The central concept determining performance here is *chunking*, so this section begins with an introduction to chunking and data layout in HDF5, followed by details of the optimization, and finally benchmark results.

First though, a brief description of the typical usage of the HDF5 writer and reader in a Chaste simulation. After loading the mesh and setting up the problem

an HDF5 writer is invoked to either create a new HDF5 file or append to an existing one as requested. In the main solve loop, results are added to the file periodically, and at the end of the simulation the file is closed. If any postprocessing has been requested (e.g. APD durations, upstroke times), an HDF5 reader is then used to extract data from the HDF5 file, the calculations are performed, and a new HDF5 writer is used to save the statistics to new dataset(s) in the same HDF5 file.

4.4.1 Dataset chunk shape

At the hardware level, digital data storage is one-dimensional, so obviously there must be a mapping between the multidimensional datasets represented in HDF5 files and the disk. The default method simply serializes in 'row-major order', which might or might not be suitable depending on the *access pattern* (the order in which each process accesses values in the dataset).

Suppose we have a two-dimensional 10 by 10 dataset using this default layout (fig. 4.1). In row-major order the data are serialized row by row, e.g. elements 1 to 10 (labelled) are within the first row, and therefore contiguous on disk, and the same for the following rows. Using conventional matrix notation indexed from 1, if a process wants to read entries (3,4) to (3,8) inclusive, i.e. the 24th to 27th elements (shaded, left), then it can do this very efficiently by reading a contiguous region on disk (solid arrows). On the contrary, if a process wants to read entries (4,3) to (8,3) inclusive, i.e. the 33rd, 43rd, 53rd, 63rd, and 73rd elements (shaded, right), then it must perform a number of relatively expensive disk seeks between each row (dotted arrows). Clearly, for good performance the data layout must be based on the access pattern.

Since the way the data are laid out on the disk has a significant influence on performance, HDF5 allows us to specify a *chunk* shape based on typical access patterns. A chunked dataset is divided into repeating units of the chunk dimensions,

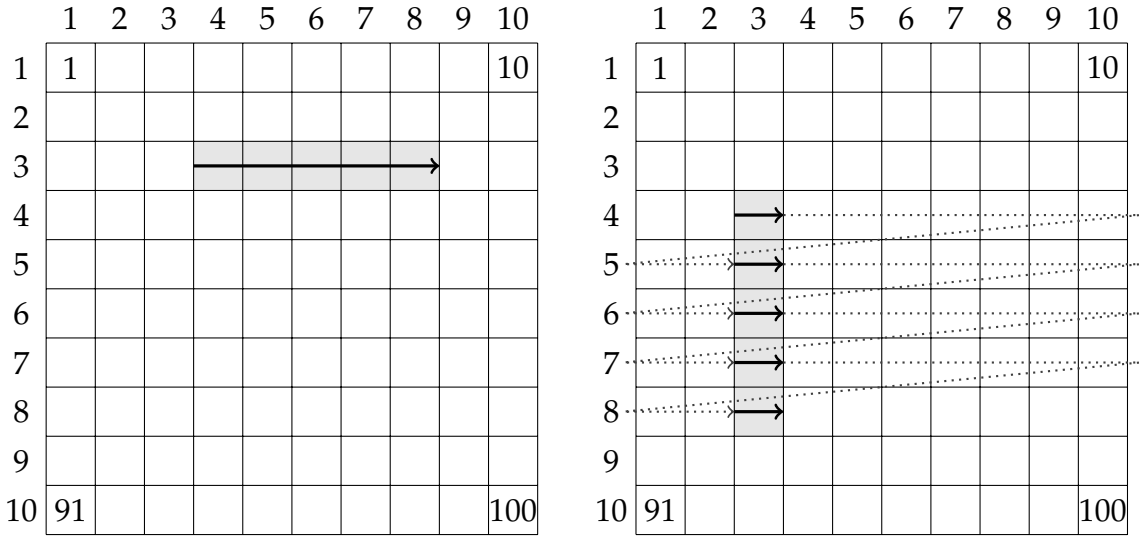

Figure 4.1: An illustration of the default (contiguous) data layout in HDF5, and the difference between reading a row (left) and a column (right): a row may be read efficiently because the elements are contiguous (solid arrow); reading a column by contrast requires a disk seek (dotted arrows) before every read. The columns and rows of the dataset are labelled 1–10. Four of the data in each dataset are labelled with their locations on disk (1, 10, 91, and 100).

and space for each chunk is allocated contiguously on the disk. In the second example above (reading part of a column) we might utilize chunking by setting the chunk dimensions to the shape of a column in the dataset (fig. 4.2). With chunks that coincide with columns, the library would lay the data out on the disk column by column, so that any access from a column becomes a single, fast, contiguous read.

A simple chunked layout was introduced early in development of the Chaste HDF5 writer because only chunked datasets may be extended beyond their initial size, e.g. when resuming a simulation from a *checkpoint*. Before describing the chunking further a brief description of (cardiac) Chaste datasets is needed. They are three-dimensional objects, over time in the first dimension, nodes in the second, and variables in the third (table 4.2).

When solving the electrocardiographic problem the time dimension is typically several hundred entries of results sampled at 1 ms intervals, such that a whole

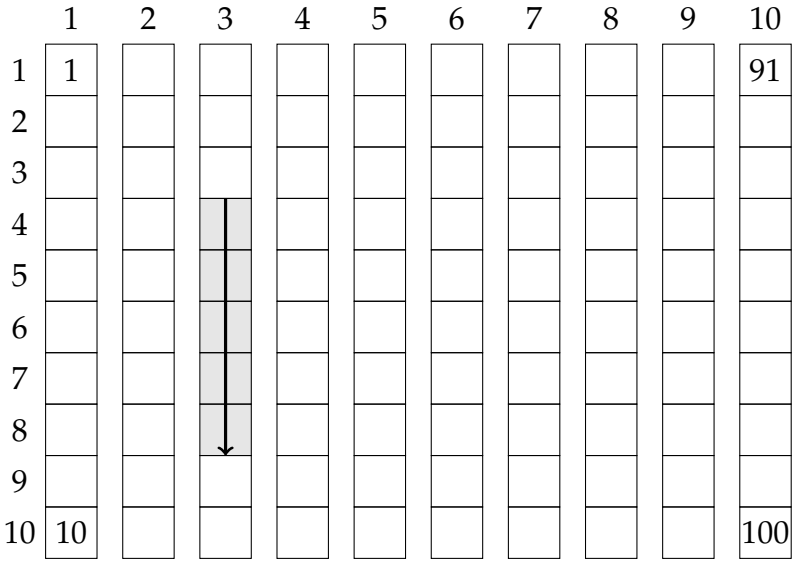

Figure 4.2: The same dataset and read operation as fig. 4.1 (right) but chunked with dimensions 1×10. With this chunked layout, reading from a column now involves reading part of a (contiguous) chunk, and is therefore much faster than with the default layout.

dataset describes one heartbeat. The nodal dimension contains a column for each of the $\mathcal{O}(10^6)$ nodes in the mesh. The 'variables' dimension contains a column for each output variable, i.e. V and ϕ_e for a bidomain simulation. Specifically, each value in the dataset is an 8 B double precision floating point number, so one time step is

$$8\,\mathrm{B} \times 3\,253\,316 \times 2 \approx 52.1\,\mathrm{MB}. \tag{4.1}$$

Simulation results for a heartbeat sampled at 1000 Hz (1 ms printing timestep) therefore run into the dozens of GB.

Formerly, a chunk written by the Chaste HDF5 writer had size $\{C_t, D_N, D_v\}$,

Dim.	Type	Typical size
1	Time	100 to 750
2	Nodes	3.3 million
3	Variables	2

Table 4.2: Types and typical sizes for the purpose of this thesis of the three dimensions of a cardiac Chaste HDF5 results dataset.

where D_N is the number of nodes in the dataset, D_v is the number of variables, and C_t is based on the number of timesteps D_t as

$$C_t = \max\left(100, \left\lfloor \frac{D_t}{100} \right\rfloor\right). \tag{4.2}$$

In other words, each chunk spanned the dataset in the second and third dimensions, and the dataset was split into about 100 chunks in the time dimension, but with no fewer than 100 elements in the time dimension. This strategy is fairly efficient for simulations with a small number of processors and nodes as each chunk will be relatively small and easy to cache. It is, however, poorly suited to parallel programs for several reasons as described below.

The most immediate problem was not in fact performance, but an upper limit on the number of nodes permitted in a mesh. HDF5 has a chunk size limit of 4 GiB, i.e. 536 870 912 double precision floating point numbers.[2] Using a chunk of size 100 in the time dimension and two variables per node (as in a bidomain simulation) implies an upper limit on the number of nodes of about 2.68 million. Since this thesis concerns problems of this order of magnitude, alleviating this restriction was an urgent priority.

Considering parallel performance, using chunks that span the node dimension is suboptimal due to the way the problem is partitioned across parallel processes. At the start of a Chaste simulation the mesh is partitioned and each process is assigned a subset of the nodes that remains unchanged for the duration of the simulation. For simplicity, the nodes are reordered (compared to their original order in the mesh node file) so that the nodes owned by each process are contiguous. A process is therefore only concerned with a block of the dataset corresponding to the

[2] A note on decimal and binary prefixes. In this chapter, the difference between 1000 (10^3) and 1024 (2^{10}) can be significant. Decimal prefixes (k, G, etc.) will generally be used for familiarity, but binary prefixes (Ki, Gi, etc.) will be used when they are more natural. In this example, the chunk size limit is exactly $4\,\mathrm{GiB} = 4 \times 2^{30}\,\mathrm{B} \approx 4.29\,\mathrm{GB}$.

nodes it owns for all time and variables. Using the earlier notation, each process will access a contiguous block of approximately size $\{D_t, D_N/j, D_v\}$ where j is the number of processes (assuming an equal partition of nodes between processes). The shape of this block of the dataset that each process will access is therefore 'orthogonal' to the chunk shape (fig. 4.3), very unlike using a chunk shape based on the access pattern.

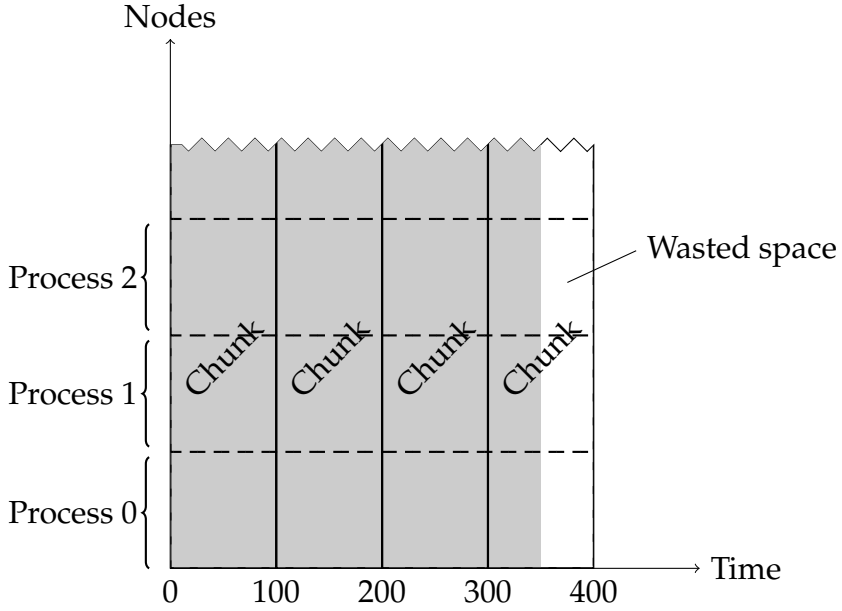

Figure 4.3: Cartoon illustrating the old chunk layout. This example depicts a dataset of size 350 in the time dimension and a large number in the nodes direction (grey). The chunks are of size $100 \times D_N$ (solid lines), so 50 elements (12.5 % of the file) are wasted at the edge (white). Each process is concerned with a slice of the dataset shaped orthogonally to the chunks (dashed lines). The third dimension has been suppressed for clarity, since the chunks and process boundaries span it.

At this point it is necessary to briefly introduce two of the HDF5 *drivers* used by Chaste, and how they differ in reading and writing a chunked dataset. When writing, Chaste uses the MPI-IO driver which is specialized for parallel applications. The chunk shape should have relatively little impact on writing because the MPI-IO driver uses direct access to the disk, and *collective writes* are used.[3] When

[3]In a collective write, data are first concentrated from every process to a smaller number of *aggregators*, which then perform the write to disk. This is desirable as a smaller number of larger writes

reading, however, (e.g. when performing post-processing or generating visualization files) the default driver is used which uses standard POSIX operations.[4] Unlike the MPI-IO driver, the default driver attempts to cache an entire chunk when data in it are accessed. Furthermore, the caching implementation in the default driver dictates that access can only be done on whole chunks. In other words, even if a process attempts to read just one entry from a chunk and it is possible for the chunk to be cached, then the process will read the entire chunk into the cache before continuing. If the chunk is too large to be cached then the cache is bypassed completely and direct access is used.

From the preceding paragraphs it should now be clear why using chunks that span the node dimension is very suboptimal. First, consider an HDF5 reader being used to perform post-processing on a dataset in parallel. Each process is expected to access all the variables for all its nodes for all time. Depending on the size of each chunk compared to the size of the chunk cache there are two possibilities.

- If the chunks are small compared to the size of the cache, then when a process requests a value in its block it must first read the entire chunk into its cache despite most of the nodes belonging to other processors and being therefore of no interest. In this case, if a process owns a proportion x of the nodes ($0 < x < 1$) then $1 - x$ of the total I/O activity would be wasted. If enough chunks can be cached and the access pattern is conducive then the overall performance might be acceptable, but it is easy to imagine circumstances where it would not be. For example, imagine the reader accesses all the time points for one node, then all the time points for the next node, etc. Unless all the chunks can be cached at once, the reader will be forced to read the entire dataset on each and every iteration, just to get the values for each node.

usually results in better bandwidth than a larger number of smaller writes. On Cray machines, the collective write algorithm is provided by ROMIO (an MPI-IO implementation included in MPI).

[4]Portable Operating System Interface.

- If the chunks are too large to fit in the cache then each process will access the dataset independently and directly. In this case, the potential performance improvement from using the cache is lost, but the requirement to read in whole chunks is dropped so only the requested entries are accessed, possibly resulting in less disk activity. Still, for optimal performance we should try to utilize caching.

Considering now a Chaste HDF5 writer, as written above, we would expect it to be less affected by the chunk settings than reading, since the MPI-IO driver only uses direct access, and the collective writing mode is designed to improve performance when writing from many processes to a single file. Nevertheless, the chunk shape results in every chunk being accessed simultaneously by every process (recall fig. 4.3), possibly resulting in increased library overhead to track modifications and maintain consistency between processes. A chunk layout more closely resembling the process boundaries would alleviate this issue.

4.4.2 A new chunk algorithm

The most methodical way to proceed would be to set the chunk shape based on an analysis of the most frequent access patterns, within some chunk capacity limits (in e.g. bytes). The chunk size is important because (1) disks are generally better at reading fewer, larger regions than more, smaller regions, and (2) it influences the size of chunk cache needed for good read performance. Unfortunately, in Chaste there are two common modes of access at different points in the simulation. When solving, the fastest varying dimension is the variable dimension, followed by the node dimension, and finally the time dimension. When performing post-processing, it is not uncommon to instead access the variables for each node over all time.

Such an approach, however, would result in highly problem-specific and/or

machine-specific algorithms that might show good performance in some applications at the expense of others. Instead, I developed a simple and general algorithm to set the chunk size ($\vec{C}$) based only on the size of the dataset ($\vec{D}$) and one parameter, the target chunk size in bytes (T_B).

Because the HDF5 writer in Chaste can in principle be used with any number of dimensions, for maximum generality I decided that the chunk algorithm should treat each dimension the same. Another design requirement was that the chunk shape should result in high storage efficiency. HDF5 allocates the minimum integer number of chunks required to contain the dataset (recall fig. 4.3). Most conceivable chunk shapes, therefore, would result in unacceptable amounts of wasted space at the edges of the dataset, because the chunk size is unlikely to be (close to) a factor of the dataset size in every dimension.

Central to my solution is a variable 'target size' T (not to be confused with T_B). The crux of the method (for a given target size) is as follows. First, for each dimension, divide the dataset size by the target size and round up, which gives the minimum number of target-sized chunks that would span the dataset. Second, divide the dataset size by this number of chunks and again round up, which gives the actual size of chunk that is closest to the target size while still being close to a factor. The problem is then finding the target size that best satisfies the chunk size requirement in bytes.

The solution may be written concisely as follows. Let the chunk size and dataset size be vectors denoted by $\vec{C}$ and $\vec{D}$, respectively:

$$\vec{C} = (C_1, C_2, \ldots, C_N) \tag{4.3}$$

$$\vec{D} = (D_1, D_2, \ldots, D_N) \tag{4.4}$$

where N is the number of dimensions in the dataset (normally 3, see table 4.2). We

get $\vec{C}$ by finding the *largest* T such that

$$8\prod_{i=1}^{N} C_i \leq T_B \tag{4.5}$$

$$C_i \leq D_i \qquad (1 \leq i \leq N) \tag{4.6}$$

$$C_i = \left\lceil \frac{D_i}{\lceil D_i/T \rceil} \right\rceil \qquad (1 \leq i \leq N) \tag{4.7}$$

where inequation 4.5 represents the chunk size constraint (recall that each element is 8 B), inequation 4.6 limits the chunk to the dataset size in each dimension, and equation 4.7 defines the chunk size given the dataset size and target size.

The method for solving the above as implemented in Chaste is described in algorithm 1. The dataset size (array $\vec{D}$) and target chunk size (T_B) are predetermined (see below for more on the latter). The first while loop (line 9) increases the target size (T) and calculates the resulting chunk dimensions (array $\vec{C}$) each iteration until the target size in bytes (T_B) is reached, or the chunk spans the entire dataset (i.e. Υ is True). After leaving the while loop, if the target chunk size in bytes has been exceeded (line 13), the algorithm goes back one step to bring the size back below the target.

Given a target size in elements, the CALCULATECHUNKDIMS function (line 19) calculates chunk dimensions that aim for the target size in each dimension while minimizing wasted space at the dataset edge as outlined above. First, it calculates the minimum number of chunks of size T that would be required to span the dataset (x, line 23). Then, given x chunks, it calculates the minimum number of elements required in each chunk to span the dataset (line 24). This function also calculates the actual chunk size in bytes (C_B) and determines Υ. The chunk size in bytes is the product of 8 B and all the elements of $\vec{C}$ (line 24). Finally, if x is equal to 1 on every iteration then one chunk spans the entire dataset and Υ is True (line 26).

Algorithm 1 HDF5 chunk size algorithm

$\vec{D}$ ▷ Dataset size in elements
$\vec{C}$ ▷ Chunk size in elements
$C_B \leftarrow 0$ ▷ Chunk size in B
$T \leftarrow 0$ ▷ Target chunk size in elements
$T_B \leftarrow 128 \times 2^{10}$ ▷ Target chunk size in B
$\Upsilon \leftarrow$ False ▷ Whether chunk spans dataset

function SETCHUNKSIZE
    while $(C_B < T_B)$ & $!\Upsilon$ **do** ▷ While chunk is smaller than target
        Increment T
        $(\vec{C}, C_B, \Upsilon) \leftarrow$ CALCULATECHUNKDIMS(T)
    end while
    if $C_B > T_B$ **then** ▷ If chunk has exceeded target
        Decrement T
        $(\vec{C}, C_B, \Upsilon) \leftarrow$ CALCULATECHUNKDIMS(T)
    end if
end function

function CALCULATECHUNKDIMS(T)
    $C_B \leftarrow 8$ ▷ 8 B per element
    $\Upsilon \leftarrow$ True
    for i in dimensions **do** ▷ For each dimension
        $x \leftarrow$ CEIL(D_i/T)
        $C_i \leftarrow$ CEIL(D_i/x)
        $C_B \leftarrow C_B \times C_i$
        $\Upsilon \leftarrow \Upsilon$ & $(x \leftrightarrow 1)$
    end for
    return $\vec{C}, C_B, \Upsilon$
end function

The choice of T_B depends on the problem size and the computer. The default value was set to 128 KiB (16 384 doubles) as this resulted in consistent performance in the small 'profiling' tests that are run regularly to monitor Chaste's performance, whereas much smaller and larger values exhibited worse performance. For large problems it was increased to 1 MiB.

The algorithm could be substantially optimized but this was not done since it is only invoked once per dataset and is very fast, consisting of only arithmetic calculations and having only a weak dependence on dataset size ($O(N)$ or smaller). For example, since the problem is essentially to find the largest T such that $C_B \leq T_B$, and $T \leq \max(\vec{D})$ (otherwise we have one spanning chunk), it would be much faster to do a binary search for T between 1 and $\max(\vec{D})$ than increment monotonically as at present.

4.4.3 Lustre striping

The preceding section described a better way of serializing an HDF5 dataset by dividing it into relatively small 'logical' chunks. In this section I describe performance improvements at the filesystem level, which first requires some background detail about the *Lustre* filesystem commonly used in HPC systems including Archer.

Data striping is a common technique in computer data storage to increase data throughput by splitting files into segments and dividing the segments amongst multiple physical storage targets (e.g. hard disk drives). A similar technology by way of analogy is RAID 0,[5] where using n identical hard disks leads to a factor of n increase in read and write performance.

There are two parameters that may be set on a file or directory basis on Lustre: the stripe size (S) and stripe count (c). The former is obviously the size (in bytes) of

[5]Redundant Array of Independent Disks, level 0

each stripe. The latter is a misnomer, it is in fact the number of OSTs over which to divide the stripes.[6] See fig. 4.4 for an example. The system defaults on Archer are 1 MiB and 4, respectively, and there are 48 OSTs at the time of writing this thesis. The Archer documentation states that each OST is capable of writing at roughly 500 MB s^{-1}.

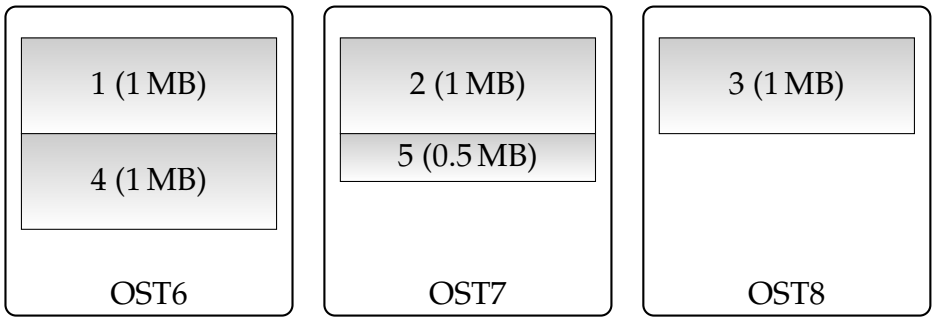

Figure 4.4: Cartoon illustrating striping on Lustre. This example depicts a 4.5 MB file striped across three OSTs with 1 MB stripe size. The first OST is chosen at random by the filesystem in order to load-balance, and in this example it is OST6. 1 MB stripes are then placed round-robin on each of the three OSTs, ending with a 0.5 MB stripe on OST7.

Optimal values for S and c can only be found through experimentation. The default stripe size on Archer (1 MiB) is generally optimal and other values will not be tried here. For reference, the Lustre documentation recommends a stripe size between 512 KiB and 4 MiB. Smaller sizes are not recommended 'because the Lustre file system sends 1 MiB chunks over the network'; more is not recommended because 'stripe sizes larger than 4 MiB may result in longer lock hold times and contention during shared file access'. Finally, the stripe size must be a multiple of the page size (Lustre software tools enforce a multiple of 64 KiB for compatibility).

The stripe count will be investigated below. Some general recommendations for setting the stripe count include the following:

- For directories of many small files (e.g. source code) set the stripe count to 1 (no striping) to reduce the overhead.

[6]Object Storage Target. Analogous to the number of hard disks in RAID 0.

- Where there is a one-to-one mapping between files and processors, set stripe count to 1 (no striping). Parallelism is provided by the multiple files, and this avoids any overhead from striping.

- For writes from many processors to a single file use a large stripe count, but not too many as this results in extra overhead for diminishing returns. The exact choice is not critical—the Lustre documentation simply states 'a file can be striped over as many [OSTs] as it takes to achieve the required peak aggregate bandwidth for that file'. A starting point is to use approximately '1 stripe per GB' to '1 stripe per 4 GB' of file size. Another is to 'load balance' by using an integral factor of the number of processors, for example one stripe per compute node so that each node gets one aggregator.

A large HDF5 file accessed by many (all) processors is obviously an example of the final category in the above list, so we should stripe it over some number of OSTs between 2 and 48 depending on the expected size. For example, a 100 ms simulation would be expected to generate a main dataset of $8\,\mathrm{B} \times 100 \times 3\,253\,316 \times 2 \approx 5.21\,\mathrm{GB}$, so c should probably be between 2 and about 6.

4.4.4 Cached writes

Regardless of the chunk shape used (see previous sections), Chaste writes to the HDF5 file every print time step of simulation time. Recall that one time step of data consumes about 50 MB on disk (eq. (4.1)). The Lustre file system on Archer has good sustained throughput as mentioned in the previous section (500 MB s^{-1} per OST), but performance for small writes is much lower. Archer states that 'Parallel file systems in general are typically optimized for high bandwidth: they work best with a small number of large, contiguous IO requests rather than a large number of small ones'. We should therefore expect several hundred 50 MB writes to show

worse performance than, say, one 40 GB write.

The chunk cache provided by HDF5 and described above might have been the answer, but it is currently not available when using the MPI-IO driver in write mode. My solution was to implement a cached mode in the `Hdf5DataWriter` class. The constructor now takes an argument specifying if the cache will be enabled. I chose this over making the cache switch-able for simplicity, as the latter would require extra logic like flushing the cache when switching. A new `std::vector` member variable acts as the cache, with a reserved size equal to $C_t \times N_n \times N_v$ where C_t is the size of a chunk in the time dimension (calculated by the new algorithm or set manually by an advanced user), N_n is the number of nodes owned by a process, and N_v is the number of variables (as above). No new public methods (to be called by the solver, postprocessor, etc.) were required. Instead, the existing methods were altered so that rather than writing directly to the HDF5 file they instead added to the cache vector. Once the number of elapsed print timesteps equals C_t, each process writes the contents of the cache to the HDF5 file. Because we are still using collective writes, the library then takes care of consolidating the data onto aggregators as normal.

4.4.5 Performance testing

In this section the HDF5 methods and techniques described in the preceding chapters will be distilled into three strategies. The performance of each of these will be measured in a benchmark, for both a small (8) and a medium (20) number of compute nodes to test parallel scaling, and for three values of stripe count (4, 24, and 42) to investigate I/O efficiency (section 4.4.3).

The first method used the former chunk shape (variable in time and spanning the other two dimensions) and the other two the new algorithm I devised. These three are:

1. Old-style chunks (section 4.4.1) of size 41 in the time dimension, i.e. chunk size $\{41, 3\,253\,316, 2\}$. 41 was chosen as it is the largest size that can be run, resulting in a chunk size of 2 134 175 296 B, or about 1.99 GiB.[7]

2. New-style chunks (section 4.4.2) of target size 1 MiB, specifically $\{151, 434, 2\}$ (1 048 544 B) in the benchmark described below.

3. As item 2, but with caching enabled (section 4.4.4).

The benchmark was a ventricles and torso bidomain-with-bath problem on 8 Archer compute nodes (192 cores). The PDE time step was 25 µs. In order to minimize the computational cost of generating the solution, the print time step was set to the PDE time step (0.025 ms), and the total run time was 150 print timesteps (3.75 ms). Most of my simulations have between 100 and 1000 time steps, so 150 steps is on the low end of this range. To include other types of data access, upstroke time map and APD map postprocessing was enabled, which are two separate steps that occur at the end of the simulation and involve reading from the main results dataset (recorded as PostProc). Finally, all datasets are converted to VTK visualization files (recorded as DataConversion).

The benchmark results are shown in fig. 4.5. The three stacked bars in each subfigure correspond to the three strategies, above. In each stacked bar the time spent (in minutes) in three categories are shown: Output (writing to disk), PostProc (performing post-processing), and DataConversion (converting to VTK visualization files). For all tests these three categories accounted for (71 ± 28) % of the total wall time. The three rows from top to bottom show results for a stripe count on the HDF5 file of 4, 24, and 42, respectively. The left and right columns show results from 8 and 20 compute nodes, respectively. Each simulation was performed three times in isolation to attempt to quantify the effect of machine load. The times

[7] Using 42 causes the chunk size to exceed 2 GiB and the program to fail. The reason for an apparent upper limit on the chunk size of 2 GiB rather than the expected 4 GiB is unknown.

shown are the means and standard error of the mean (SEM) of the three repeats. Performing more trials was unfeasible due to the time requirements.

The principal results of this study are as follows. In all cases, the old-style chunk shape (1) spent little time in Output and substantial time in DataConversion and PostProc. The former point might be due to the small number of chunks (just three) resulting in very little overhead when coordinating collective writes. The latter point illustrates the high cost of reading results from a poorly laid out dataset. Conversely, the new style (2) spent the vast majority of time in Output and little in DataConversion or PostProc. Writing a single row at a time over 7497 chunks is evidently slow, but the smaller, squarer chunks allow the postprocessing and VTK conversion steps to happen extremely quickly.

The effect of enabling the chunk cache on the new-style method is striking (3). In this case, the benefits of the new chunk shape to DataConversion and PostProc are retained and the time spent in Output is reduced to under 30 s in all cases. Clearly the new algorithm with cached writes offers superior performance on Archer compared to the alternatives.

Figure 4.5 also illustrates the relationships between the performance of each method with respect to node count and stripe count, which are of importance for performance optimization on Archer. Firstly, by comparing each row we can see the effect of the stripe count on each algorithm. Secondly, by comparing the two columns of subplots we can see the parallel scaling of each algorithm for a given stripe count on the HDF5 file.

First, comparing rows, for the old-style chunk shape (1), the time spent in DataConversion was greater with 24 or 42 stripes compared to 4 stripes with both 8 and 20 nodes. Clearly the DataConversion process is unable to leverage the extra bandwidth offered by the large stripe counts, probably either due to the overhead of communicating with many OSTs, or possibly due to a technical advantage from

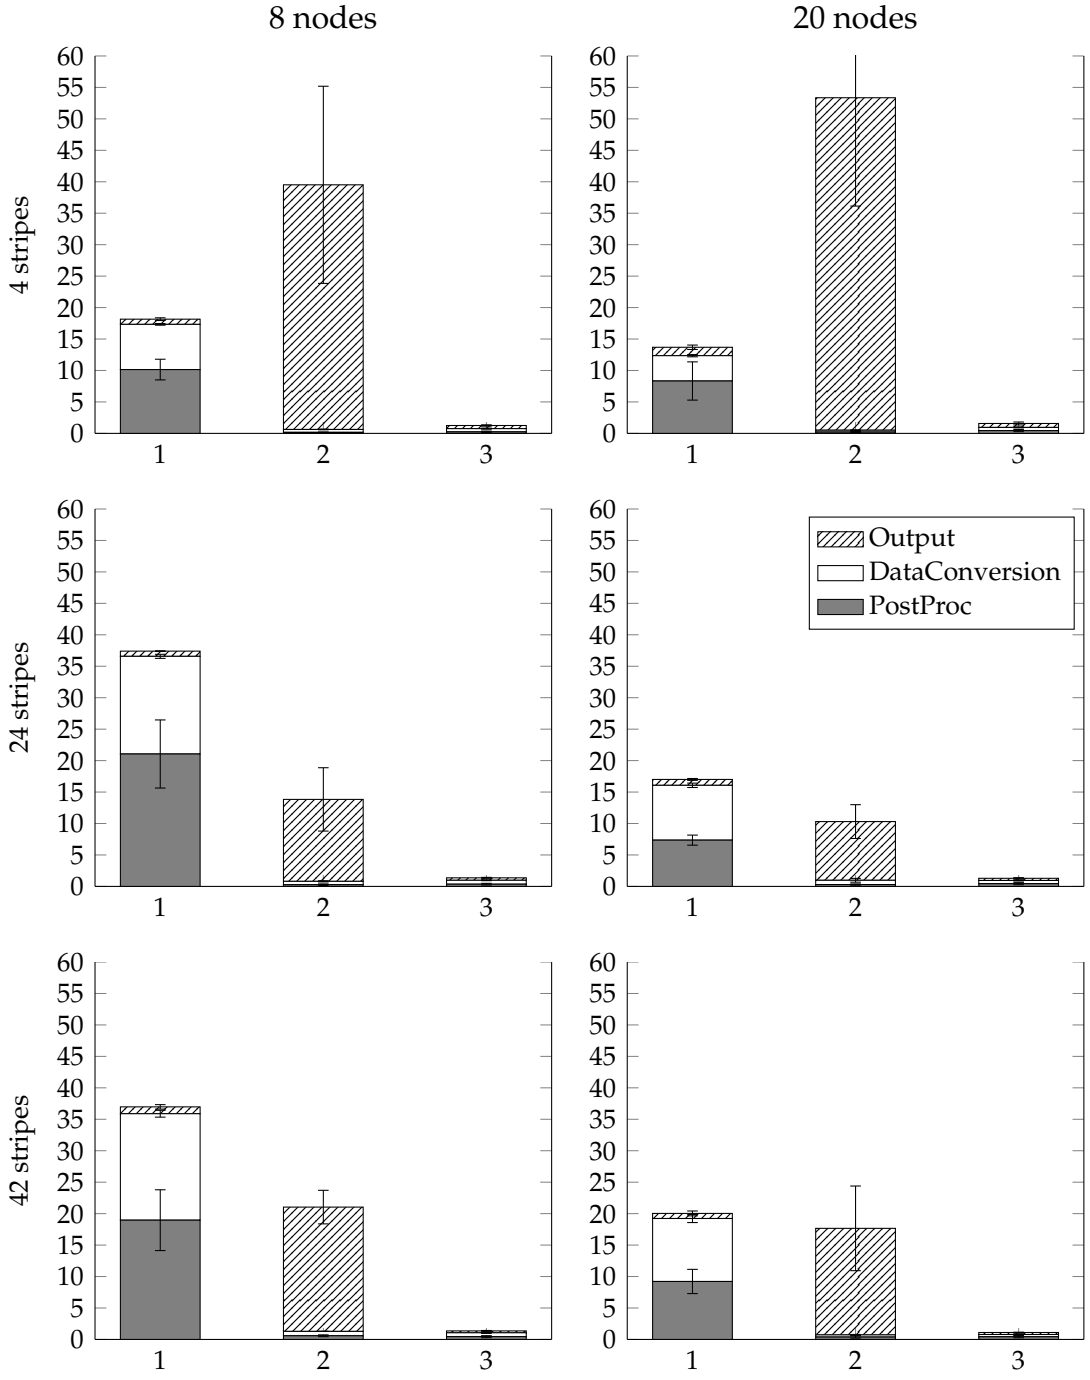

Figure 4.5: Time spent writing to disk (Output), converting to VTK visualisation files (DataConversion), and performing post-processing (PostProc) by each of the four methods: (1) old-style chunks with $t = 41$, (2) new-style chunks with a target size of 1 MiB, and (3) new-style chunks with caching. The two columns of figures are for simulations on 8 and 20 compute nodes, respectively. The three rows of figures are for simulations using stripe counts of 4, 24, and 42, respectively. Vertical axes are time in minutes. Times are means from two repeats. Error bars show the SEM.

concentrating on a small number of OSTs (e.g. internal caching within an OST). PostProc showed the same trend. In contrast, the uncached new shape (2) was faster with 24 or 42 stripes than 4, showing the performance benefit of parallel I/O. The severe bottleneck in Output is alleviated with a larger number of stripes, suggesting that the performance with 4 stripes is either limited by the OSTs or due to overwhelming the four threads assigned to aggregators. If we interpret the large error bars on Output as a sign of sensitivity to the machine load then the answer is probably the former. Finally, this method performed better with 24 stripes than 4 or 42, supporting the rule of thumb that a single large file written to by many processors should be striped, but not too much, to avoid incurring large overheads. There were no significant differences for the cached method (3) between stripe counts.

Second, comparing columns, performance is improved for the old-style chunk method (1) going from 8 to 20 nodes, but only DataConversion was materially faster. This suggests that despite the sub-optimal I/O performance of these chunk layouts, the postprocessing stage is fairly well able to utilize extra processors. The un-cached new algorithm (2), however, scales poorly at best, showing no significant difference between 8 and 20 nodes. In the case of 4 stripes, Output is still slow, possibly even exacerbated by the larger number of nodes. Again, there were no significant differences for the cached method (3) between the numbers of nodes.

Method	Output/s	DataConversion/s	PostProc/s
1 Old shape	49.5 ± 22.5	431.2 ± 13.8	608.7 ± 170.6
2 New shape	2333.5 ± 1628.8	27.5 ± 5.6	10.0 ± 5.1
3 Cached	29.9 ± 11.3	30.6 ± 8.7	14.6 ± 12.6

Table 4.3: Benchmark times shown in fig. 4.5 for 4 stripes on 8 compute nodes (top-left). Times are given in seconds (mean ± S.D.). The old chunk shape (1) is fast in Output but slow in the other two categories. The new shape (2) is the opposite. With caching of writes enabled on the new shape (3) the time spent in all three areas is low.

4.4.6 Chunk alignment on Lustre

Section 4.4.3 introduced Lustre and section 4.4.5 measured the effect of the HDF5 chunk shape and stripe count. Despite the excellent performance of the algorithm I introduced, there is a final small optimization to describe that combines the concepts of chunks and stripes.

As described in section 4.4.2, the chunking algorithm is unlikely to produce chunks of exactly the target size. Also, by default the chunks are located at irregular locations within the file. These statements together imply that a given chunk is unlikely to align with the stripe boundaries of the file. In such circumstances, accessing a chunk requires reading stripes from more than one OST. For example, if the stripe size and chunk target size are 1 MiB and the true chunk size is slightly less than 1 MiB, then reading a chunk is likely to involve requests to two OSTs.

It would be preferable when the chunk and stripe sizes are similar for each chunk to be padded slightly with empty space so that each chunk starts on a stripe boundary. Fortunately HDF5 has a command that does exactly this: `H5Pset_alignment`. Note that this function should be used with caution to avoid excessive wasted space.

For convenience I added a public method on the `AbstractCardiacProblem` class (from which the monodomain and bidomain problems inherit) that takes a numerical argument in bytes and sets the target chunk size and alignment parameters to it.[8]

The benchmark problem from section 4.4.5 was used to test the performance with and without alignment. The new algorithm with cache (method 4 in the previous section) was used for the unaligned case; the aligned case was identical apart from a call to `H5Pset_alignment(1048576)`. The results are shown in table 4.4. The number of repetitions for the unaligned and aligned cases were 110 and 113, re-

[8] `SetHdf5DataWriterTargetChunkSizeAndAlignment(hsize_t size)`

spectively. The mean and standard deviation (S.D.) are given in seconds, and statistical significance was calculated using the two-tailed Welch's *t*-test. Statistical power was calculated according to Cohen (1988) at the 5 % level.

For this problem, the aligned option is significantly slower than unaligned in Output, PostProc, and Total times. The explanation for this might be that the processes regularly read and write across chunk boundaries (i.e. the partition boundaries rarely fall exactly on chunk boundaries) so placing each chunk into its own stripe rarely results in a reduction in the number of OSTs accessed. Also, enabling alignment might introduce small gaps into otherwise contiguous data, reducing performance slightly. There was no significant difference for DataConversion, but this is probably a type II error since the statistical power was only about 7 %.

It should be noted that the absolute difference of several seconds is negligible in a simulation taking of the order of one hour. With a larger problem the theoretical advantage of aligning chunks to stripes might become apparent.

	Output	DataConversion	PostProc	Total
Unaligned	(20.7 ± 4.3) s	(30.2 ± 13.6) s	(7.5 ± 1.9) s	(195.0 ± 18.5) s
Aligned	(23.8 ± 12.5) s	(31.0 ± 13.8) s	(8.4 ± 2.6) s	(203.9 ± 23.9) s
p	0.013	0.648	0.004	0.002
Stat. power	0.701	0.074	0.828	0.869

Table 4.4: Timings (mean ± S.D.) for a test problem with and without HDF5 alignment set to the stripe size. The combined sample size was 223. Times are given in seconds. Welch's *t*-test was used to test for significant differences in the means for each category independently (p). The aligned method was significantly slower than unaligned at the 5 % level in Output, PostProc, and Total. The statistical power is over 70 % for these three categories.

4.5 CVODE solver stability

Time integration of the cell models at the mesh nodes is one of the two components of the main solve loop in Chaste (along with solving the linear system) and there-

fore a top target for optimization. Chaste provides various numerical schemes for doing this, mainly Runge-Kutta methods including RK4 and the backward Euler method. In 2012, support was added for CVODE, a 'solver for stiff and nonstiff ordinary differential equation (ODE) systems' utilizing 'variable-order, variable-step multistep methods' (Hindmarsh et al. 2005). CVODE is sophisticated and fast, but for about two years it was unreliable in the kind of large tissue simulations involved in this thesis. Furthermore, it displayed the worrying characteristic of producing different results depending on the *print timestep* (the interval at which results were sampled and saved to disk in a simulation).

Three improvements were made to overcome these issues. Firstly, we discovered that the CVODE solver in tissue simulations was allowed to extend its timestep to solve beyond the next PDE solve (the point at which the bidomain solution is updated). In these circumstances the following CVODE call would be from a time in its 'past', and with a different membrane voltage due to the PDE solve. We corrected this by setting the maximum timestep in tissue simulations to the PDE timestep, and utilizing `CVodeSetStopTime` to inform the solver of the desired time to solve to, i.e. the time of the next PDE solve. Secondly, the test I created to demonstrate the link between the results and the print timestep led to architectural improvements in the *timestepper* (the class that tracks and increments the simulation time). These two actions made CVODE sufficiently stable for regular use in tissue simulations. Thirdly, the rare occasions of CVODE failing to reach a solution are now detected, the solver state reset, and a new solve requested. These occasional failures are related to the sudden inconsistency between the solver state and the system state following a PDE solve step, which is an unavoidable consequence of iteratively solving the coupled systems of equations. Forcing the CVODE solver to start afresh enables the simulation to complete normally.

The performance advantage of CVODE over backward Euler was measured

in the same benchmark problem as section 4.4.5, but for a complete heartbeat (750 ms). The relative and absolute CVODE tolerances were 1×10^{-5} and 1×10^{-7}, respectively (Chaste defaults). The backward Euler ODE time step was set to 1.25 µs to achieve comparable accuracy to CVODE with the aforementioned tolerances in a single-cell test of the same cell model. No output was produced and no post-processing performed in order to best isolate the main 'solve loop'. The time spent in 'ODE' was almost 75 % less with CVODE than backward Euler, reducing the total wall time by over 50 % (fig. 4.6). Also, the superiority of CVODE is a function of the PDE timestep because of the equality of the PDE and CVODE maximum timesteps. This relationship will be beneficial in section 4.6. Note that 'KSP' refers to solving the bidomain equations, and 'AssRhs' refers to assembling the 'right hand side' of the linear system (see Pathmanathan et al. (2010, §3.3)).

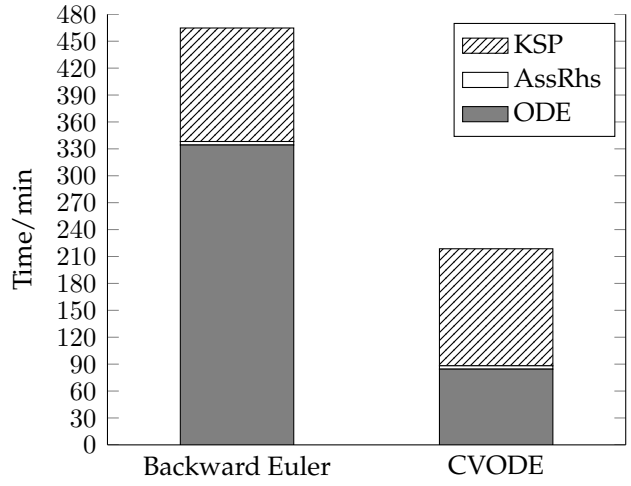

Figure 4.6: Timings for one heartbeat (750 ms) in a bidomain heart-torso simulation comparing ODE solver performance. 'KSP' refers to solving the bidomain equations, 'AssRhs' to assembling the 'right hand side' of the linear system, and 'ODE' to solving the cell model equations. Only categories with at least 1% of the total time are shown. The CVODE solver was almost four times faster than backward Euler in ODE (5065 s vs. 20 059 s), reducing the total wall time from 468 min to 222 min. The other categories were within 4% of each other (AssRhs=225 s vs. 233 s, KSP=7827 s vs. 7593 s).

4.6 Adaptive timestepper

Section 4.5 demonstrated the superior performance of CVODE over a commonly used alternative. In that section, the benchmark problem was entirely typical, including the PDE time step of 25 µs. This small time step is necessary while the activation wave is travelling through the tissue due to the steep voltage gradient at the wavefront. Once the tissue is activated, however, the electrical potentials throughout the domain are relatively smoothly varying, and we can increase the PDE time step without loss of accuracy in the body surface potentials.

Chaste provides an abstract 'time adaptivity controller' class that can be implemented by the user to alter the PDE time step dynamically. I developed one such controller that changes the time step once all myocardial nodes have undergone an upstroke (fig. 4.7). Rather than use a sophisticated error metric, this simple controller exploits the fact that the cell models are drastically less stiff once the tissue is entirely depolarized. It does this by starting with a small PDE time step (usually 25 µs). Each time step it examines the minimum V for all nodes. While the minimum is less than zero, small time steps are used. Once the minimum crosses zero, at least a further 5 ms of small timesteps are used to ensure all cell models are in the plateau phase and thus relatively slowly varying, until the simulation time is a multiple of the print timestep (usually 1 ms).[9] This ensures that all cells have passed the upstroke in the AP, and that each PDE timestep will fall on a factor of the print timestep. At this point the controller switches to large PDE timesteps (usually 1 ms).

The controller was later extended to return to small timesteps in simulations with multiple beats. If informed of a pacing time, the controller will reset its state and return to small timesteps when the simulation time is a multiple of the pacing

[9] Because all nodes are sampled for the minimum V, including bath nodes, once the tissue is entirely activated the minimum V actually equals 0, because V for bath nodes is always 0 for implementation reasons.

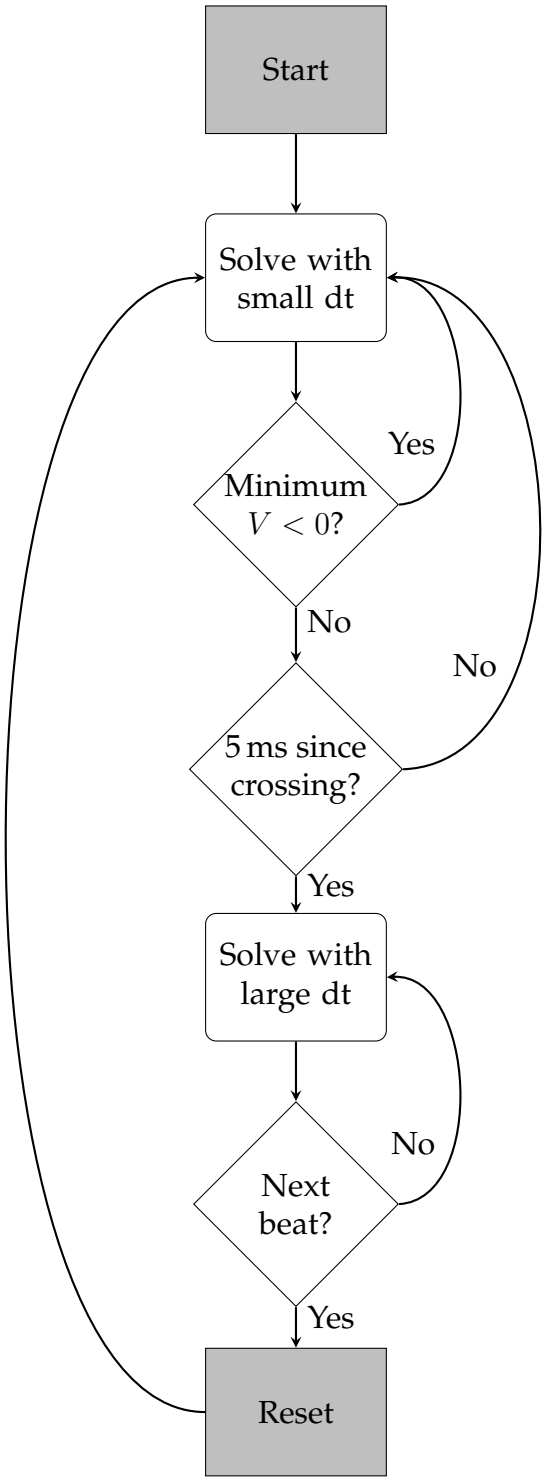

Figure 4.7: Flow chart of the "post-upstroke" time controller I developed to accelerate simulations of complete heartbeats. A small time step is used while any nodes are at resting potential ($V < 0$). Once all have crossed zero, 5 ms more of small time steps are used, then the controller switches to large timesteps. If the simulation time reaches a multiple of the pacing frequency the controller is reset to its initial state.

time.

To illustrate the performance increase with adaptive PDE timesteps the same benchmark was used as in section 4.5, switching to a larger dt of 1 ms or 0.5 ms. The results are shown in fig. 4.8. With the larger steps, performance in solving a complete heart beat is improved by a factor of about 4.5 when using the adaptive timestepper, the time spent in PDE was reduced by about 75 %, and the time spent in ODE was reduced (by 86 %). The 0.5 ms case was only about 3 % slower than the 1 ms one.

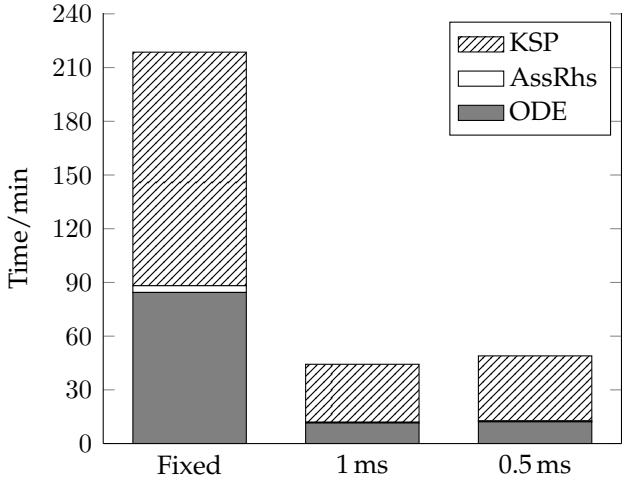

Figure 4.8: Timings for a bidomain simulation with a fixed or adaptive PDE timestep. The 'Fixed' bar is identical to the 'CVODE' bar in fig. 4.6, and the categories have the same meanings. The other bars show the time taken to perform the same 750 ms simulation but with either 1 ms or 0.5 ms PDE timesteps after complete activation of the ventricles (after about 61 ms of simulation time). Total wall time was reduced from about 221 min to about 50 min for the adaptive cases, a reduction of over 77 %. The 0.5 ms case took 3 % longer than the 1 ms case.

Simulated ECGs were examined for each scheme to assess convergence. The T wave of lead V4 is shown in fig. 4.9 for example. The use of the larger adaptive time step of 1 ms resulted in a small shift in the T wave peak. Cross-correlation was used to estimate the shift for all leads in this simulation as (1.18 ± 0.07) ms. With the smaller 0.5 ms adaptive time step the shift in T wave was negligible. Based on these and the above results, a large time step of 0.5 ms offers an excellent compro-

mise of accuracy and speed in this test case.

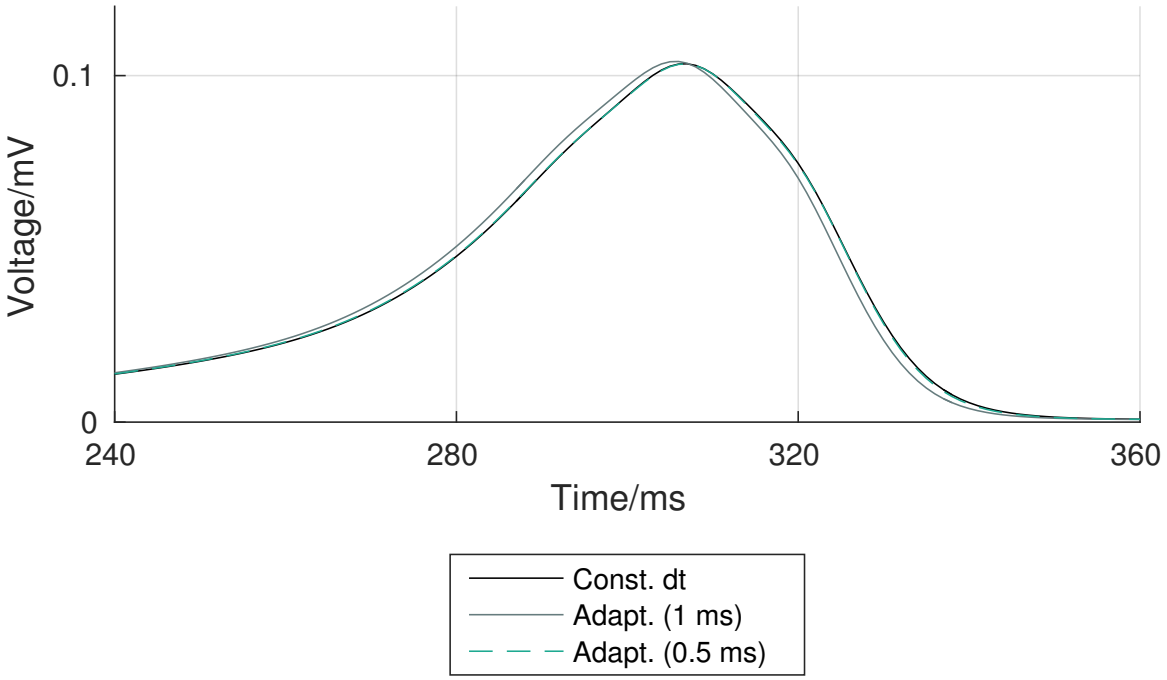

Figure 4.9: Detail of simulated body surface potential at the V4 electrode for the constant and adaptive time step schemes. The shapes are similar but the adaptive simulation with the larger dt of 1 ms slightly accelerates repolarization compared to the constant time step, shifting the T wave forward in time by about one millisecond. The difference between the constant and smaller adaptive time step (0.5 ms) was negligible. The other electrodes showed a similar difference.

4.7 Conclusion

Before the improvements described in this chapter, Chaste's computational efficiency in large (over one million node) simulations was inadequate. In particular, substantial fractions of the total computer time were spent in reading the mesh, writing results to disk, and inefficient numerical solvers. An increase in performance was required to enhance the usefulness of simulations to researchers and simultaneously conserve resources.

The main result of this work is in enabling large bidomain simulations, of the order of seconds of simulation time, to be performed in a tolerable amount of wall

time (typically under one hour). While such simulations themselves were not impossible before, performance analysis was typically done without considering the complete framework, e.g. by excluding post-processing in Bernabeu et al. (2014). Furthermore, the distribution of time accounted for in each of Chaste's categories continues to indicate good efficiency: typically about 83 % of wall time is spent solving, and only 14 % and 2 % spent in post-processing and set-up, respectively.

Chapter 5

Simulating the human ventricular activation sequence and electrocardiographic QRS complex

Chapter 4 detailed performance improvements made to Chaste when simulating large bidomain models with HPC, substantially accelerating the simulations performed for this chapter. This chapter describes techniques for modelling the activation sequence of the ventricles based on human data, and the generation of realistic QRS complexes in healthy and pathological cases. It presents a detailed investigation on how variability in the main determinants of the human activation sequence and passive myocardial and body conduction properties translate into variability in clinical QRS biomarkers.

5.1 Introduction

Since its invention by Einthoven at the beginning of the 20th century, the body-surface electrocardiogram (ECG) remains as the most extensively used clinical tool

for the non-invasive diagnosis of cardiac disorders. A specific challenge is to discriminate between the effects on the ECG biomarkers of the different components of the ventricular activation and repolarization sequences (involving simultaneous transmural, apico-basal, posterior-anterior and inter-ventricular propagation), and/or the contribution of concurrent pathological states. This hampers our ability to effectively extract the maximum information about the heart from the ECG.

In this context, multiscale human heart models are powerful platforms to integrate electrophysiological and structural information from the ionic to the whole organ levels, and to differentiate between key factors determining ECG biomarkers in disease (Wilhelms, Dossel, and Seemann 2011) or under pharmacological action (Zemzemi and Rodriguez 2015). Such an improved knowledge may further aid in the development of more selective biomarkers for specific diseased conditions or drug-induced risk stratification (Mincholé et al. 2015). Most computational studies to date, however, have mainly concentrated on replicating the repolarization sequence of the human ventricles and placed less attention on the activation sequence, often resulting in poorly recovered QRS complex morphology and lead polarities (see for example Wilhelms, Dossel, and Seemann (2011) and Zemzemi and Rodriguez (2015)). This highlights the need for accurate representations of the human activation sequence in ECG simulations, due to its direct influence in determining repolarization, and to properly address the computational study of cardiac conduction disorders.

In this work, we investigate the effect of variability in characteristics of the human ventricular activation sequence in the simulated QRS complex using a human torso-heart bidomain model. The simulated activation sequences are built and validated using experimentally derived human activation data, ranging from ex-vivo microelectrode recordings, non-invasive in-vivo electromechanical wave and electrocardiographic imaging, up to the body-surface ECG. Quantitative investi-

gations on variability in activation sequence and passive conduction properties are then presented to gain novel insights on the impact of tissue-level propagation in clinical QRS biomarkers, as well as to demonstrate the feasibility of our approach to model pathophysiology of the human conduction system. Our results can therefore have important implications in advancing the realistic modelling of cardiac conduction under healthy and diseased conditions, and aid in unravelling the role of variability in modulating response to therapy at the population level.

5.2 Methods

5.2.1 Experimental data

The present investigations build on three main sources of human experimental data. Firstly, ventricular activation sequences were based on ex-vivo microelectrode recordings by Durrer et al. (1970) (fig. 5.1A), in-vivo electromechanical wave imaging by Provost et al. (2013) (fig. 5.1B), and in-vivo epicardial activation sequences as reported by Ramanathan et al. (2006) through non-invasive electrocardiographic imaging (fig. 5.1C). Secondly, distributions of body-surface potentials (BSPs) during ventricular activation in healthy conditions were analysed from Taccardi (1963) (fig. 5.1D). Thirdly, 12-lead body-surface ECG signals (fig. 5.1E) in healthy and bundle branch block conditions were obtained from the PTB Diagnostic ECG Database (Bousseljot, Kreiseler, and Schnabel 2009), freely available in Physionet (Goldberger et al. 2000), whilst hemiblock ECGs (not present in the PTB Diagnostic ECG Database) were analysed from Elizari, Acunzo, and Ferreiro (2007).

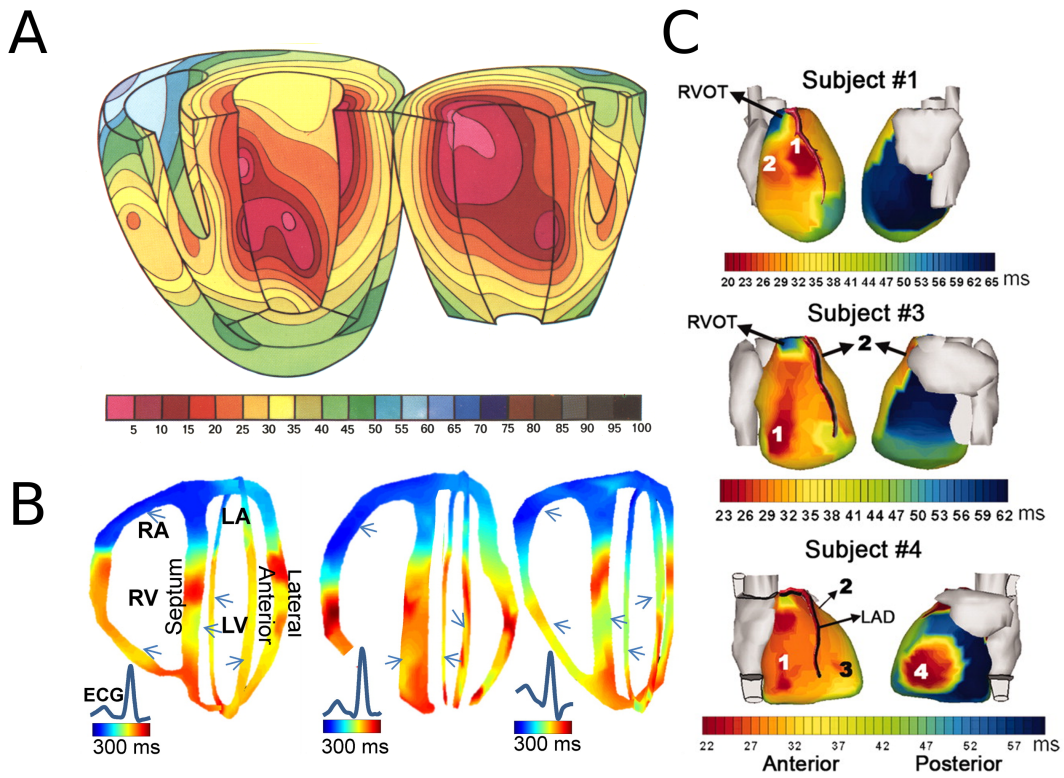
A
5 10 15 20 25 30 35 40 45 50 55 60 65 70 75 80 85 90 95 100
B
RA
LA
RV
Septum
Lateral
Anterior
LV
ECG
300 ms
300 ms
300 ms
C
Subject #1
RVOT
1
2
20 23 26 29 32 35 38 41 44 47 50 53 56 59 62 65 ms
Subject #3
RVOT
2
1
23 26 29 32 35 38 41 44 47 50 53 56 59 62 ms
Subject #4
2
LAD
1
3
4
22 27 32 37 42 47 52 57 ms
Anterior
Posterior
D
E
I
II
III
aVR
aVL
aVF
V1
V2
V3
V4
V5
V6

5.2.2 Activation system models

As in Keller, Kalayciyan, et al. (2010), and due to the variability and lack of high-resolution data on which to build anatomically detailed models of the human free-running Purkinje network, we assume that its branches couple with and excite the endocardial layer at several sites of earliest activation (root points) from which excitation quickly progresses. We base the positions of these root points on the notion of the human trifascicular conduction system (Rosenbaum et al. 1969) (left anterior-superior, left posterior-inferior and right lateral fascicular branches). To model the tightly packed endocardial Purkinje network, pairwise distances between all endocardial surface nodes were pre-processed using Dijkstra's algorithm. At runtime, each endocardial surface node is assigned a stimulus time proportional to the distance to its closest root node. The constant of proportionality allows a straightforward control of conduction velocity in the abstracted endocardial Purkinje network. This was adjusted to yield a distribution of endocardial activation times in accordance with the reported ex-vivo microelectrode recordings by Durrer et al. (1970). The implementation is illustrated by example in fig. 5.2.

Figure 5.1 *(preceding page)*: Non-diseased human ventricular activation data considered in this study. (A) Ex-vivo microelectrode recordings of activation sequence in a healthy subject. (B) Non-invasive in-vivo electromechanic wave imaging isochrones of three normal subjects. Arrows indicate ventricular sites of early activation. (C) Non-invasive in-vivo electrocardiographic imaging reconstruction of epicardial activation sequence in three healthy subjects. (D) Body surface potential maps in a healthy subject during progression of ventricular activation. (E) QRS complexes in the 12-lead ECG of a representative healthy individual from the PTB Diagnostic ECG Database (Bousseljot, Kreiseler, and Schnabel 2009, subject #198). Clinical ECG grid resolution: 40 ms/0.1 mV. Panels A to D reproduced with permission from Durrer et al. (1970) (Wolters Kluwer Health, Inc.), Provost et al. (2013) (Elsevier), Ramanathan et al. (2006) (© 2006 National Academy of Sciences), and Taccardi (1963) (Wolters Kluwer Health, Inc.), respectively.

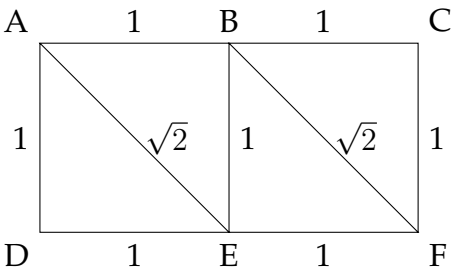

(a) Example triangular mesh. Six nodes A–F are connected by nine edges labelled with their edge lengths.

	A	B	C	D	E	F
A	0.0	1.0	2.0	1.0	1.4	2.4
B	1.0	0.0	1.0	2.0	1.0	1.4
C	2.0	1.0	0.0	3.0	2.0	1.0
D	1.0	2.0	3.0	0.0	1.0	2.0
E	1.4	1.0	2.0	1.0	0.0	1.0
F	2.4	1.4	1.0	2.0	1.0	0.0

i \ j	0	1	2	3	4	5
0	-	-	-	-	-	-
1	0	-	-	-	-	-
2	1	2	-	-	-	-
3	3	4	5	-	-	-
4	6	7	8	9	-	-
5	10	11	12	13	14	-

(b) Table of pairwise minimum distance to 2 s.f. (left), and the corresponding serialized indices in the distance file (right).

Figure 5.2: An example to illustrate the graph-based activation model. (a) A triangular mesh. (b) (Left) Table of shortest distances between nodes in the example mesh in (a). The table is symmetrical because the edges are undirected, so only one triangle is saved to disk. If we save only the lower triangle to disk (highlighted) and serialize row-by-row, we get the indices shown on the right. (Right) The indices of the distances corresponding to the shaded area in the serialized distances file. The nodes are indexed in alphabetical order. Given node indices i and j ($j < i$) the serialized index is $\left(\sum_{n=1}^{i-1} n\right) + j = (i(i-1))/2 + j$.

5.3 Results

5.3.1 Model construction and validation

Human ventricular activation data exhibit a high level of inter-subject variability in the sites of earliest activation and epicardial breakthrough, as illustrated in the multi-modality data comparison summarized in fig. 5.1. The following commonalities were however identified across the studied human datasets: (1) activation within the left ventricle (LV) usually begins on the basal anterior paraseptal, the mid septum and the posterior apex regions, and then progresses transmurally (fig. 5.1A–B); (2) earliest activation in the free wall of the right ventricle (RV) occurs 'near the insertion of the anterior papillary muscle' (Durrer et al. 1970), the position of which is known to be highly variable (fig. 5.1B–C); and (3) activation of the RV epicardium usually occurs before that of the LV epicardium (fig. 5.1B–C).

Based on the above, we designed a baseline configuration of the human activation sequence (see fig. 5.3A, black arrows) consisting of four LV earliest activation sites (LV mid septum, LV basal anterior paraseptal, and two LV mid-posterior) and three in the RV (RV mid septum, two RV free wall). Such a configuration allows for a multiscale investigation (up to the body-surface ECG) of inter-subject variability in activation speed and passive tissue conductivities, as well as in the position and number of earliest activation sites, without favouring a specific data modality in our analysis.

Simulated ventricular activation times for the baseline configuration are presented in fig. 5.3A. The LV endocardial surface is fully activated within 30 ms, transmural propagation occurs in 35 ms in the LV basal free wall, with transseptal propagation of 25 ms. This is in accordance with the ex-vivo microelectrode recordings by Durrer et al. (1970) (fig. 5.1A), reporting 35 ms in endocardial LV activation, and 25 ms and 35 ms in transseptal and transmural propagation, respec-

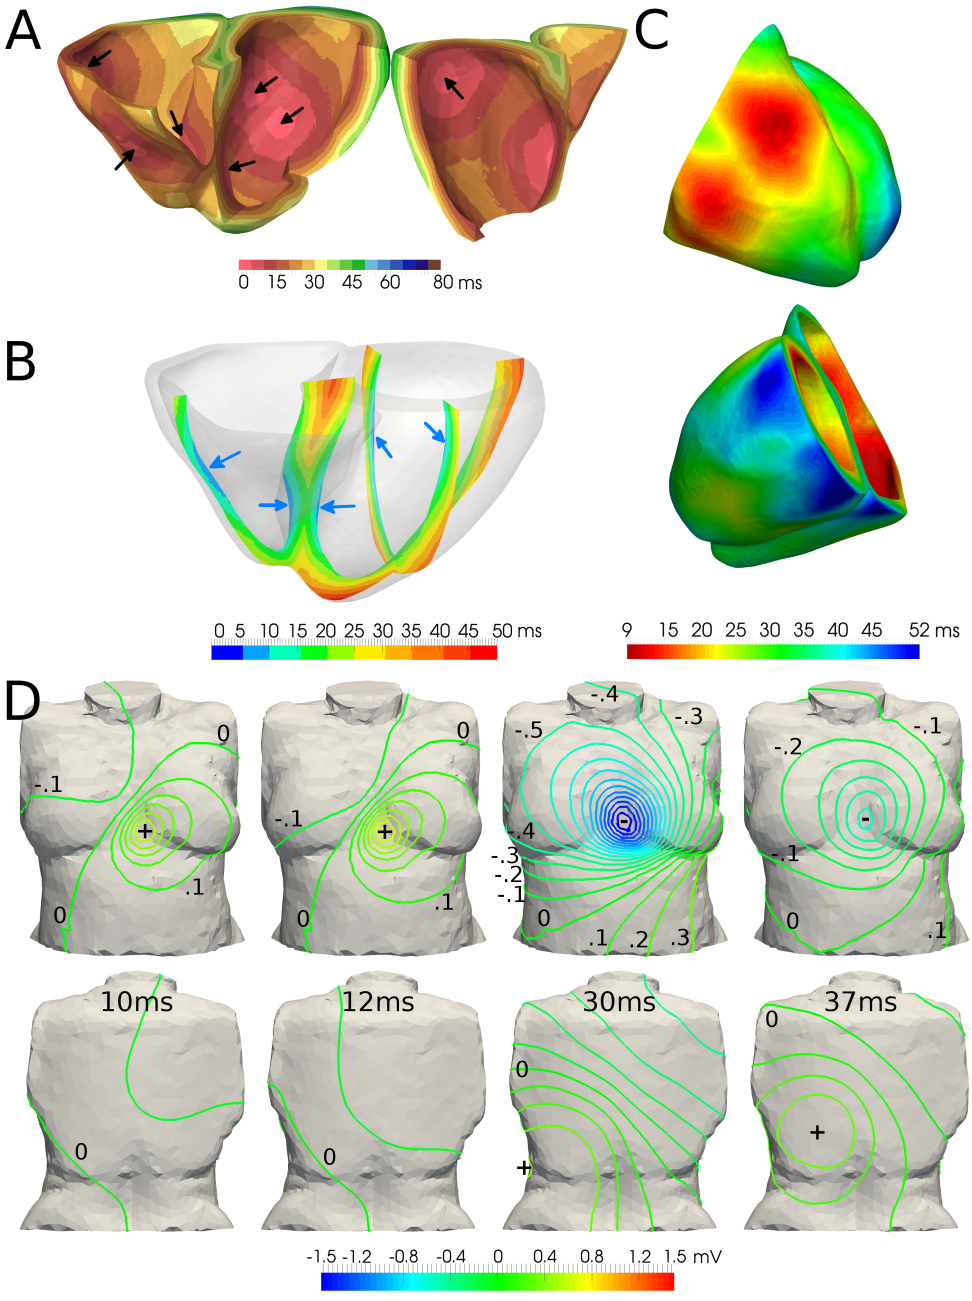
A
0 15 30 45 60 80 ms
C
B
0 5 10 15 20 25 30 35 40 45 50 ms
9 15 20 25 30 35 40 45 52 ms
D
10ms
12ms
30ms
37ms
-1.5 -1.2 -0.8 -0.4 0 0.4 0.8 1.2 1.5 mV
E
I
II
III
aVR
aVL
aVF
V1
V2
V3
V4
V5
V6

tively. Slightly shorter activation times are obtained in the RV base in simulations, possibly due to uncertainties in the atria-ventricular segmentation.

Cross-sectional slices of ventricular activation times are shown in fig. 5.3B for comparison with in-vivo electromechanical wave imaging data by Provost et al. (2013) (fig. 5.1B), indicating agreement in the location of earliest and latest activation sites. Specifically, early activation occurred on either side of the septum, RV free wall, and anterior and posterior LV paraseptal sites. Latest activation took place in the basal septum, basal lateral LV wall, and apex.

Anterior and posterior views of epicardial activation times are provided in fig. 5.3C for ease of comparison with the in-vivo electrocardiographic imaging data of Ramanathan et al. (2006). Simulation results are in close agreement with subjects 1 and 3 in fig. 5.1C, with earliest breakthrough in the upper paraseptal anterior RV, apical LV activation 20 ms later, and basal LV activation 28 ms after that. The activation of the RV epicardium completed in 30 ms, compared to 25 ms in healthy subjects. The LV epicardial breakthrough occurred about 15 ms after the RV breakthrough, consistent with both Subjects 1 and 3. The pattern of LV epicardial activation is well-reproduced due to the lateral basal LV root point (leading to the epicardial breakthrough on the LV anterior paraseptal region) and the thickening of the LV free wall from apex to base (which generates activation times that lengthen from apex to base). Figure 5.3D shows isopotential maps of BSPs, exhibiting the same spatiotemporal evolution during ventricular activation as those reported in

Figure 5.3 *(preceding page)*: Simulation results of the healthy human ventricular activation sequence. (A) Volumetric isochrones of ventricular activation times. Arrows indicate activation system root points. (B) Cross-sectional slices of ventricular activation times. Arrows indicate ventricular sites of early activation. (C) Anterior and posterior views of epicardial activation times. Different colour maps are shown in panels A–C to facilitate the comparison with experimental data in fig. 5.1. (D) Isopotential body surface potential maps during progression of ventricular activation. (E) QRS complexes in the simulated 12-lead ECG. Simulated ECG grid resolution: 40 ms/0.1 mV.

healthy subjects by Taccardi (1963) (fig. 5.1D). Initially a local maximum appears on the chest due to the first RV epicardial breakthrough (10 ms). Once the activation of the RV is completed, this is replaced by a negative region. The positive pole moves leftwards and downwards as activation progresses transmurally through the LV. This forms the characteristic dipole pattern on the front of the chest with a zero isopotential between the lower-right flank and the left shoulder (20 ms to 30 ms). The maximum is finally drawn towards the latest area to activate, the posterior paraseptal LV (40 ms).

As a consequence of the distributions of BSPs described above, a frontal plane QRS axis (average of global ventricular depolarization) between −30° and 90° is considered to be normal, where 0° is leftwards in the transverse plane. This determines the limb (I, II, III) and augmented (aVR, aVL, aVF) leads, and yields negative precordial leads in V1, positive in V4 to V6, with R-wave progression in between, as illustrated in fig. 5.1E for a representative healthy subject. For comparison, fig. 5.3E illustrates the simulated QRS complexes associated with our baseline activation sequence, exhibiting a QRS width of 70 ms within the normal range (60 ms to 100 ms) and a QRS axis of 60°, both within the healthy range, and typical lead polarities in all limb, augmented limb and precordial leads, as well as R-wave progression in the latter.

5.3.2 Variability in subendocardial activation speed and myocardial conductivities modulate QRS width and amplitude

Detailed simulation studies were conducted to investigate how variability in the different determinants of distributions of BSPs modulates the human QRS complex. Whilst the extent of this variability in the human population is difficult to determine, we hypothesize that variability of ±50 % in activation sequence and tissue properties encapsulates a major fraction of the observed clinical variability in

the healthy human QRS complex (Muszkiewicz et al. 2016). Changes on baseline QRS biomarkers are hence presented in the following sections (averages reported across all leads) for these bounds of plausible physiological variability. In fact, this could be a modest estimate of total variability across the population, since variability bounds from 2 to 12-fold have been reported in tissue conductivities in human (Keller, Weber, et al. 2010), including blood (×2.3), lungs (×3.4), fat (×5.5), soft muscle (×7.5), heart (×9.0), and bone (×12).

Variability in endocardial activation speed directly impacts the full ventricular activation sequence (fig. 5.4A, left panel), affecting the width and amplitude of the reconstructed QRS complexes. Increased activation speeds (fig. 5.4A, solid red traces) resulted in earlier R-wave times-to-peak, shorter QRS widths, and increased R-wave amplitudes ((75 ± 3) %, (90 ± 4) % and (125 ± 30) % of baseline, respectively). Decreased activation speeds had the opposite effects (fig. 5.4A, solid blue traces), exhibiting protracted R-wave times-to-peak, wider QRS widths, together with decreased R-wave amplitudes ((178 ± 7) %, (182 ± 12) % and (83 ± 8) % of baseline, respectively).

Variability in intracellular myocardial conductivities (as modelled by globally scaled conductivity tensors) substantially influences the transmural dispersion of ventricular activation times (fig. 5.4B, left panel). This yielded a predominant scaling effect on R-wave amplitudes ((121 ± 13) % and (65 ± 12) % of baseline for increased/decreased intracellular conductivities, respectively), with a smaller impact on times to R-peak ((95 ± 3) % and (105 ± 3) %, respectively). Wider QRS widths ((127 ± 7) %) and less symmetric QRS complexes were also associated with conditions of decreased intracellular coupling, in particular in the precordial leads (fig. 5.4B, solid blue traces).

On the other hand, variability in extracellular myocardial conductivities has a purely scaling role on the QRS complex, with negligible effects on the ventricu-

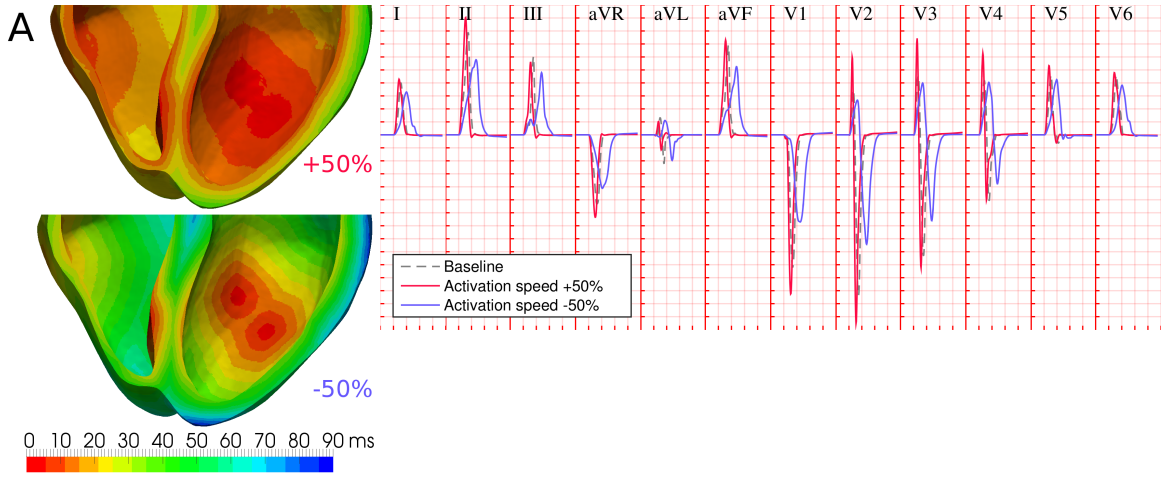
A
+50%
-50%
0 10 20 30 40 50 60 70 80 90 ms
I II III aVR aVL aVF V1 V2 V3 V4 V5 V6
Baseline
Activation speed +50%
Activation speed -50%

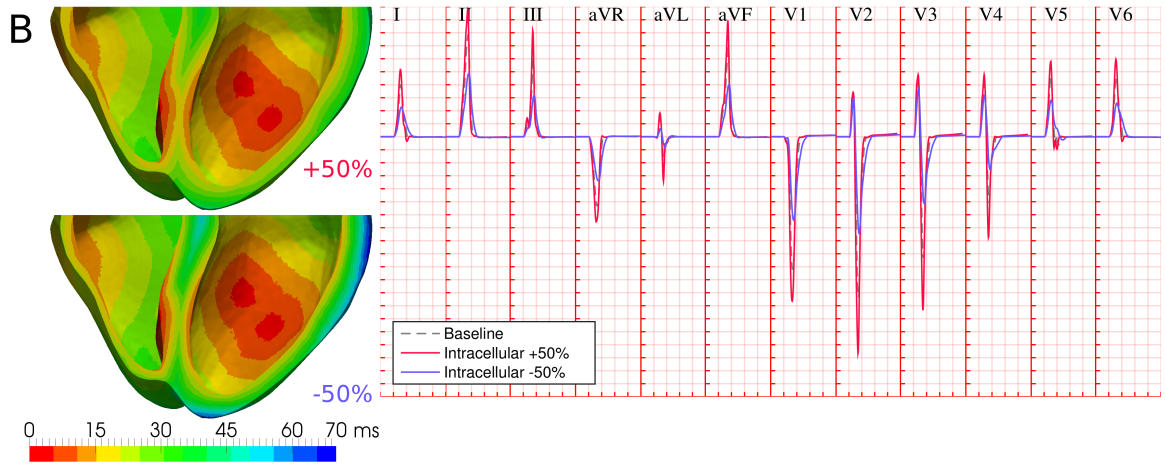
B
+50%
-50%
0 15 30 45 60 70 ms
I II III aVR aVL aVF V1 V2 V3 V4 V5 V6
Baseline
Intracellular +50%
Intracellular -50%

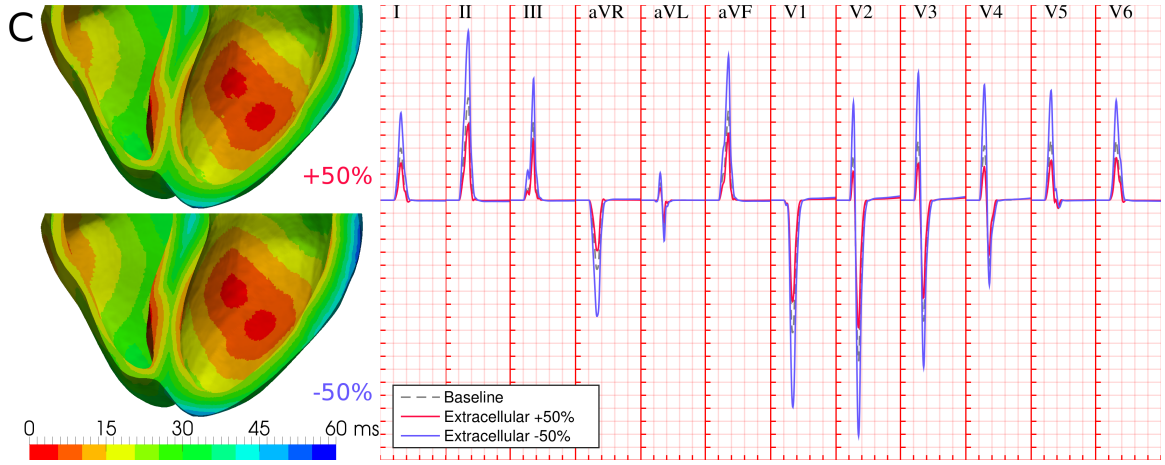
C
+50%
-50%
0 15 30 45 60 ms
I II III aVR aVL aVF V1 V2 V3 V4 V5 V6
Baseline
Extracellular +50%
Extracellular -50%

lar activation sequence (fig. 5.4C). Across all leads, peak R-wave amplitudes were $(69 \pm 6)\,\%$ and $(183 \pm 25)\,\%$ of baseline for a $50\,\%$ increase and decrease in extracellular conductivity tensors, respectively. Times to R-peak were within $2\,\%$ in all cases, with negligible changes to QRS widths.

The observed effects can be explained by recalling a simplified core-conductor model of electrical propagation in myocardial fibers (Plonsey and Malmivuo 1995), with extracellular potentials given by

$$\phi_e\left(t, \vec{x}\right) \propto \frac{\sigma_i}{\sigma_e} \int_{\Omega} \nabla \phi_i\left(t\right) \cdot \nabla \left(\frac{1}{\vec{r}}\right) \; d\Omega \tag{5.1}$$

where: ϕ_e is the extracellular potential at time t and body surface point $\vec{x}$; σ_i and σ_e are the intracellular and extracellular conductivities, respectively; and the integral represents the sum of the contribution of all gradients in intracellular potentials ($\nabla \phi_i$) in the volume of the heart (Ω), weighted by their respective distances to the measure point ($\vec{r}$). Changes in activation speed therefore affect the time course of endocardial and transmural gradients in intracellular potentials, hastening or protracting ventricular depolarization as reflected in the QRS width and times to R-peak, while their spatial distribution (more spread gradients for larger activation speeds) impacts R-wave amplitudes (fig. 5.4A). In contrast, variability in intracellular conductivities under the imposition of the same endocardial activation

Figure 5.4 *(preceding page)*: Variability in activation speed and myocardial conductivities modulate QRS width and amplitude. (A) Variability in activation speed impacts the ventricular activation sequence, affecting R-wave times-to-peak, R-wave amplitudes, and QRS width. (B) Variability in intracellular myocardial conductivities influences the transmural dispersion of ventricular activation times, yielding scaled R-wave amplitudes and wider and less-symmetric QRS complexes under conditions of decreased intracellular coupling. (C) Variability in extracellular myocardial conductivity yield a scaling effects on QRS amplitudes. In all cases, dashed grey ECG traces represent baseline conditions, whereas solid red and blue ECG traces represent conditions of increased and decreased myocardial parameters, respectively. Simulated ECG grid resolutions: 40 ms/0.1 mV.

sequence only modulates the transmural propagation of intracellular potentials, hence impacting the terminal part of the QRS complex with moderate influence on times to R-peak, together with a proportional scaling of extracellular potentials (fig. 5.4B). Correspondingly, extracellular potentials are inversely proportional to extracellular conductivities, modulating R-wave amplitudes with negligible effects on intracellular gradients (fig. 5.4C). Changes in QRS width by myocardial conductivities can also be interpreted by considering the effective conductivity tensor, $\sigma_m = \sigma_i\sigma_e/(\sigma_i + \sigma_e)$, and its impact on myocardial propagation speed, given by the square root of the quotient between altered and baseline effective conductivities. For our choice of tissue properties, variations of +50 % and −50 % in intracellular conductivities respectively yield changes of +20 % and −27.5 % in transmural propagation speed, yet only of +1.7 % and −4.6 % for equivalent variations in the extracellular tensor, which therefore explains the minor contribution of the latter to the width of the QRS complex.

5.3.3 Variability in body conductivities modulates QRS amplitudes

Similar to the role of extracellular myocardial conductivities, variability in body conductivities (torso/lungs/bones) exerted a scaling effect on QRS amplitudes by affecting the magnitudes of BSPs. Results on the effects of this variability in torso conductivity are illustrated in fig. 5.5, as this exhibited the largest contribution in the human QRS complex compared to lung and bone conductivities (fig. 5.6).

Although the spatial patterns of BSPs remained mostly unaltered under variations of body conduction properties, the larger diffusion of extracellular potentials associated with larger conductivities yielded BSPs with lower magnitudes at times for R-peak, whilst smaller body conductivities accentuated the magnitude of the observed gradients (fig. 5.5, left panel). This was hence translated in inversely pro-

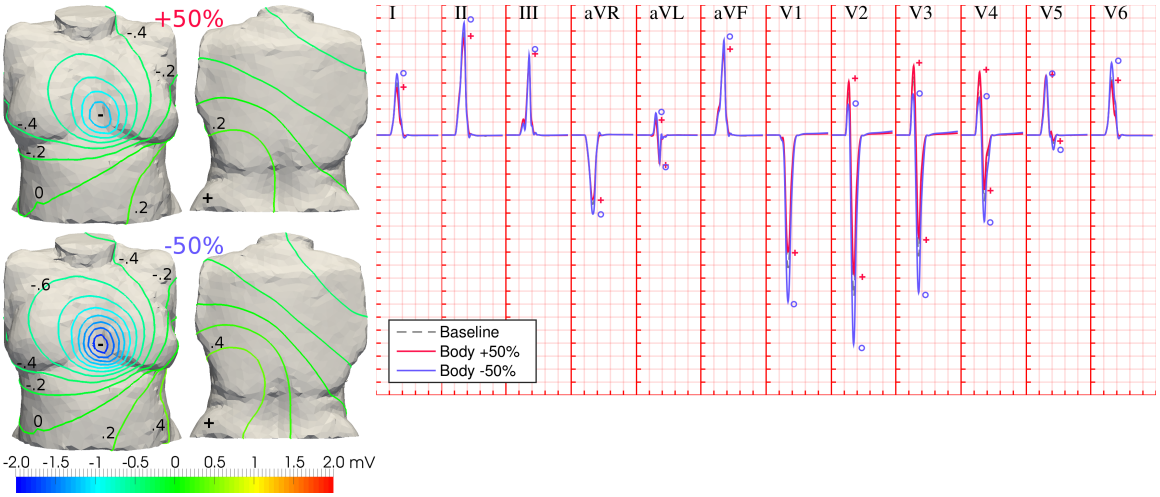

Figure 5.5: Variability in body conductivities modulates QRS amplitudes. Differences in the diffusion of electrical potential across the human torso affect the magnitudes of BSPs, without significantly altering their spatial pattern (left panel; BSP distributions shown at time for R-peak in lead II). This is translated in inversely proportional QRS amplitudes to body conductivities, and increased R-wave progression (right panel). Dashed grey ECG traces: baseline conditions; solid red and blue ECG traces: increased and decreased torso conductivity, respectively. Symbols indicate maximum wave amplitudes under variability in torso conductivity (same colour code). Simulated ECG grid resolution: 40 ms/0.1 mV.

portional QRS amplitudes to body conductivities in all leads (R-peak amplitudes of (92 ± 6) % and (112 ± 10) % of baseline for increased/decreased torso conductivities, respectively). In the precordial leads, reduced body conductivities increased R-wave progression (fig. 5.5, solid blue traces), and vice versa.

5.3.4 Variability in location of endocardial root points affects S-wave progression

Additional simulation studies were performed in order to evaluate how anatomical variability in the location of endocardial activation root points (at coupling sites of the free-running Purkinje system) impacts the human QRS complex. Eight activation sequences were designed based on the anatomical variability reported on the human trifascicular system (Elizari, Acunzo, and Ferreiro 2007), as well as the

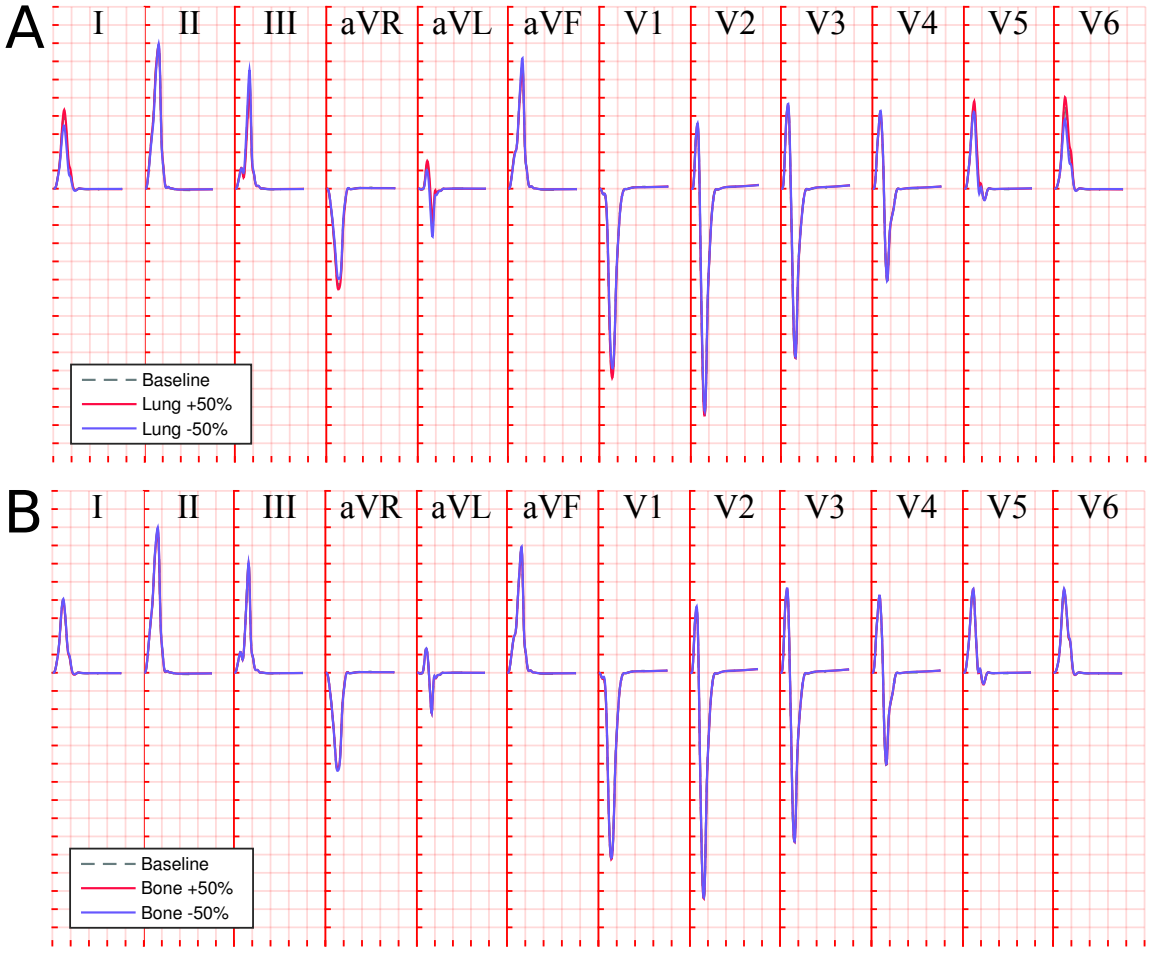

Figure 5.6: Effects of variability in lung and bone conductivities on QRS amplitudes. A: Lung conductivity. B: Bone conductivity. Dashed grey ECG traces: baseline conditions; solid red and blue ECG traces: increased and decreased conductivities, respectively. ECG grid resolution: 40 ms/0.1 mV.

early activation sites in the ex-vivo microelectrode studies by Durrer et al. (1970) and in-vivo electromechanical wave imaging by Provost et al. (2013). These included six activation sequences altering the position of LV anterior and posterior coupling sites towards more apical and basal locations, and two for more apical and/or lateral RV bundle coupling sites.

Variability in the location of early activation on the anterior LV wall, from more basal towards more apical positions (fig. 5.7A, top to bottom), was associated with increasing S-wave progression in the limb and augmented limb leads (fig. 5.7A, blue arrows). This progression was especially marked in leads II, III and aVF,

exhibiting increasingly more negative S-waves. Variability in the apico-basal location of the LV posterior wall site was similarly correlated with increasing S-wave progression in precordial leads V5 and V6 (fig. 5.7B, red arrows), as these are the only electrodes with a solid angle view of the LV posterior wall. Some of the activation sequences also presented notched QRS complexes in precordial and limb leads (fig. 5.7A–B, black arrows), as a consequence of variability in the positions of simultaneous LV earliest epicardial breakthroughs. Such QRS features were also observed in healthy individuals of the PTB Diagnostic ECG Database, as shown in the two ECG excerpts presented in fig. 5.7C. On the other hand, activation sequences with more apical and/or lateral RV sites led to similar QRS morphologies compared to baseline (fig. 5.8). Simultaneous variations of more than one LV/RV coupling sites led to an additive effect of the above discussed contributions.

5.3.5 Knockout of root nodes reproduce QRS phenotype caused by intraventricular conduction defects

Due to their direct correspondence with the trifascicular activation system, variability in the number of root nodes can be used to reproduce pathophysiology of the human conduction system. This is illustrated in figs. 5.9 and 5.10 for LV anterior/posterior hemiblocks and complete LV/RV bundle branch blocks, respectively.

Knocking out the activation of the LV anterior wall (fig. 5.9A, black circle) resulted in the late activation of the LV basal free wall. While the negative pole on the chest was almost unvaried, this late region caused the positive pole in BSPs to move up from the left hip to the centre of the back, yielding a late QRS vector mainly oriented from front to rear torso surfaces and slightly leftwards and upwards. As a consequence, the reconstructed QRS complex exhibits the clinical manifestations of a left anterior fascicular block (Elizari, Acunzo, and Ferreiro 2007) (LAFB; lower-right panel of fig. 5.9A): QRS width smaller than 120 ms; left-


A
B
C
I II III aVR aVL aVF V1 V2 V3 V4 V5 V6

deviated late QRS axis to −50°; qR pattern in aVL and RS/Rs patterns in leads II, III and aVF, with wave amplitudes satisfying RII>RIII and SIII>SII; and delayed intrinsicoid deflection (time from Q to R peaks) in aVL of 42 ms, close to the clinical threshold of 45 ms.

Similarly, the knockout of LV posterior wall activation (fig. 5.9B, black circle) was associated with a delayed activation of the posterior LV basal wall, severely affecting the time course and position of the poles in BSPs on the chest. Such a distribution of BSPs recovered the clinical phenotype of a left posterior fascicular block (Elizari, Acunzo, and Ferreiro 2007) (LPFB; lower-right panel of fig. 5.9B): QRS width smaller than 120 ms; right-deviated late QRS axis to 100°; Rs or RS patterns in leads I and aVL; qR patterns in leads III and aVF; and S waves in all precordial leads, with intrinsicoid deflection times in V6 (24 ms) and aVF (34 ms) greater than in aVL (21 ms). The less-symmetric (trapezoidal-like) morphology in the last deflection of the QRS complex of all limb and augmented limb leads observed in the clinical recordings is also recovered.

Finally, a complete knockout of both LV anterior and posterior branches led to a lateral progression of LV activation, significantly delaying septal to basal LV stimulation (fig. 5.10A). This retrieved the characteristic signature of left bundle branch block (LBBB) in precordial leads: QRS complex of at least 120 ms; marked QS pattern in V1; and notched R-wave in lead V6. Similarly, the knockout of cou-

Figure 5.7 *(preceding page)*: Variability in the location of LV endocardial activation sites affects S-wave progression. (A) More apical coupling sites of the LV anterior fascicular branch were associated with increasingly more negative S-waves in leads II, III and aVF (blue arrows). (B) Similar trends were present for a more basal coupling of the LV posterior fascicular branch in precordial leads V5 and V6 (red arrows). Variability in LV stimulation also resulted in notched QRS complexes in precordial and limb leads (black arrows). (C) QRS complexes of two healthy individuals in the PTB Diagnostic ECG Database (Bousseljot, Kreiseler, and Schnabel 2009) (subjects 237 and 117, respectively), exhibiting similar QRS features. Simulated and clinical ECG grid resolutions: 40 ms/0.1 mV.

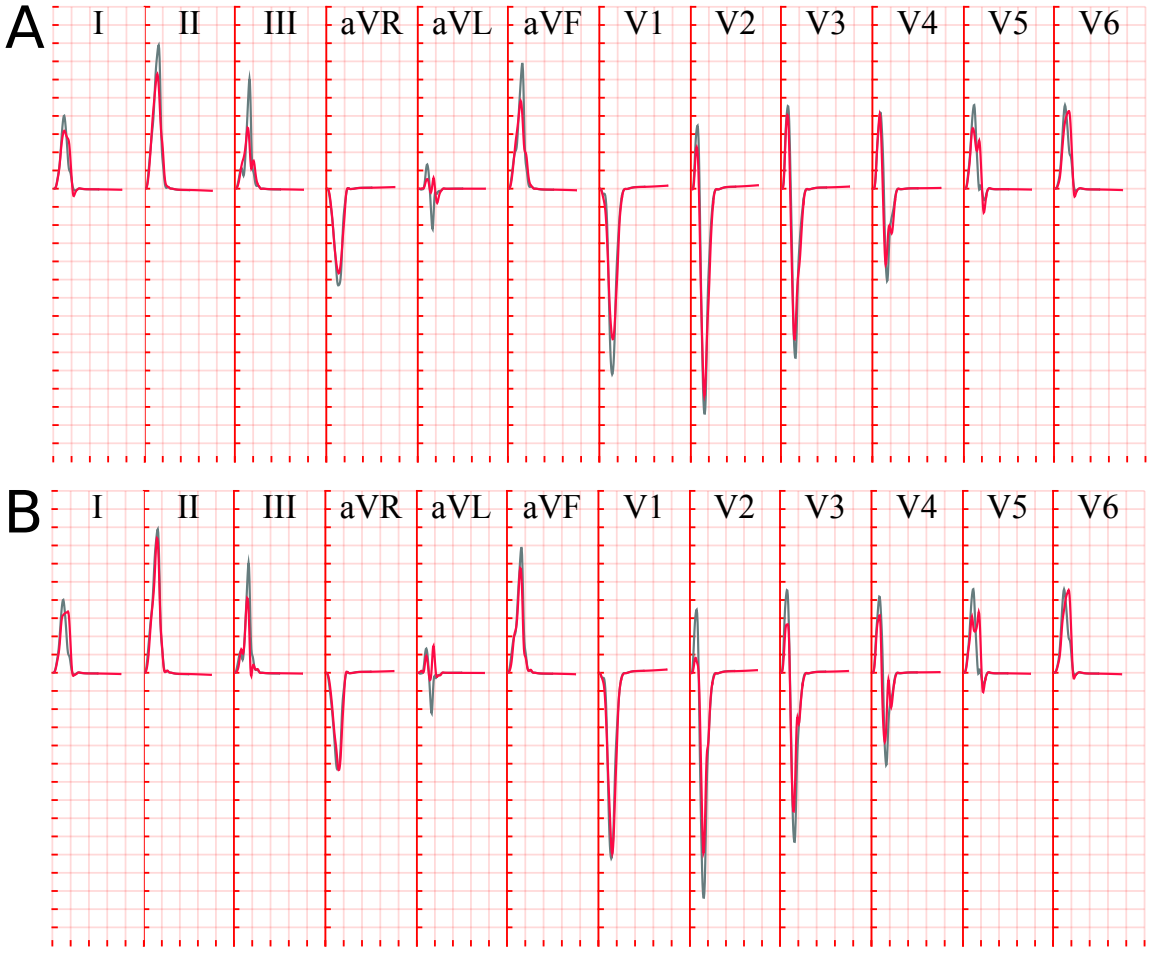

Figure 5.8: Variability in the location of RV free-running Purkinje coupling sites. A: More apical RV anterior activation site. B: More apical RV posterior activation site. ECG grid resolution: 40 ms/0.1 mV.

pling sites on the RV led to a lateral sequence of RV activation (fig. 5.10B), yielding the manifestation of right bundle branch block (RBBB): QRS duration of at least 120 ms; terminal R-wave in lead V1; slurred S-waves in leads I and V6.

5.4 Discussion

This study presents a detailed investigation on how variability in the main determinants of the human activation sequence and passive myocardial and body conduction properties translates into variability in clinical QRS biomarkers. Our

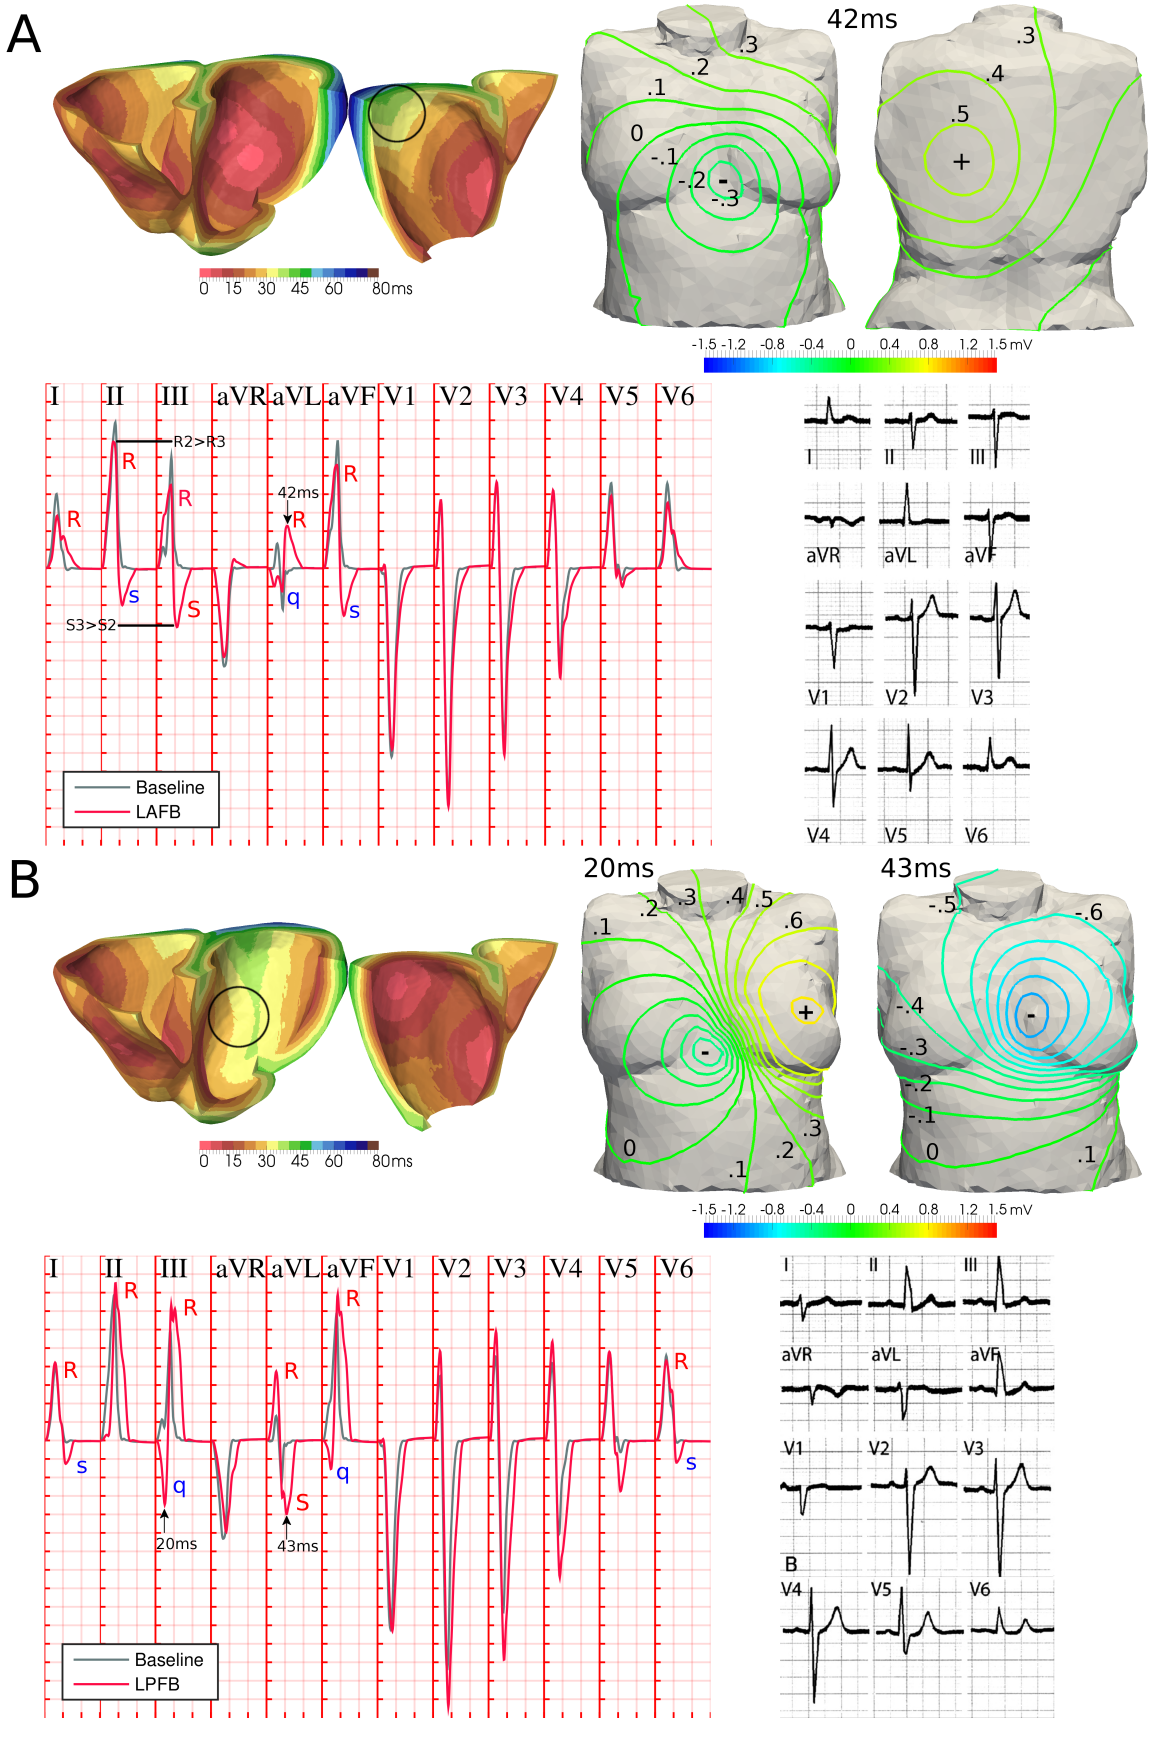
A
42ms
0 15 30 45 60 80ms
-1.5 -1.2 -0.8 -0.4 0 0.4 0.8 1.2 1.5 mV
I II III aVR aVL aVF V1 V2 V3 V4 V5 V6
R2>R3
S3>S2
42ms
Baseline
LAFB
B
20ms
43ms
0 15 30 45 60 80ms
-1.5 -1.2 -0.8 -0.4 0 0.4 0.8 1.2 1.5 mV
20ms
43ms
Baseline
LPFB

approach tightly couples state-of-the-art anatomically based multiscale bidomain simulations of human ventricular electrophysiology with ex-vivo and in-vivo human activation data. This allows augmenting the information attainable from scarce and limited experimental recordings in human (usually available for only small numbers of subjects, of low-resolution, and highly variable between individuals) and gaining novel physiological knowledge on the information enclosed in the human QRS complex.

R-peak amplitudes exhibit the largest amount of variability, as these are shown to be simultaneously modulated by endocardial activation speed, and by myocardial intracellular and extracellular conductivities. Variability in torso and organ conductivities shows an impact on QRS magnitude, in agreement with the literature (Keller, Weber, et al. 2010). On the contrary, QRS width was only regulated by activation speed and intracellular myocardial conductivities, whereas intrinsicoid deflection (QR interval) is only affected by the endocardial activation. In particular, there is a high impact in reducing the activation sequence speed on QRS duration. These insights can be of clinical relevance in order to separate from the ECG the contribution of these three factors, as well as to make progress towards more personalized activation sequences in computational studies.

Variability in the anatomical locations of activation root sites (corresponding to

Figure 5.9 *(preceding page)*: Hemiblock of the LV human conduction system. (A) Knockout of the LV anterior branch resulted in the late activation of the LV basal free wall, moving the positive pole in body surface potentials from the left hip to the centre of the back, and recovering the clinical manifestations of left anterior fascicular block (LAFB). (B) Similarly, LV posterior wall stimulation knockout resulted in delayed LV basal paraseptal activation, affecting the time course and position of the negative pole in the front of the torso, with the associated clinical manifestations of left posterior fascicular block (LPFB). Representative ECGs of LAFB and LPFB are reproduced with permission from Elizari, Acunzo, and Ferreiro (2007). ECG grid resolutions: 40 ms/0.1 mV (simulation); 200 ms/0.5 mV (clinical).

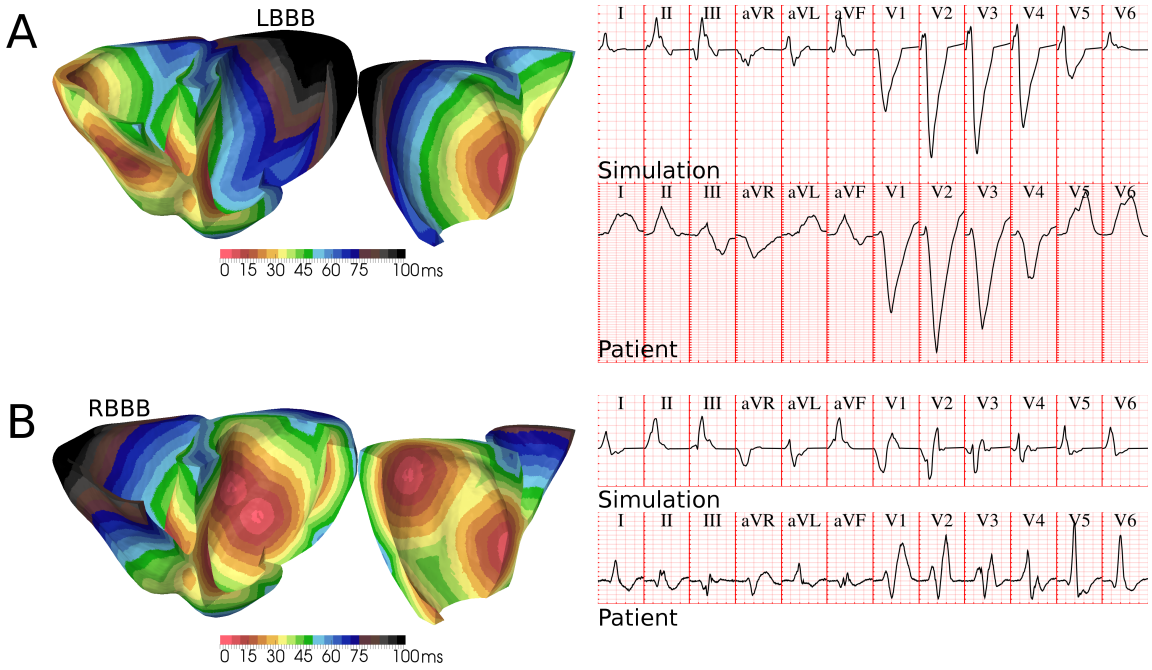

Figure 5.10: Bundle branch block of the human conduction system. (A) A lateral progression of activation in LV due to the simultaneous knockout of the LV anterior/posterior branches closely mimics the diagnostic criteria of left bundle branch block (LBBB) in precordial leads V1 to V6. (B) Similar results for lateral progression of activation in RV, associated with right bundle branch block (RBBB). QRS complexes of LBBB and RBBB patients in the PTB Diagnostic ECG Database (Bousseljot, Kreiseler, and Schnabel 2009) (subjects 208 and 209, respectively) are provided for comparison. Clinical and simulated ECG grids resolution: 40 ms/0.1 mV.

the coupling of the LV fascicular bundles with the myocardium) was associated with S-wave progression in the limb and precordial leads, and occasional notched QRS complexes in the precordial derivations. Whereas insights in S-wave progression could be used to further particularize patient-specific activation sequences, it is important to remark that the notched QRS complexes occurred in the absence of structural defects in the ventricles, in spite of their frequent association. Finally, variability in the number of activation sites recovered the pathological landmarks introduced by disease in the human conduction system in the QRS complex, both under partial and complete bundle block conditions.

Different modelling approaches have been used to date to study the ventricular activation sequence on the human ECG. These vary in generality, from specifying

activation times analytically by means of a parameterized sequence (Boulakia et al. 2010), to the creation of idealized Purkinje networks including explicit Purkinje-muscle junctions (Okada et al. 2013; Sahli Costabal, Hurtado, and Kuhl 2015). They also differ in computational complexity, from the semi-automatic generation of activation profiles (Keller et al. 2011; Potse et al. 2014), to more iterative and interactive parametrization processes to converge on a personalized activation sequence (Okada et al. 2013). The activation sequences utilized here encapsulate the minimal information needed to investigate the implication on the QRS complex of anatomical variability in the human trifascicular conduction system. This allows for an accurate, robust and efficient simulation of the human QRS complex, requiring minimum user input beyond the specification of Purkinje-branch end points.

Compared to previous studies, our work represents a detailed investigation on how variability at the conduction level (progression of the activation sequence, passive tissue conductivities, and anatomical location of Purkinje-myocardial coupling) impacts the variability observed at the human QRS complex. Although in this work we have explicitly concentrated our investigations in exploring variability at the tissue level, the tools presented here could be easily combined with experimentally calibrated populations of models to capture human variability at the cellular electrophysiological level (Muszkiewicz et al. 2016). This may considerably advance the field of computational cardiac electrophysiology towards a more predictive technology for risk assessment at the population level, as both variability in conduction and cellular electrophysiology modulate individual responses to pharmacological and electrical therapy.

An accurate modelling of pathological conditions of the human activation system might also yield a better understanding of clinical ECG biomarkers of disease. Our activation sequences have been shown to replicate the main diagnostic manifestations of left anterior fascicular, left posterior fascicular, left bundle branch and

right bundle branch block conditions. Previous computational studies have also shown agreement in diseased QRS biomarkers associated with left bundle branch block, using either a similar modelling approach to ours for the human activation sequence (Potse et al. 2014), or more detailed anatomical representations of the human Purkinje network (Okada et al. 2013). Sahli Costabal, Hurtado, and Kuhl (2015) also addressed the modelling of right bundle branch block using detailed Purkinje trees, however without probing its impact on the precordial leads, which are the derivations used in its clinical diagnosis. To the best of our knowledge, modelling of fascicular hemiblocks on the human QRS complex has not been previously addressed in the literature. The adoption of the activation sequences proposed here could therefore also promote the computational modelling of other pathological conditions (myocardial infarction, ischaemia, heart failure, cardiomyopathies), frequently affected by comorbidities in the human conduction system.

Variability in the location of activation sites from our baseline configuration had only a small influence on the QRS axis. This may explain the preponderance of a QRS axis around 60° in healthy individuals, as reflected by the PTB Diagnostic ECG Database (Bousseljot, Kreiseler, and Schnabel 2009) where 65 % (33/51) of the healthy cohort was classified as having QRS axes between 30° and 70°. It is known that a vertical heart orientation (tall, thin individuals) is manifested in a right QRS axis shift, whilst a left QRS axis is exhibited under more horizontal heart orientations (short, broad individuals) (Elizari, Acunzo, and Ferreiro 2007). The simulation results therefore reveal that the location of stimulation sites and the other varied parameters cannot explain the large variability in the healthy QRS axis observed clinically. Therefore, other factors such as the relative orientation of the human heart within the thoracic cavity, ventricular mass, and torso geometry might be the main determinants of variability in the healthy human QRS complex (Hoekema, Uijen, and Oosterom 2001). Testing of these hypotheses would require

the expensive generation of new heart-torso anatomical models with variable heart and torso volumes and ventricular orientation, and will therefore be addressed in future investigations. Corrections to account for this geometrical variability have been proposed (Hoekema, Uijen, and Oosterom 2001), which will also be investigated in future work.

The activation sequences presented here have been shown to be a feasible approach to capture the main intraventricular conduction defects of the human conduction system. However, their underlying simplifications may limit their applicability to model other pathophysiological contributions of the Purkinje system (Vigmond and Stuyvers 2016). This may include ventricular extrasystoles arising from Purkinje fibres in idiopathic intrafascicular tachycardia, or retrograde propagation through the Purkinje network. Whereas the former could *a priori* be easily addressed by timing additional ectopic foci in conjunction with the default activation sequence, the latter is likely to require a full coupling of the ventricles with a detailed anatomical representation of the Purkinje system. In addition, our results have not been compared against those obtained using more detailed anatomical representations of the Purkinje system, due to the lack of high-resolution data in human in which to build such models. This may constitute an interesting line of future research, should the experimental data in human become available, as well as an in-depth comparison against different tissue descriptions of myocardial conduction.

Chapter 6

Repolarization gradients in human ventricular electrophysiology and the T wave

In chapter 4, performance improvements were made in Chaste that enabled much more resource-efficient simulations to be performed of ventricle-torso bidomain models of the ECG. Such simulations were utilized in the preceding chapter to demonstrate a new model of the activation system, generating realistic ECG waveforms of normal and pathological cases. This chapter builds directly on the normal model of activation, by considering the subsequent repolarization of the ventricles in the presence of regional heterogeneity in electrophysiology and evaluating the success of each simulation in recovering normal ECG waveforms as defined by a control population.

6.1 Introduction

The T wave is an ECG feature resulting from repolarization of the ventricles. The local repolarization time throughout the ventricles is determined by the activation time (AT) and the action potential duration (APD), therefore realistically simulating the T wave depends on modelling these two factors. Simulations of activation times are presented in the previous chapter. This chapter concerns the second aspect: the effect of APD heterogeneities on the T wave given a normal activation profile.

Following experimental and clinical evidence on heterogeneous distributions of ion channel density and APD in the human ventricles (Boukens et al. 2015; Bueno-Orovio et al. 2012; Glukhov et al. 2010; Janse et al. 2012; Opthof, Coronel, and Janse 2009; Opthof et al. 2007, 2016; Patel et al. 2009; Péréon et al. 2000; Szabó et al. 2005; Szentadrassy et al. 2005; Taggart et al. 2003; Zicha et al. 2004) as well as in other mammalian species (Janse et al. 2012; Opthof, Coronel, and Janse 2009; Opthof et al. 2016; Patel et al. 2009; Szabó et al. 2005; Szentadrassy et al. 2005; Zicha et al. 2004; Zygmunt, Goodrow, and Antzelevitch 2000; Zygmunt et al. 2001), three idealized 'axes' of dispersion will be considered: transmural, apicobasal, and interventricular. The approach is therefore similar to previous research in the field (Keller et al. 2012; Nguyên et al. 2015; Okada et al. 2011; Zheng et al. 2016), although with an emphasis in basing our *in silico* investigations on updated experimental data from human. All combinations of aforementioned axes will be tested in simulation with the same activation protocol, and the T wave characteristics compared to each other and to a database of ECGs from 52 healthy controls (Bousseljot, Kreiseler, and Schnabel 2009).

6.2 Methods

6.2.1 Human ventricular APD dispersion

The human dispersion data used in this work are summarized in table 6.1. The references therein report that APD is longer (1) in the (sub)endocardium than (sub)epicardium, (2) in the base than the apex, and (3) in the LV than the RV. Because these are inconsistently reported as either APD90, APD80, or ARI (a surrogate of APD), the published data have been used to calculate the dispersion as a percentage of the mean APD (or ARI) for straightforward translation to the mathematical model.

	Mean APD (ms)	Dispersion (ms)	Dispersion (%)
Transmural			
Boukens et al. (2015)	339 ± 13	42 ± 6	12.4
Glukhov et al. (2010)	350 ± 13	67 ± 13	19.2
Näbauer et al. (1996)	455 ± 19	62 ± 19	13.7
Apico-basal			
Ramanathan et al. (2006)	235	42	17.9
Interventricular			
Bueno-Orovio et al. (2012)	202 ± 15	11 ± 15	5.5
Ramanathan et al. (2006)	256	32	12.5

Table 6.1: Human ventricular APD dispersion measurements used to parameterize the heterogeneous models.

6.2.2 Modelling APD heterogeneity

As in the previous chapter, the endocardial variant of the human ventricular model by ten Tusscher and Panfilov (2006) was used to model cardiomyocyte transmembrane voltage dynamics. Each of the heterogeneities—transmural, apico-basal, and interventricular—were applied independently for a total of eight simulated cases. The simplest case was of homogeneous (HOM) ventricles, with no regional varia-

tion in ionic current densities.

The heterogeneous-APD cases were achieved by varying the slow delayed rectifier potassium current (I_{Ks}) conductance g_{Ks}, which is known to strongly influence APD (Viswanathan, Shaw, and Rudy 1999). This mechanism was used for the transmural heterogeneity rather than utilizing the published epicardial and M-cell model variants for several reasons: (1) using multiple cell types would require making assumptions about the distribution of the cell types, e.g. in layers or 'islands' (as in Keller et al. (2012)); (2) adjustment of parameters would still be needed to generate the desired dispersion of APD because the default APDs for each model variant are not necessarily suitable for direct inclusion when considering other human data; (3) the majority of differences between the variants relate to irrelevant behaviour for this application, e.g. gating variables and channel kinetics; (4) the functional role of M-cells in the intact human ventricular wall remains controversial (Opthof, Coronel, and Janse 2009; Taggart et al. 2001).

In the absence of high resolution APD data, the transmural and apico-basal cases were modelled with smooth gradients in g_{Ks}, with slopes informed by *in vivo* measurements. In order to produce a symmetric APD distribution about the default APD, the conductivity parameter was set with the following form:

$$g_i(x) = c_i^x \cdot g_{\text{default}}, \quad -1 \leq x \leq 1, \quad 0 < c_i \leq 1, \tag{6.1}$$

where subscript i denotes the heterogeneity (transmural, etc.), c is a (constant) parameter that determines the magnitude of a given APD gradient, x represents the linear distance along the gradient, and g_{default} is the default value ($0.392\,\text{nS}\,\text{pF}^{-1}$). The result is a smooth change in g_{Ks} from c to $1/c$ times the default value. In the transmural case, x equals 1 on the endocardium and -1 on the epicardium; in the apico-basal case it is -1 at the apex and 1 at the base. The interventricular case was

similar, but with x taking a discrete value of -1 in the right ventricle or 1 in the left ventricle. The septum was treated as a continuation of the LV free wall, following Konarzewska, Peeters, and Sanguinetti (1995).

Heterogeneities were combined by averaging the individual c^x 'scaling factors', i.e.

$$g_{\text{Ks}} = \frac{g_{\text{default}}}{N} \sum_i c_i^{x_i} \tag{6.2}$$

where N is the number of heterogeneities applied (1, 2, or 3).

The required scaling factor for a given APD dispersion was estimated using a 1-D fibre model. This is because the APD dispersion in a fibre or a wedge is not simply the difference in APDs that would be seen in isolated cell models with given parameters; it is also affected by (1) cell-coupling and (2) edge effects. The former refers to the tendency of nearby regions of tissue to remain smoothly varying due to the non-zero tissue conductivity. The latter refers to small deviations in the expected tissue properties at the edges of the domain due to the reduction in current sinks. When considering these confounders, the scaling factor for a given (desired) APD dispersion is therefore also a function of fibre length. To illustrate this relationship, APD dispersion was calculated empirically for various fibre lengths and scaling factors, as shown in fig. 6.1.

As shown in fig. 6.1, for long fibres the coupling effect is negligible and the maximum dispersion is produced for a given scaling factor. Also note that the dispersion with unity scaling factor is always greater than zero, due to the edge effect. A helpful way to conceptualize the relationship shown is to observe that, as the fibre shortens, for a given APD dispersion we need a larger gradient in g_{Ks} (i.e. a scaling factor closer to zero) to overcome the increasing normalizing strength of tissue coupling.

For the apico-basal and interventricular heterogeneities the 'long fibre' value may be used, because for the former the apex-base distance is long (about 7 cm),

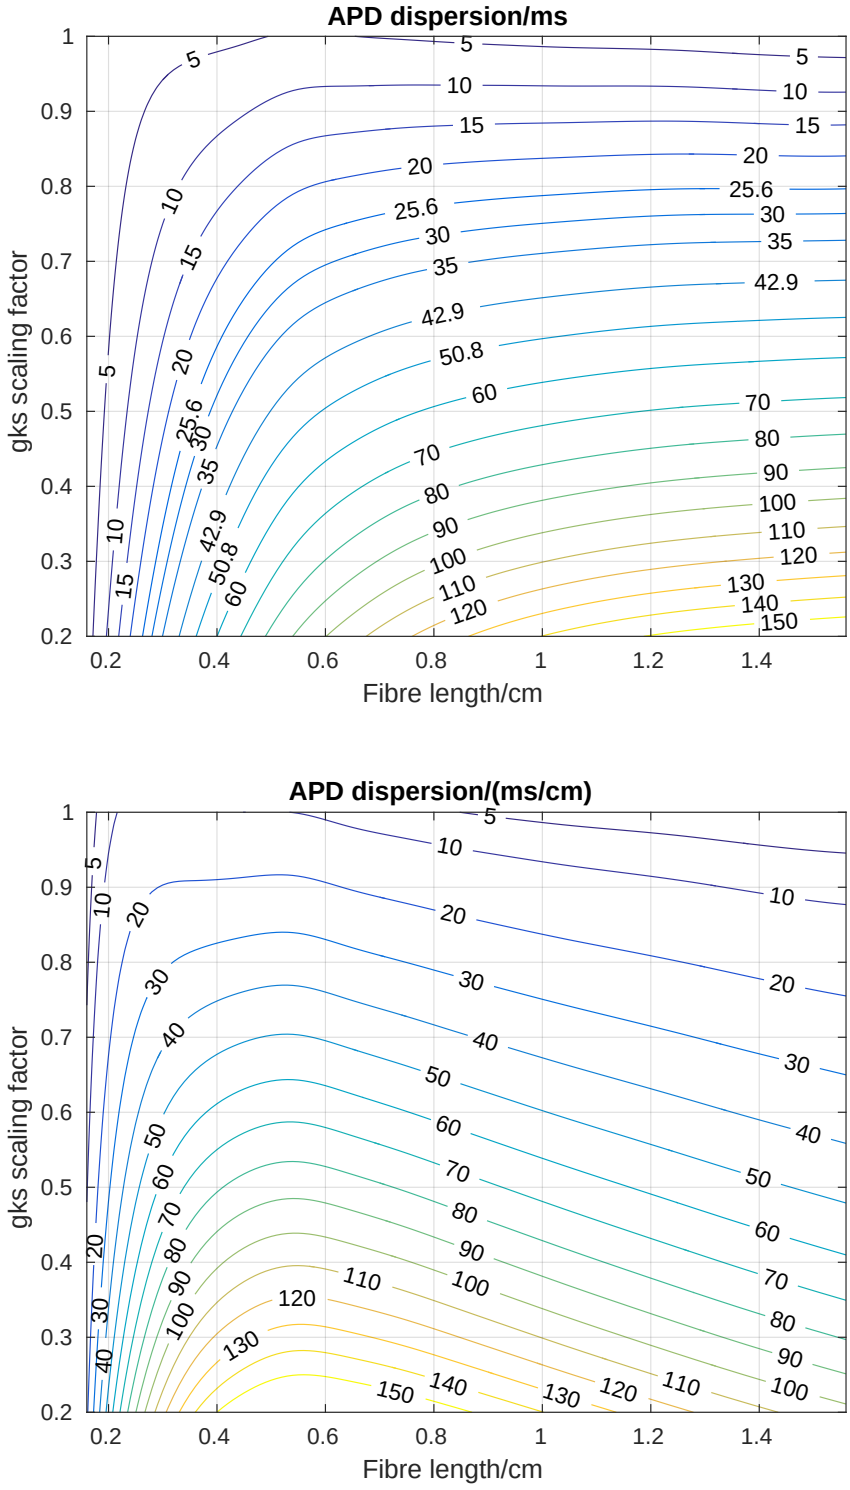

Figure 6.1: The dispersion of APD measured in a 1-D fibre model with a range of g_{Ks} scaling factors (the c in eq. (6.1)). Top: dispersion in ms. Bottom: the same data as the panel above, normalized by fibre length (i.e. dispersion in ms cm^{-1}).

and for the latter the factors are applied discretely to two large contiguous volumes. The relatively thin and variable thickness of the transmural wall, however, presents a subtle modelling choice, three alternatives for which are discussed below.

The first and simplest option is to use the same scaling factor everywhere for the transmural heterogeneity. Due to the varying wall thickness, the result of this would be less APD dispersion across thinner parts of the wall (more common in the RV) than in thicker parts (more common in the LV). For example, a scaling factor of 0.5 in fibres of the median LV and RV wall widths (fig. 6.2) would result in APD dispersions of about 66.8 ms and 31.5 ms, respectively. This is a substantial difference considering that the cell models at the edges would have the same parameters in either case. The physiological interpretation of this method might be that the expression of ion channels at the endocardial and epicardial surfaces are the same throughout the ventricles.

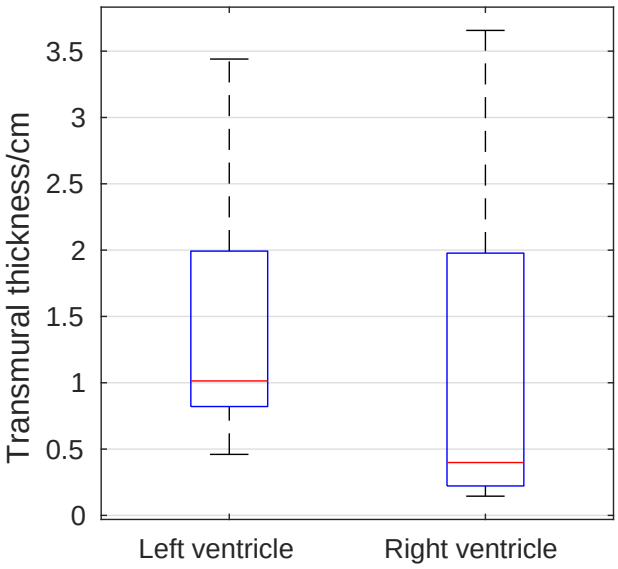

Figure 6.2: Boxplots of transmural thickness in the mesh (section 3.3.2). The median thickness is 1.014 cm in the LV and 0.399 cm in the RV.

Alternatively, to achieve a constant transmural APD dispersion regardless of

the wall thickness would require modifying the scaling factor according to the thickness, specifically using a larger gradient where the walls are thinner and a smaller gradient where it is thicker. As shown in fig. 6.1 this would require imposition of extreme gradients in the thinner regions, and some very thin regions might be unable to support such a dispersion at all. The physiological interpretation of this is that ion channel expression conspires to generate a fixed APD dispersion regardless of wall thickness, which is unlikely.

A third approach is to impose a g_{Ks} gradient that models a constant dispersion *per unit distance*, as in Zheng et al. (2016). The lower panel in fig. 6.1 shows the same data in the panel above normalized by fibre length. As in the previous scheme, the scaling factor would be adjusted depending on the wall thickness, taking a maximum value at about 0.55 cm (smallest gradient) and reducing (larger gradient) at other widths to maintain a constant dispersion per distance. Since wavefronts normally travel transmurally, the physiological interpretation might be that APD changes as a function of activation time, alternative mathematical frameworks for which exist already (Bueno-Orovio et al. 2014).

Based on the former approaches and their respective interpretations, the first method was chosen. Table 6.2 shows the mean dispersion (as a percentage of the mean APD) for each type of dispersion for the data in table 6.1. From these the desired model dispersion was calculated based on the default model APD of 284.5 ms. Finally the data presented in fig. 6.1 were used to determine the required scaling factor to use for each heterogeneity (assuming a typical tissue length of 1.4 cm). Note the contours in fig. 6.1 (upper panel) at 25.6 ms, 42.9 ms, and 50.8 ms, corresponding to the desired dispersions. Cross sections of the ventricles showing the scaling factors for all eight systems are shown in fig. 6.3.

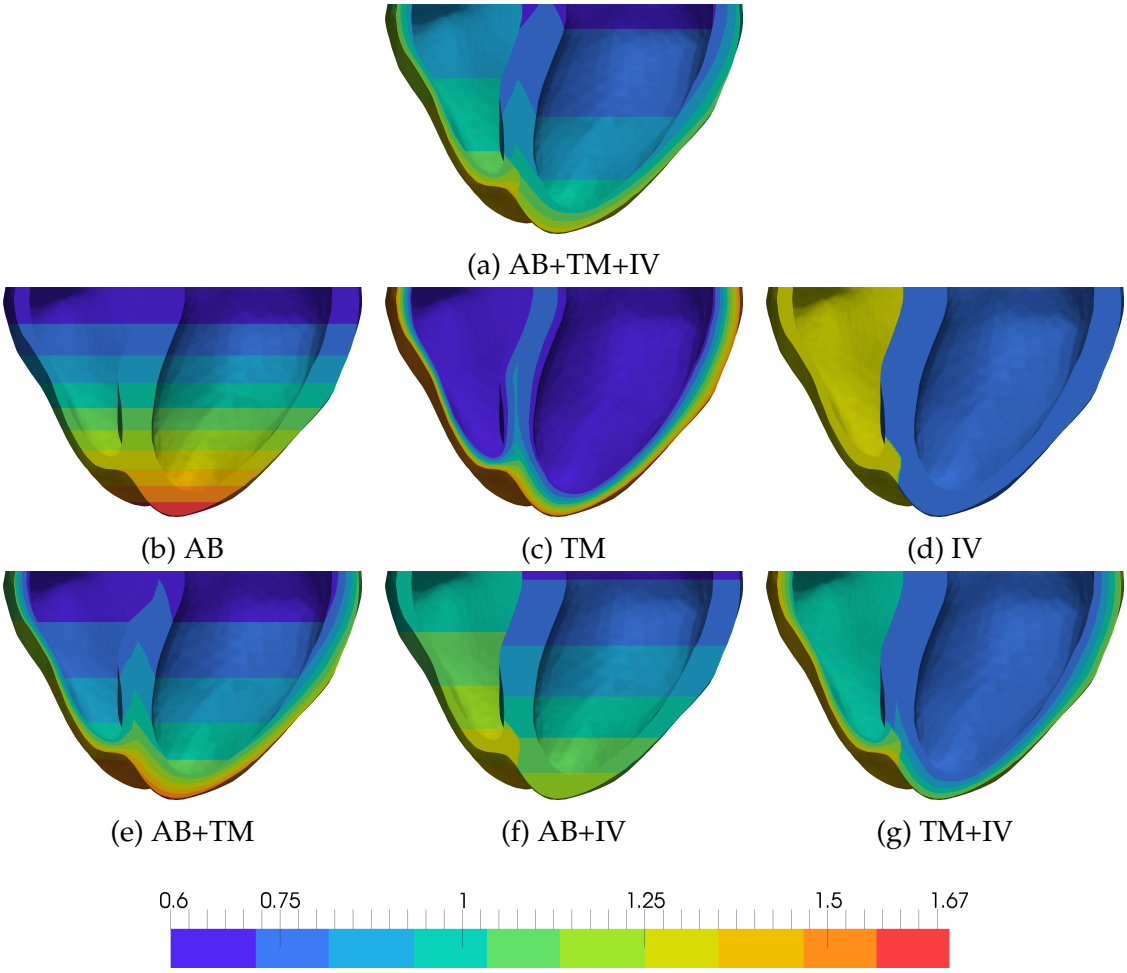

Figure 6.3: Cross sections of the ventricles showing g_{Ks} 'scaling factors' for the seven non-homogeneous simulations. AB = apico-basal, TM = transmural, IV = interventricular. The individual scaling factors (b–d) are combined by taking the mean to produce the combined heterogeneities (a, e–g). The larger ventricle is the left ventricle.

	Mean dispersion/%	Model dispersion/ms	Scaling factor
Transmural	15.1	42.9	0.67
Apico-basal	17.9	50.8	0.62
Interventricular	9.0	25.6	0.80

Table 6.2: Mean experimentally determined dispersion (from table 6.1), the desired model dispersion based on the default APD90 (284.5 ms), and the scaling factor used, for each heterogeneity.

6.2.3 Control ECG database

ECGs for 52 healthy controls were extracted from the Physikalisch-Technische Bundesanstalt (PTB) ECG database (Bousseljot, Kreiseler, and Schnabel 2009) provided by PhysioNet (Goldberger et al. 2000) for comparison to the simulated ECGs. Patients #276 and #284 were excluded due to the presence of multiple QRS and T wave abnormalities, and mislabelling of the leads, respectively.

The T waves of the remaining patients were classified automatically as positively deflected, negatively deflected, low amplitude, and other, using a custom algorithm as follows. A 350 ms sample was extracted from each lead beginning 140 ms after the R peak in lead II. Let $\vec{X}$ denote the signal as a vector of samples

$$\vec{X} = (X_0, X_1, \ldots, X_{350}). \tag{6.3}$$

Baseline wander was managed by offsetting the signal by the mean of the first and last samples, i.e. replacing $\vec{X}$ with

$$\vec{X} - \frac{X_0 + X_{350}}{2}. \tag{6.4}$$

The T wave was then classified in the following order.

$$\text{1: low amplitude if} \qquad \max \vec{X} - \min \vec{X} < S \tag{6.5}$$

$$\text{2: positively deflected if} \qquad \left|\max \vec{X}\right| > 3 \cdot \left|\min \vec{X}\right| \tag{6.6}$$

$$\text{3: negatively deflected if} \qquad \left|\min \vec{X}\right| > 3 \cdot \left|\max \vec{X}\right| \tag{6.7}$$

$$\text{4: other} \tag{6.8}$$

where S was 0.1 mV for the limb leads or 0.2 mV for the precordial leads. The results of the classification are shown in table 6.3. The majority of T waves were assigned to the first three categories; only 5.5 % (33/600) were categorized as other. Most of these were difficult to classify due to baseline drift (fig. 6.4), the correction of which is an active area of research and outside the scope of this thesis. Nearly half of these (15/33) were from two patients (#131 and #184), probably indicating low-quality recordings. The algorithm is expected to also place biphasic waves into this category, though from visual inspection there is only one candidate for a biphasic wave (dark blue, eighth from bottom, V1 in patient #165).

	I	II	III	aVR	aVL	aVF	V1	V2	V3	V4	V5	V6
Small	1	-	23	-	14	7	29	8	4	3	1	9
Positive	47	47	15	-	30	37	9	42	45	46	48	41
Negative	-	-	3	48	1	-	9	-	-	-	-	-
Other	2	3	9	2	5	6	3	-	1	1	1	-

Table 6.3: T wave shapes of 50 healthy controls as classified with a custom algorithm.

6.2.4 Calculation of simulated whole-ventricular APD dispersion

Apico-basal dispersion was calculated using simple linear regression of the LV epicardial APDs. Transmural dispersion was calculated as the mean LV *epi*cardial APD minus the mean LV *endo*cardial APD, i.e. a positive dispersion implies a

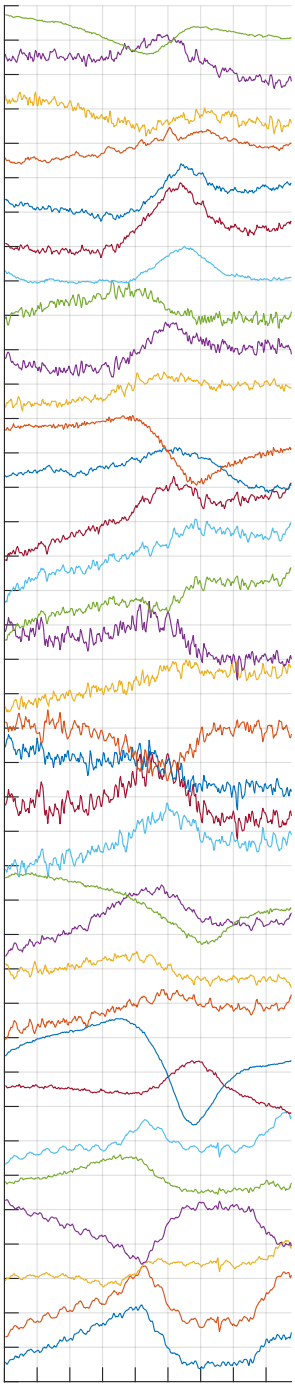

Figure 6.4: Thirty-three 'T waves' that were difficult to classify and therefore labelled 'other'. The signals are noisy and the majority suffer from baseline drift. Grid spacing: 40 ms, 0.1 mV.

longer APD at the epicardium than at the endocardium, and *vice versa*. Finally, the interventricular dispersion was calculated as the mean LV epicardial APD minus the mean RV epicardial APD, i.e. a positive dispersion implies a longer APD in the LV than the RV, and *vice versa*.

6.3 Results

6.3.1 Summary of APD dispersions

Cross-sectional visualizations of APD90 for the eight simulated cases are shown in fig. 6.5, and repolarization time (approximated as upstroke time plus APD90) in fig. 6.6. A summary of simulation APD statistics is provided in table 6.4.

	Apico-basal		Transmural		Interventricular	
	ms	%	ms	%	ms	%
HOM	0.3	0.1	6.0	2.1	−0.3	−0.1
AB	62.4	21.3	10.4	3.6	−5.8	−2.0
TM	−2.9	−1.0	41.5	14.4	−10.2	−3.5
IV	0.4	0.1	6.8	2.3	28.0	9.5
AB+TM	25.5	8.8	26.6	9.2	−8.3	−2.9
AB+IV	34.9	11.9	8.9	3.0	11.6	4.0
TM+IV	−1.6	−0.5	26.2	9.0	7.2	2.5
All	19.2	6.6	20.8	7.2	3.3	1.1

Table 6.4: Summary of APD dispersion statistics for the eight simulations, in milliseconds and as a percentage of the mean APD in the simulation.

The method of combining heterogeneities by averaging scaling factors is evident in the reduction in dispersion of APD with multiple heterogeneities. For example, the apico-basal dispersion in the AB+TM case (25.5 ms) is close to the average of the individual AB and TM cases (62.4 ms and −2.9 ms). The same can be seen in the other statistics for the other heterogeneity combinations.

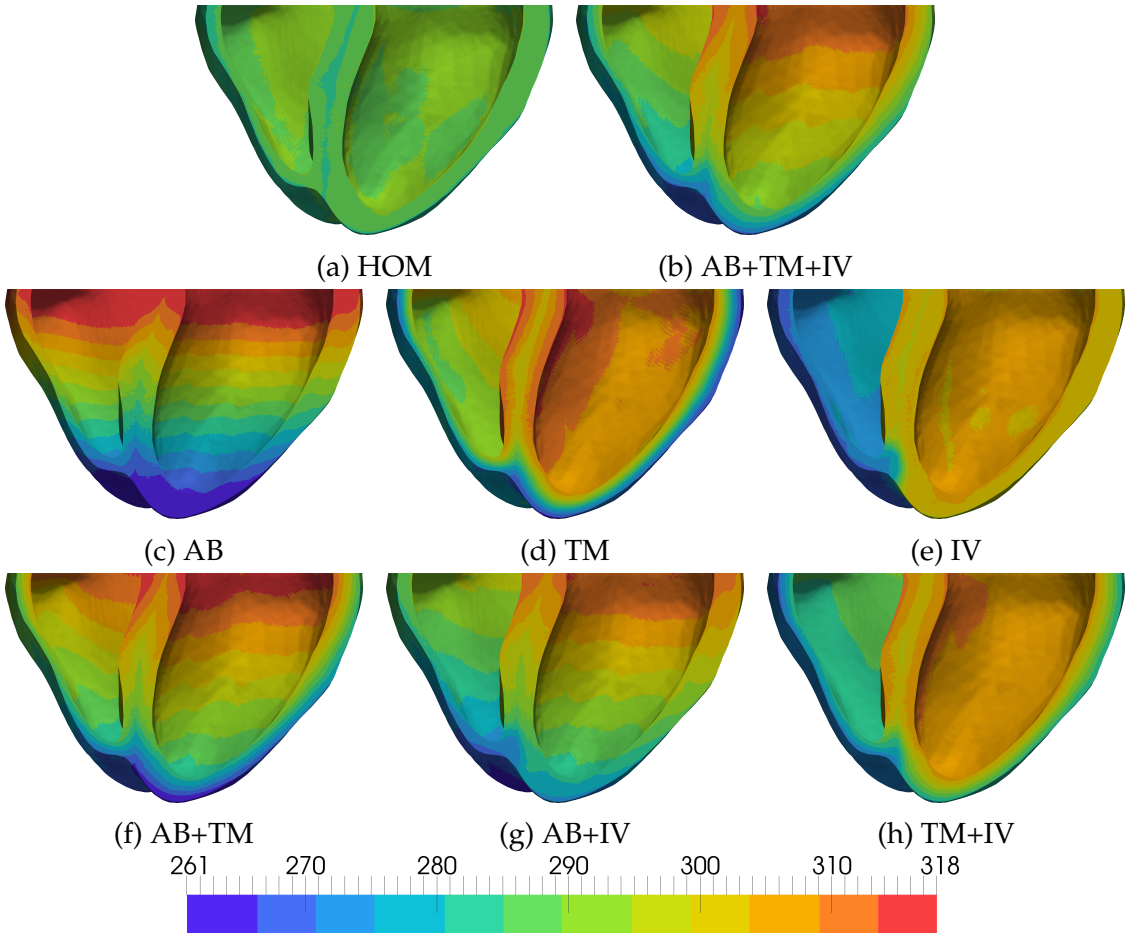

Figure 6.5: Cross sections of the ventricles showing APD90 (in ms) for the eight simulations. HOM = homogenous, AB = apico-basal, TM = transmural, IV = inter-ventricular. The larger ventricle is the left ventricle.

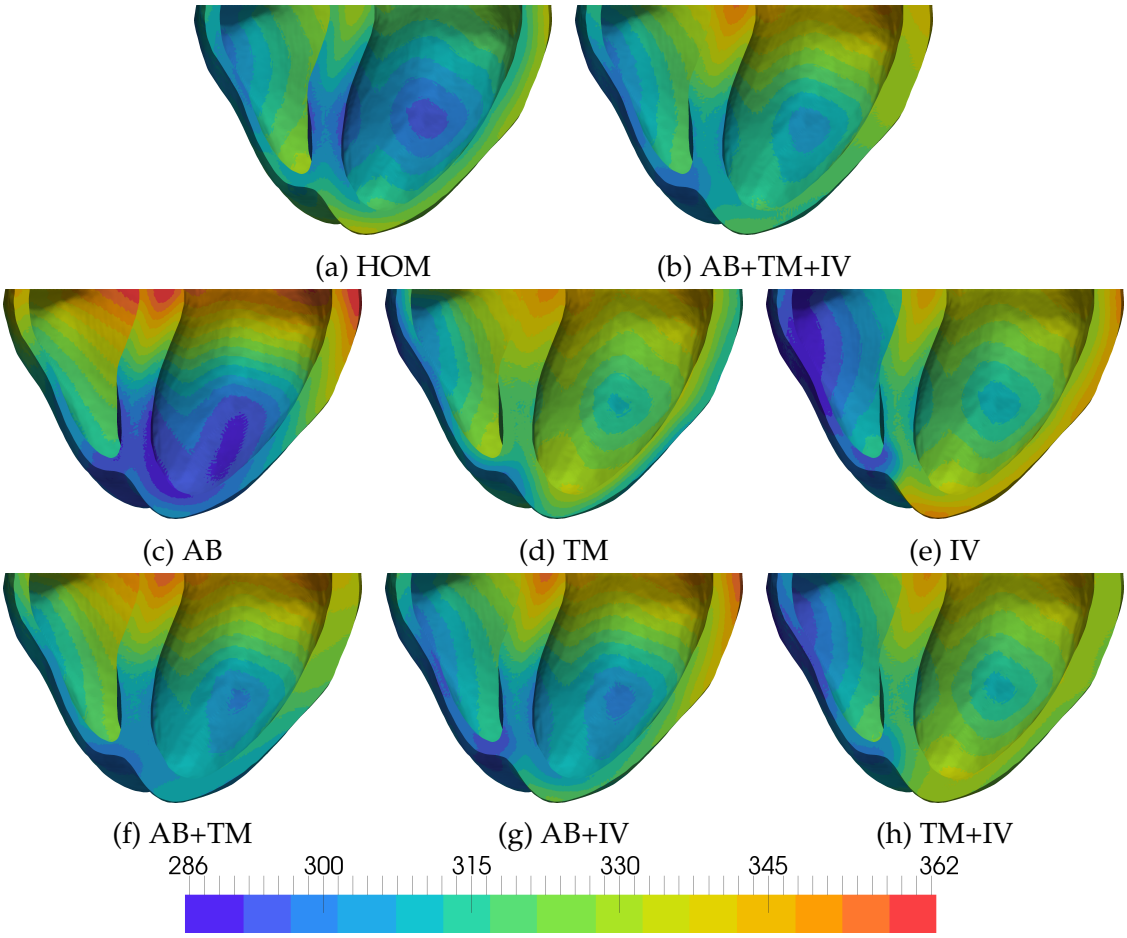

Figure 6.6: Cross sections of the ventricles showing repolarization time (in ms) for the eight simulations. HOM = homogenous, AB = apico-basal, TM = transmural, IV = interventricular. The larger ventricle is the left ventricle.

6.3.2 HOM and two of the IV cases had inverted T waves in several leads

Contrary to the expected repolarization sequence in a healthy human ECG, the HOM, IV, and AB+IV simulations generated inverted (negative) T waves in leads I and V6 (fig. 6.7). The HOM case exhibited a negative T wave in the other lateral leads (II and aVF). In V5, the HOM and IV cases had slightly biphasic T waves, and the AB+IV case was slightly notched. Finally, in aVR, the HOM and IV cases had inverted (positive) T waves. None of these features were observed in any of the healthy controls (table 6.3). Note that in the controls, no biphasic waves were seen in V5; the single lead classified as other suffered from baseline drift.

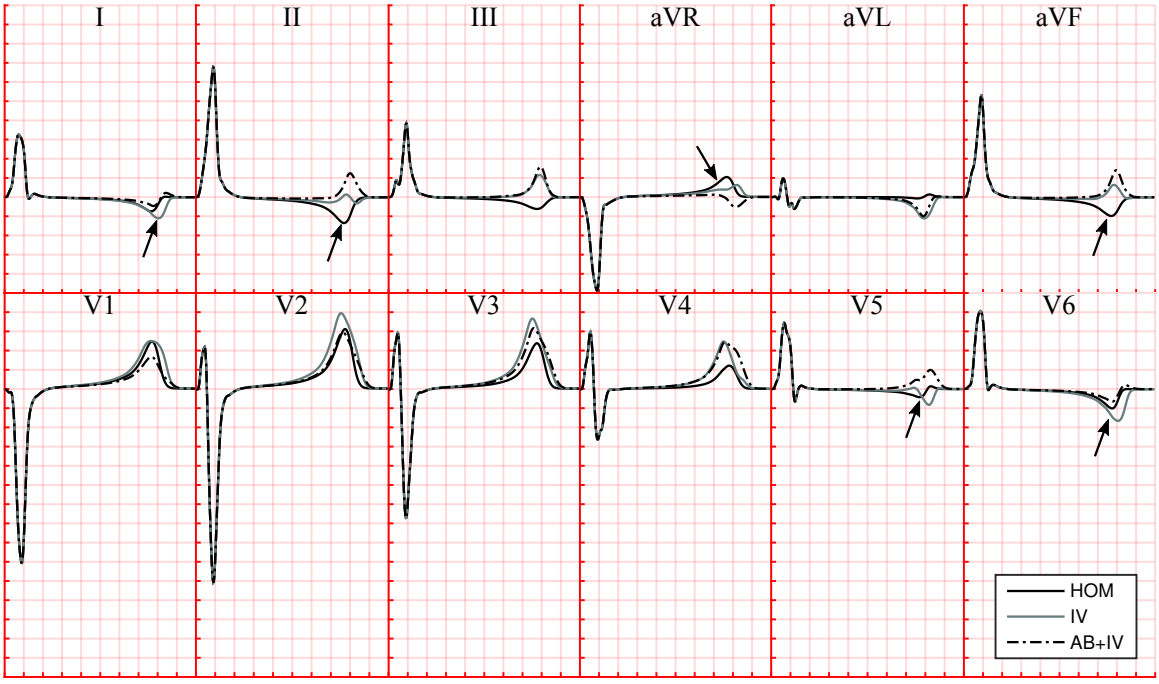

Figure 6.7: For homogeneous ventricles and IV gradients in the absence of transmural APD shortening, inverted T waves were observed in multiple limb leads and V6 as well as biphasic/notched waveforms in V5. Grid spacing: 40 ms, 0.1 mV.

6.3.3 Other simulations with an IV heterogeneity produced notched T waves in V5 and limb leads

The remaining two systems with the imposed interventricular heterogeneity (TM+IV and the system with all three heterogeneities) generated notched (bifid) T waves in V5 (fig. 6.8). Small notches are also visible in leads I, III, and aVF, although the T waves in these leads are relatively small in magnitude. No notches were seen in these leads in the controls.

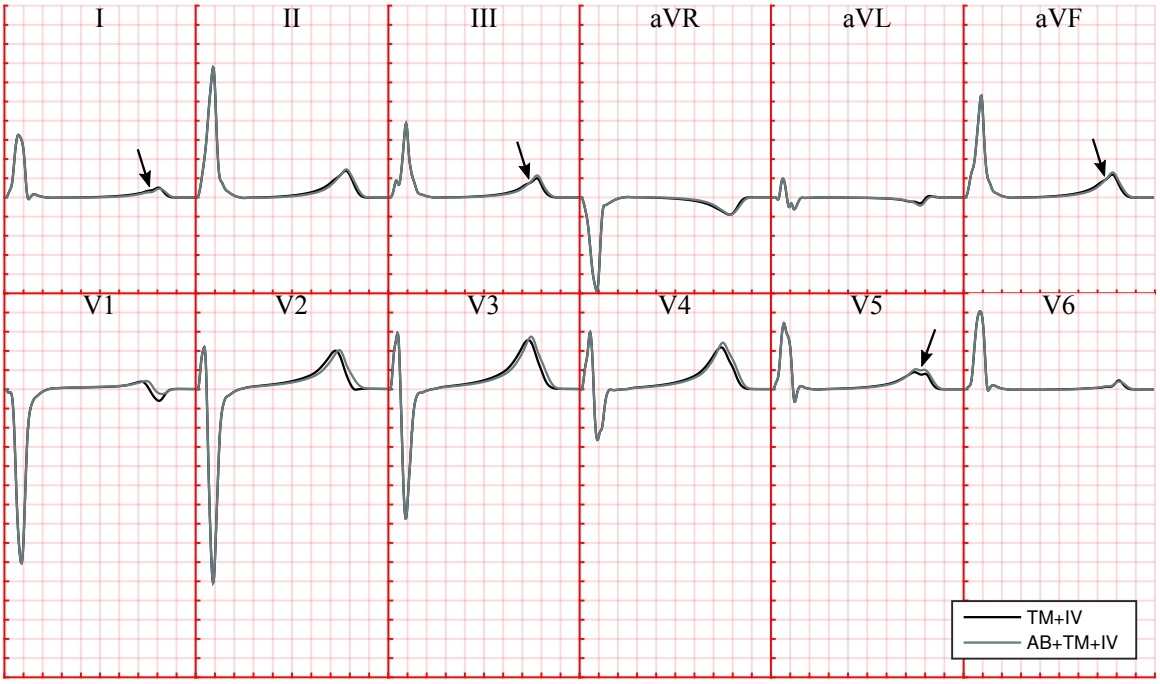

Figure 6.8: IV and TM gradients with and without an apico-basal gradient produced abnormal T wave features, specifically notched T waves in the inferior leads and V5. Grid spacing: 40 ms, 0.1 mV.

6.3.4 The AB and TM cases have irregular septal leads

The AB and TM simulations generated abnormalities not seen in the healthy controls in the septal leads, V1 and/or V2 (fig. 6.9). The apico-basal heterogeneity on its own (AB) produced 'camel hump' T waves in V1, and a broad, flat-topped

T wave in V2. The transmural heterogeneity on its own (TM) was normal in V1 but generated a fairly large (about 0.3 mV peak-to-peak) biphasic T wave in V2.

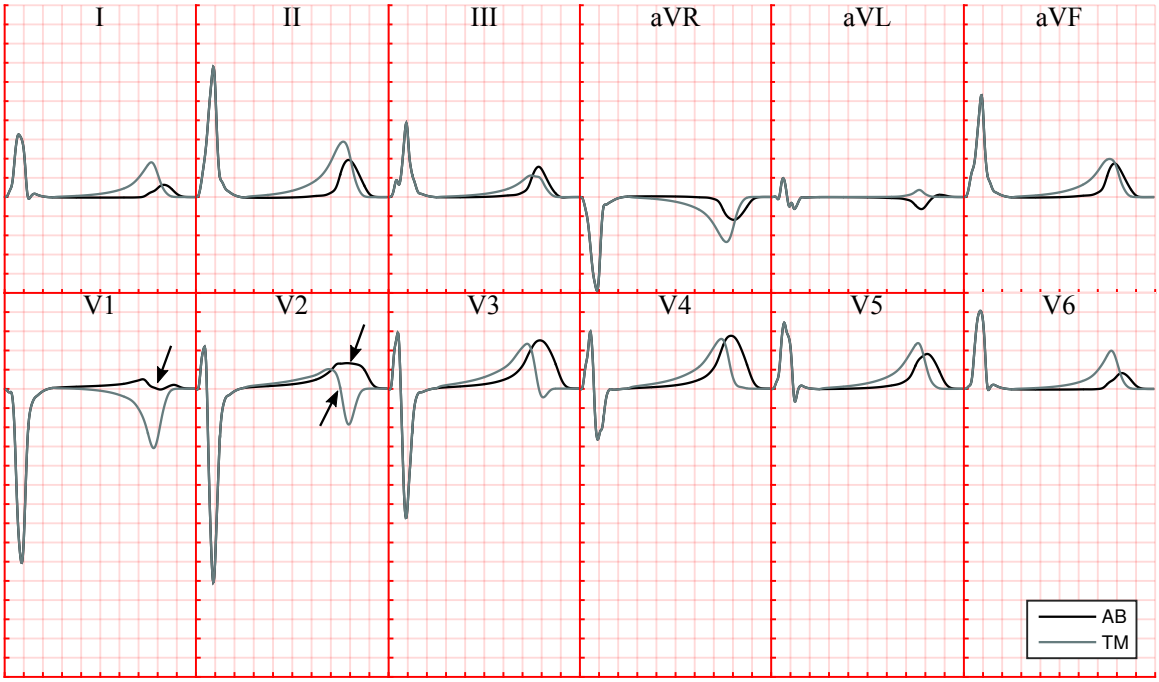

Figure 6.9: A 'camel hump' and flattened T wave are seen for the AB simulation in leads V1 and V2, respectively. The TM simulation generated a biphasic T wave in V2. Grid spacing: 40 ms, 0.1 mV.

6.3.5 The AB+TM system is normal

The one remaining APD distribution, featuring apico-basal and transmural heterogeneities (AB+TM), produced T waves morphologically typical according to the control set (fig. 6.10). The T waves are upright in all leads apart from aVR and V1, as normal. The ECG for control patient #242 has similar T waves and is reproduced for comparison.

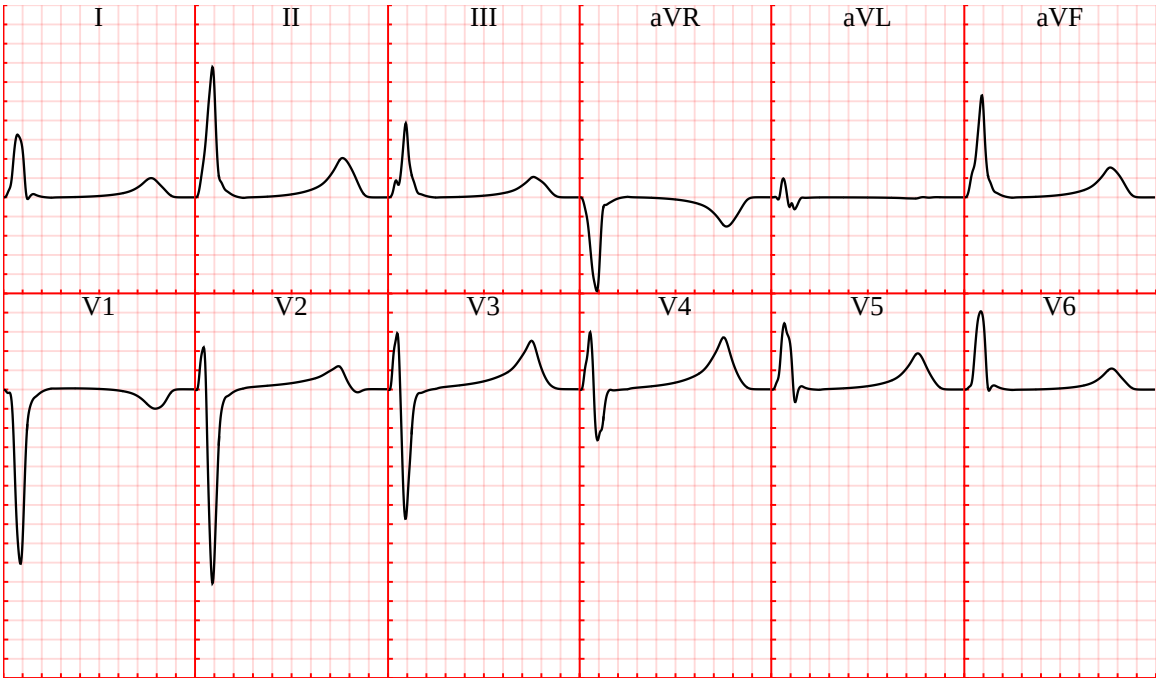

(a) Simulation. Grid spacing: 40 ms, 0.1 mV.

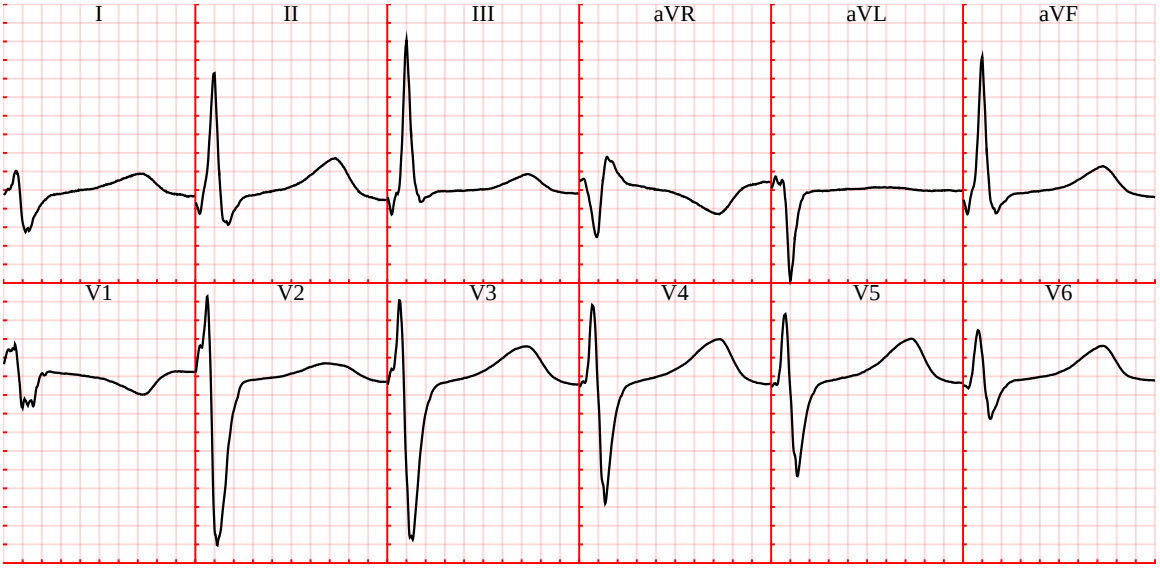

(b) Recording. Grid spacing: 40 ms, 0.2 mV.

Figure 6.10: (a) In the presence of apico-basal and transmural heterogeneities in APD, the resulting T waves appear normal. (b) An ECG recorded from patient #242 in Bousseljot, Kreiseler, and Schnabel (2009) for comparison to the simulation. Note the similar T wave polarities and relative amplitudes.

6.4 Discussion

Distributions of action potential duration in the ventricles were constructed according to human data and applied to a multiscale human ventricular ECG model. The resulting T waves were examined for morphological similarity with those of 50 control patients. The method presented of regionally adjusting the APD (eqs. (6.1) and (6.2)) produced APD distributions that were symmetrical about the default APD and combined intuitively (figs. 6.3 and 6.5 and table 6.4).

All but one of the systems had at least one T wave with a morphology not seen in the control set and/or typically pathological. Unsurprisingly when considering the underlying physics, the homogeneous system (fig. 6.7) produced T waves of opposite polarity to the dominant wave in the QRS complex in every lead. The IV and AB+IV cases (fig. 6.7) also had abnormal T waves in leads I and V6, but had normal polarity in most of the other limb leads. The positive T wave in V1 seen in all three of these simulations (fig. 6.7) is unusual but not unprecedented, occurring in a minority of control patients (9/50, table 6.3).

The other two simulations with an interventricular heterogeneity produced bifid T waves in V5 (fig. 6.8), as well as subtle notches in some limb leads. A similar notch is seen in V5 in the AB+IV simulation (fig. 6.7). The interventricular heterogeneity as incorporated here appears to be implicated in notched T waves. This is supported by a recent canine study in which bifid T waves were observed when interventricular dispersion was increased using dofetilide (Meijborg et al. 2015).

More unusual abnormalities were seen in the AB and TM simulations (fig. 6.9). Here, T waves were normal apart from in the septal leads (V1 and V2). With only the apico-basal gradient, small camel hump T waves are seen in V1. While a small, positive T wave is not unusual in V1 (as discussed above), two peaks such as these are not observed in the control set. Camel hump T waves are characteristic of hypokalemia or 'hidden U waves' (Mattu 2003). In V2 this simulation generated a

broad, slightly flattened T wave which is nonspecific. Also in V2, the TM simulation exhibited a biphasic T wave reminiscent of Wellen's syndrome (de Zwaan, Bär, and Wellens 1982).

The configuration that included apico-basal and transmural gradients in APD was the only one able to produce normal T waves in every lead (fig. 6.10). The only remarkable lead is perhaps aVL which was electrically silent, as seen in several controls to varying degrees, especially patient #242 as reproduced for comparison. Despite differing methodologies, the finding that the most realistic T waves are generated in the presence of an apico-basal and transmural gradient in APD has been shown in other computational studies, including those by Seemann et al. (2013) and Okada et al. (2011). The simulations in this chapter show the same result with a realistic activation sequence.

In contrast, Arteyeva et al. (2015) report that the further addition of an anterior-posterior gradient, or an anterior-posterior and a left-right gradient, produced the best body surface potentials. Similarly, Zheng et al. (2016) concluded that all four gradients are required. Interestingly, in both of these studies the interventricular gradient has the opposite direction to the ones used in this chapter and by Seemann et al. (2013). The former draws on human data presented above. The publication by the latter contains a literature review of four animal studies in which APD is reportedly shorter in the RV than the LV, and only one to the contrary (a finding inferred indirectly using a computer simulation).

Future applications of this model might further refine the relationship between APD heterogeneity and the T wave in various ways. Performing simulations with a larger number of heterogeneity gradient magnitudes might result in models capable of generating T waves resembling a larger sample of the control patients, particularly in terms of more subtle waveform properties like T wave asymmetry. Similarly, the importance of other heterogeneities such as anterior-posterior could

be evaluated. The model could be used to examine the impact of the activation sequence, and especially conduction abnormalities (as in the previous chapter), on the T wave. Finally, an interesting area of enquiry could be around the supposed relationship between APD and AT, evidence for which is inconsistent (Hanson et al. 2009; Orini et al. 2013), e.g. by imposing proportionality in the model.

6.5 Limitations

The T wave is the result of many factors, some of which may be relatively important and were not investigated here. In particularly, the activation sequence probably has a similar impact as the APD distribution (compare fig. 6.5 and fig. 6.6), and although the activation sequence was carefully constructed to generate a realistic QRS complex (see chapter 5), it is nevertheless only one possible sequence and the sensitivity of the T wave to the activation sequence is unknown. Also, the position and angle of the ventricles in the torso are expected to impact the T wave by modifying the proximity of each electrode to the regions of the ventricles (as shown by Nguyên et al. (2015)), but these were not investigated here due to the technical challenges. Finally, simulation studies have shown that contraction and deformation of the ventricles effects T wave amplitude and/or time to T peak (Keller et al. 2011; Smith, Buist, and Pullan 2003).

Chapter 7

Simulating the early repolarization pattern

Models of activation and repolarization of the human ventricles and the resultant ECGs were shown in chapter 5 and chapter 6, respectively, both benefiting from the performance improvements made in the software environment in chapter 4. This chapter builds directly on the models of these earlier chapters to investigate the mechanism of an ECG sign called early repolarization pattern. It presents a study of the impact on the ECG waveforms of a localized region of altered electrophysiology based on the most commonly concomitant genetic mutations.

7.1 Introduction

The early repolarization (ER) pattern (ERP) was first described eighty years ago in terms of slurring and notching in the QRS complex (Shipley and Hallaran 1936). In the intervening years a variety of ECG features have been described as ERP, including J point elevation, ST segment elevation (without chest pain), terminal QRS slur, and terminal QRS notch (Maury and Rollin 2013). In the presence of other fac-

tors such as unexplained ventricular fibrillation (VF), an ERP might suggest early repolarization syndrome (ERS).

Until recently the ERP was considered benign, but in the past decade a correlation has been detected between ER and negative outcomes. This discovery began with a landmark study that reported the ERP was more common in survivors of idiopathic ventricular fibrillation than in control patients (Haïssaguerre et al. 2008). Many studies subsequently showed a link between ERP and an increased risk of one or more of the following: cardiac arrhythmia, unexpected death, cardiac death, and all-cause death (Haruta et al. 2011; Junttila et al. 2014; Olson et al. 2011; Rollin et al. 2012; Rosso et al. 2008; Sinner et al. 2010; Tikkanen et al. 2009, 2011).

Four recent consensus reports made recommendations to address the problem of multiple definitions of the ERP (Antzelevitch et al. 2016; Macfarlane et al. 2015; Patton et al. 2016; Priori et al. 2013). Priori and colleagues recommended that ERP can be diagnosed in the presence of J point elevation $\geq$1 mm in two or more contiguous inferior and/or lateral leads. Macfarlane and colleagues proposed that ERP is present when there is an end-QRS notch or slur, the J point is at least 0.1 mV in two or more contiguous leads (excluding V1 to V3), and the QRS width is less than 120 ms. Patton and colleagues proposed that ERP is 'any of the following: ST segment elevation in the absence of chest pain, terminal QRS slurring, or terminal QRS notching'. Antzelevitch and colleagues proposed a score system for diagnosis of ERS with criteria including a notch or J wave $\geq$0.1 mV in two or more inferior and/or lateral leads (excl. V1–V3), and horizontal/descending ST segment, as well as additional clinical outcomes including cardiac arrest/syncope or family history.

The prevalence of ERP varies according to the definition and population, but has been reported between 1 % and 18 % (Patton et al. 2016). Evidence suggests that the morphology of the ST segment indicates relative risk: a J wave followed

by an ascending ST segment is considered benign but a horizontal or descending ST segment is associated with arrhythmic risk (Adler et al. 2013; Tikkanen et al. 2011).

There is still uncertainty in the biological basis of the ERP, specifically whether the abnormality in the terminal part of the QRS complex (J-wave) is indeed due to early repolarization or in fact an area of delayed conduction (Meijborg et al. 2016; Wellens 2008). The repolarization hypothesis suggests that the ERP is due to a heterogeneity in the refractoriness gradient due to increased net outward current, probably due to decreased inward sodium (I_{Na}) or calcium current (I_{CaL}) and/or increased transient outward potassium current (I_{to}) (Patton et al. 2016). ERS has been associated with mutation(s) in: KCNJ8, a gene involved in the $I_{\mathrm{K,ATP}}$ channel (Haïssaguerre et al. 2009; Medeiros-Domingo et al. 2010); three of the genes encoding 'L type calcium' current (I_{CaL}) subunits (Burashnikov et al. 2010); and SCN5A which encodes the I_{Na} channel (Watanabe et al. 2011). The depolarization hypothesis proposes that slowed conduction, especially in the right ventricular outflow tract, might be responsible for the J-wave (Patton et al. 2016).

The transmural I_{to} gradient has been suggested to be responsible for the J wave (Yan and Antzelevitch 1996). Experiments in canine LV wedge support the hypothesis that an I_{to} gradient is responsible for the ERP by accentuating the AP notch (Koncz et al. 2014). Similarly, accentuated J waves were measured in the inferior LV wall in patients, where I_{to} is more expressed (Nakagawa et al. 2014). Electrocardiographic imaging (ECGI) of two patients with ERP (Ghosh et al. 2010) showed relatively normal activation of the ventricles, although with possible regions of conduction delay in the mid-lateral RV. Both patients had localized areas of abnormally short action-recovery intervals (ARIs, a surrogate of APD) in the mid-lateral RV and inferior-basal RV.

In this chapter I will investigate the effect on the ECG of including localized re-

gions of abnormal APD shortening in the human ventricles and whether these are sufficient to recover the clinical ERP phenotype. I will do this by beginning with a computational human ventricle-torso ECG-forward-problem model with realistic activation sequence (chapter 5) and dispersion of repolarization (chapter 6) to generate a baseline ECG in healthy conditions. Such a multiscale heart-torso model will then be utilized by introducing regions of altered electrophysiology based on ER regions as reported by Ghosh et al. (2010) and repolarization abnormalities based on the most prevalent genetic phenotypes of the disease, and analyze their effects on the human ECG.

7.2 Materials and methods

7.2.1 Model construction

The modelling techniques have been previously described in the preceding chapters of this thesis. The tetrahedral mesh and mathematical and software techniques were as described in chapter 3. The activation sequence was as described in chapter 5. Heterogeneous electrophysiology in repolarization was included in the ventricles following the optimal reconstruction of T wave morphology and polarities as described in the previous chapter, specifically apico-basal and transmural gradients in APD (APD longer in the endocardium than the epicardium, and in the base than the apex).

Additionally, for the aims of this chapter, a transmural gradient in g_{to} was included to reproduce the differing degree of AP notch between endocardium and epicardium. A linear gradient in this parameter was imposed between the published values for the cell model (ten Tusscher and Panfilov 2006), specifically $0.073\,\text{nS}\,\text{pF}^{-1}$ on the endocardium and $0.294\,\text{nS}\,\text{pF}^{-1}$ on the epicardium. The inclusion of this transmural gradient in I_{to} in addition to heterogeneities in I_{Ks} did not

alter the results of ventricular repolarization described in the previous chapter. As shown below, it only slightly affects the APD at the cellular level.

7.2.2 Defining a region with abnormal electrophysiology

The localized region of altered (pathological) electrophysiology was selected in the right ventricle (fig. 7.1) approximately corresponding to the area of short APD detected in Ghosh et al. (2010). The two patients exhibited similar position of the ER region. The volume was a wedge in the basal RV spanning the majority of the free wall circumferentially, and spanning the free wall transmurally. It contained about 4.3 % of the nodes in the ventricles.

In order to evaluate the role of I_{to} in accentuating the AP notch and its impact on the clinical manifestation of ERP, the I_{to} conductance (g_{to}) was set to either its default value, four times the default value, or ten times its default value. To account for additional APD shortening in the diseased cases, a reduced I_{CaL} current was considered due to its prevalence as a genetic factor underpinning the disease (Antzelevitch et al. 2016), with g_{CaL} set to either the default value, 50 % of the default value, or 25 % of the default value. These factors were chosen to generate a range of action potential notch depths and shortening in APD (figs. 7.6 and 7.9). Exploring the parameter space more widely was not possible due to the computational requirement of each simulation instance.

7.2.3 ECG signal analysis

Simulated 12-lead ECGs with each set of g_{to} and g_{CaL} parameters are compared to a 'baseline' ECG from an identical simulation with no ER region. The differences between each lead of an ECG and the corresponding lead in the baseline lead were calculated and are referred to below as 'residuals'. Also, the total difference be-

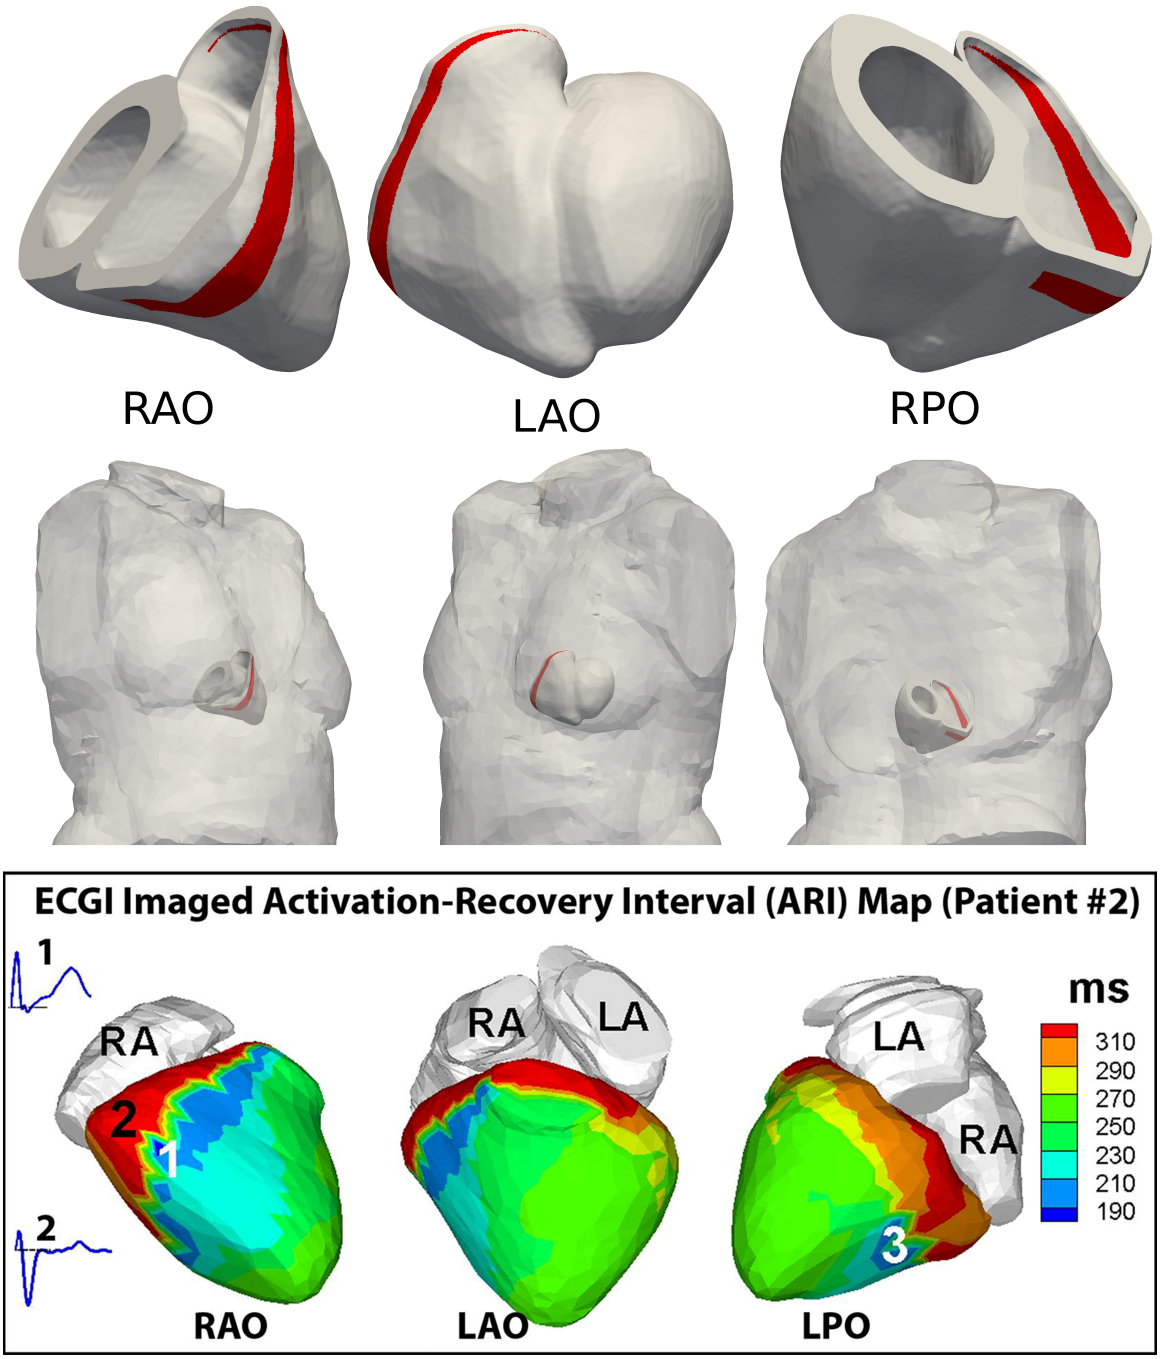

Figure 7.1: Above: Visualizations of the volume in the model ventricles demarcated as having altered electrophysiology (red). About 4.3 % of the ventricular nodes were selected. The volume spans the RV free wall. RPO=right posterior oblique, RAO=right anterior oblique, LAO=left anterior oblique. Below: ARI map for patient #2 from Ghosh et al. (2010, fig. 4) (reproduced with permission from Elsevier). The region of short APD in the RV free wall (blue) was translated to the model. LPO=left posterior.

tween the signals was quantified using the root mean square deviation (RMSD), a measure of sample standard deviation.

The residuals and RMSD are sensitive to time misalignment of the signals, so interpolation (MATLAB `interp1` function with a spline) followed by cross-correlation (MATLAB `xcorr` function) was used to estimate the lead/lag of an ECG where necessary. The cross-correlation was calculated for all 12 leads simultaneously to produce a single estimate of lead/lag with maximal correlation in all leads.

7.3 Results

7.3.1 Increased $I_{\mathbf{to}}$ in regions of localized ERP does not account for J wave elevation, but minor QRS slurring

QRS complexes are shown for a four-fold and ten-fold increase in g_{to} in the 'ER region' in fig. 7.2a. Without aligning the signals, the residuals are dominated by a small lag in the increased-g_{to} signals (fig. 7.3a). Cross-correlation suggested a lag of about 0.7 ms (table 7.1) in both signals. The residuals after alignment of the signals by these lags are shown in fig. 7.2b. The largest residuals were in V2 and neighbouring leads, the maximum being about 0.05 mV (g_{to}×10). RMSDs were of the order of 1 µV for the ECGs from both simulations (fig. 7.4). The mean residual for all leads was positive for both simulations, of the order of 1 µV.

		g_{to}		
		×1	×4	×10
g_{CaL}	×1		−0.7	−0.7
	×0.5	−0.7	−0.7	0.0
	×0.25	−0.7	0.0	−0.6

Table 7.1: Leads/lags of QRS complexes (in ms to 1 d.p.) for each simulation over the baseline, estimated using cross-correlation. Negative numbers indicate lags.

(a) Grid spacing: 40 ms, 0.1 mV

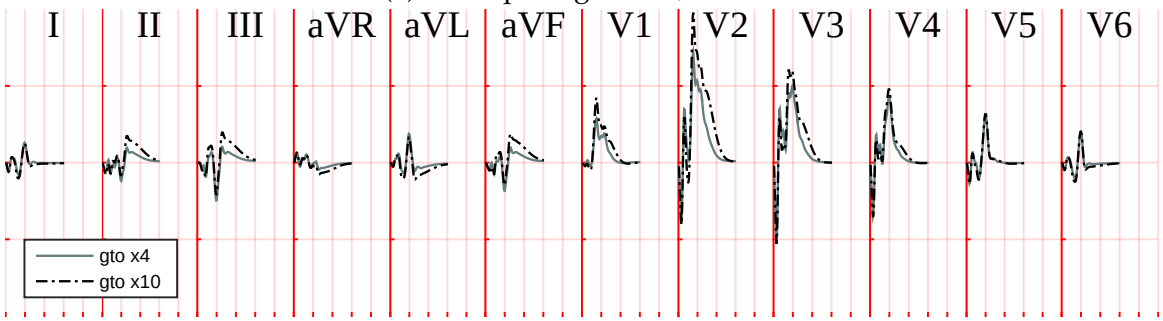

(b) Grid spacing: 40 ms, 0.025 mV

Figure 7.2: (a) QRS complexes for a baseline simulation, with a region with $4\times g_{\text{to}}$, and with a region with $10\times g_{\text{to}}$. (b) Difference between signal in each increased-g_{to} simulation and the baseline. The signals were first advanced by the lags shown in table 7.1.

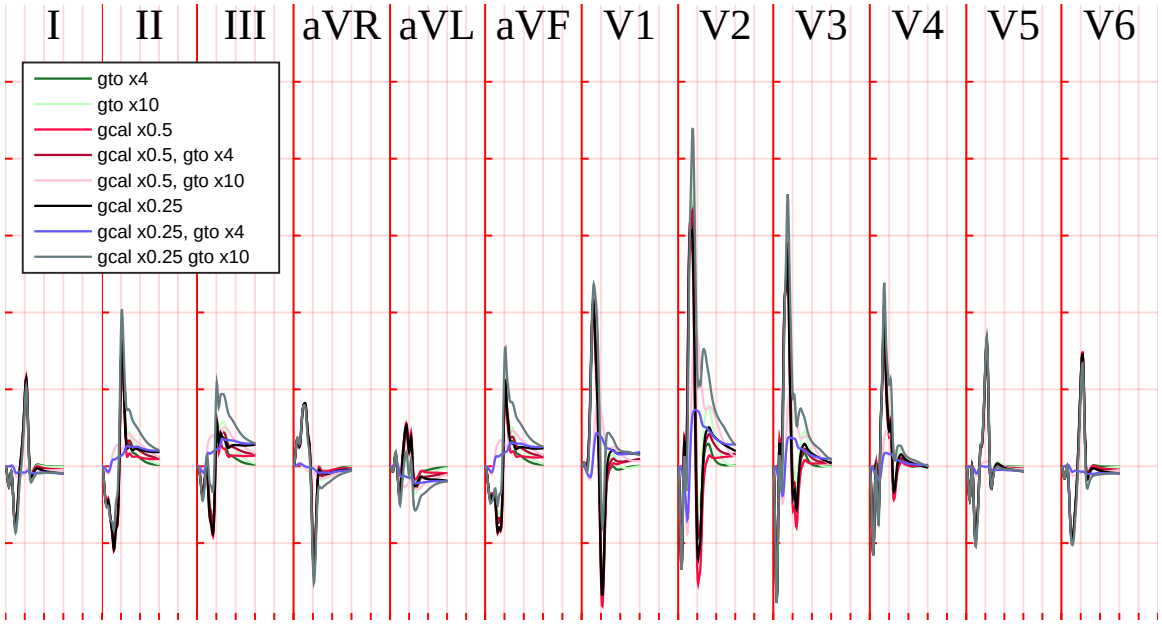

(a) Unaligned

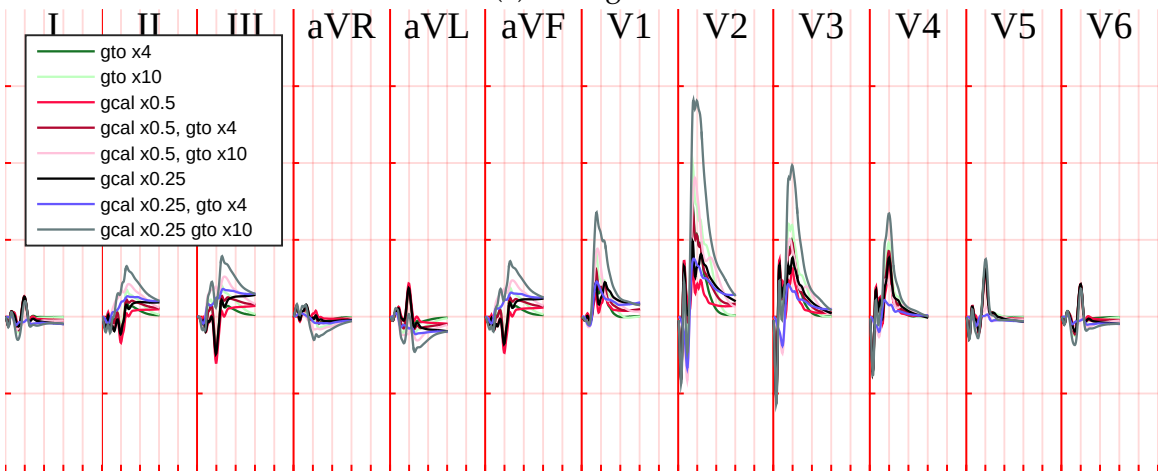

(b) Aligned

Figure 7.3: Difference between ECGs simulated with a region of altered g_{to} and/or g_{CaL} and baseline. Only the first 120 ms are shown to highlight the QRS complex. The residuals for 'g_{CaL} x0.25, g_{to} x4' and 'g_{CaL} x0.25, g_{to} x10' (lilac and grey) suggest a time lag between the signals and baseline. Compare with the residuals after alignment shown in fig. 7.8b. Grid spacing: 40 ms, 0.025 mV.

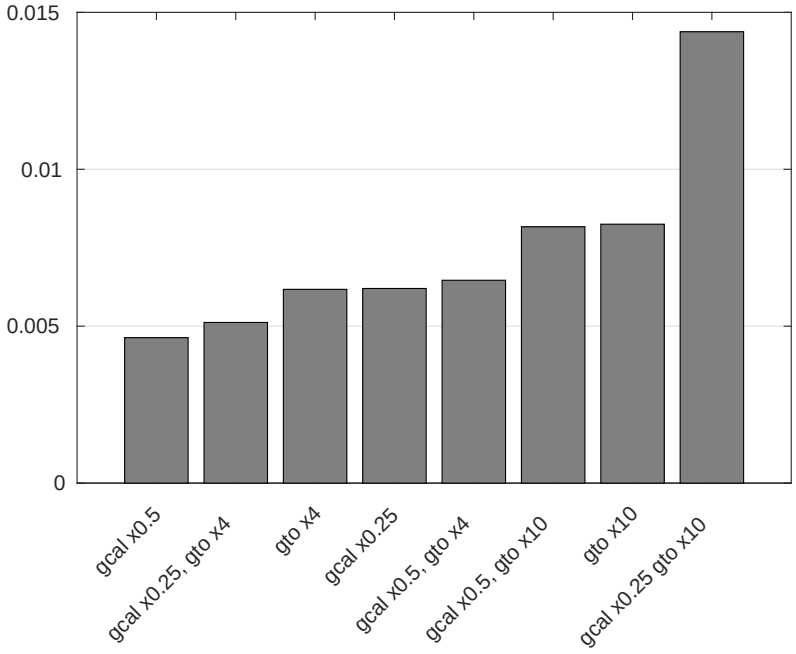

Figure 7.4: RMSD (mV) between the QRS complexes in each simulation and the baseline after aligning the time series using cross-correlation. The three largest deviations occur in simulations with a ten-fold increase in g_{to}.

Action potentials for these simulations recorded at a pair of locations spanning the RV wall in the ER region (fig. 7.5) are shown in fig. 7.6. Larger settings for the g_{to} parameter in the ER region were associated with a deeper phase 1 notch and a smaller 'spike' voltage. The other phases (upstroke, plateau, and repolarization) were relatively unaffected.

7.3.2 Deeper AP notches due to secondary I_{CaL} further increase late-QRS manifestations of ERP

Under all the considered scenarios, including decreased g_{CaL} as well as increased g_{to}, differences between the simulations with altered electrophysiology in the basal RV region and the baseline were small in the QRS complex (fig. 7.7). A closer examination of signal residuals (fig. 7.3a) suggested that the QRS complexes in several of the simulations were delayed compared to the baseline.

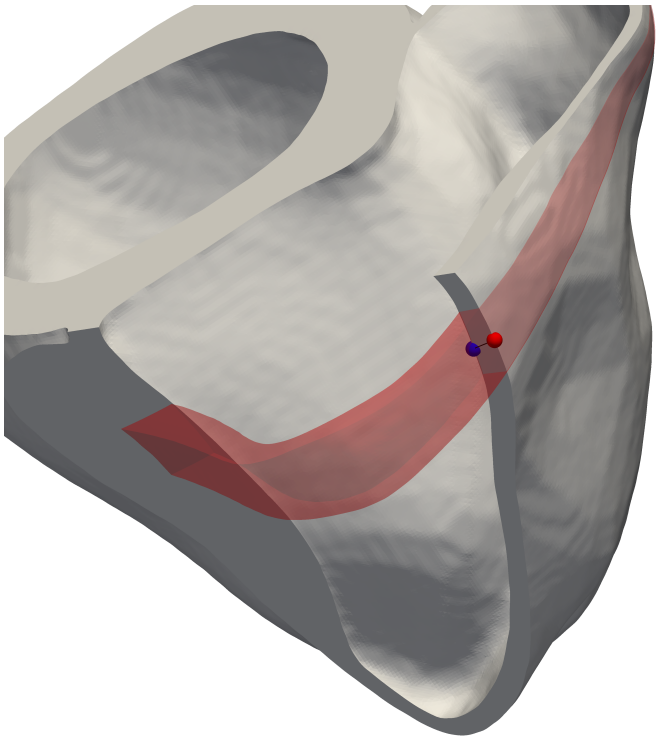

Figure 7.5: Visualization of the locations of the endocardial (blue sphere) and epicardial (red sphere) cells in the RV free wall from which APs were sampled. The ER region is shown in red, and some of the RV wall has been removed for clarity.

After interpolation, cross-correlation of the twelve leads suggested the leads and lags in table 7.1. Residuals after alignment of the signals with these lags are shown in fig. 7.8. As above, when only increasing g_{to}, the later parts of the QRS complex (after about 35 ms, when the waveforms are returning to zero) were slightly more positive in most leads. Both increasing g_{to} and reducing g_{CaL} contributed to this sign. The largest residuals were again in V2 (max. 0.07 mV) and neighbouring leads.

Action potentials from the ER region (fig. 7.5) for the six simulations with reduced g_{CaL} are shown in fig. 7.9. Reduced g_{CaL} causes a reduction of APD and phase 2 (plateau) potential. The effect of larger g_{to} on the phase 1 notch is more pronounced in combination with a greater reduction in g_{CaL}.

The root mean square deviations (RMSD) between each aligned QRS complex

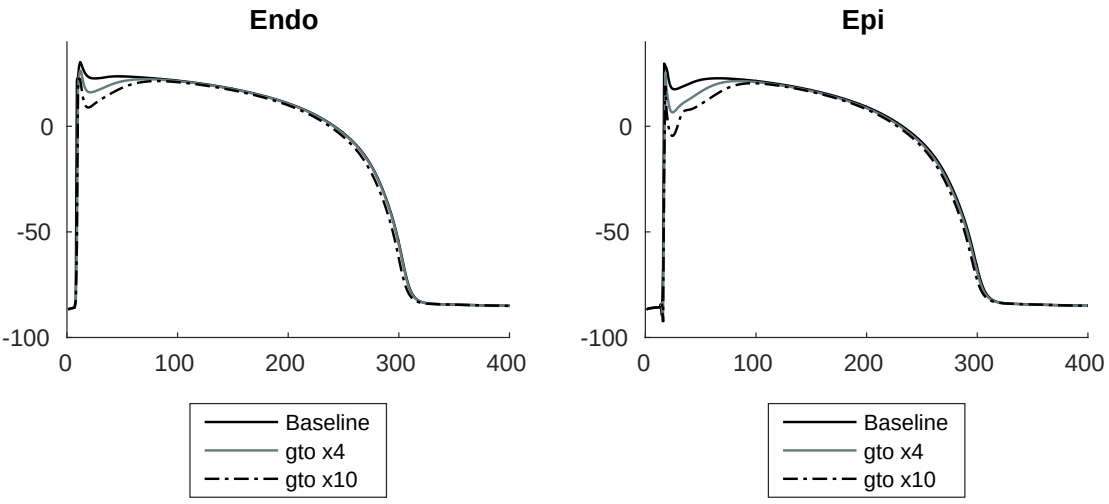

Figure 7.6: Action potentials for the baseline and two simulations with increased g_{to}, on the endocardium (left) and epicardium (right). See fig. 7.5 for the sampling locations. Horizontal axes represent time in ms, vertical axes transmembrane voltage in mV.

and the baseline are shown in fig. 7.4. In agreement with the AP recordings, the manifestations of ER on the QRS complex (as quantified by the RMSD) increased as the severities of both electrophysiological changes in the impacted region were increased, and reflected in the deeper AP notches in the epicardial region.

7.3.3 Additional APD shortening due to reduced I_{CaL} in ER impacts T wave shape and ST segment

Simulated T waves are shown in fig. 7.10. Changes to T wave shape are most apparent in leads V1 and V2. Times of T peak were practically unchanged, the largest mean shift being only (0.14 ± 0.11) ms (in 'g_{CaL}×0.25, g_{to}×10'). Some ST segment elevation is visible, especially in leads II and III. RMSDs for each signal compared to baseline are shown in fig. 7.11. The grouping of the results depending on the value of g_{CaL} in the ER region suggest that the ST segment and T wave are predominantly impacted by it, not g_{to}.

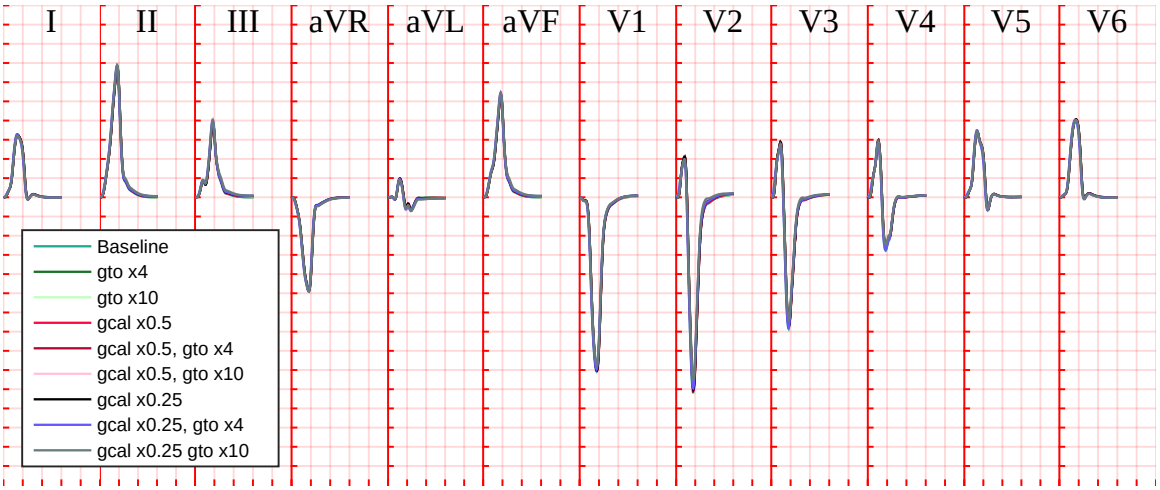

Figure 7.7: QRS complexes for a baseline simulation and eight including a region of altered electrophysiology in g_{to} and g_{CaL}.

The residuals (fig. 7.12) show that in most leads the T wave is impacted mainly before the T peak (which occurs at about 170 ms to 190 ms in the detail shown, 290 ms to 310 ms in the simulation). There is generally an increase (more positive) in ST segment and T wave (7/12 leads) proportional to the degree of g_{CaL} reduction. The simulation with g_{CaL}×0.25 (black) has a flatter ST segment in V2–V4 than the baseline. Note the negative residuals, peaking at about 200 ms. Examining the APs in detail (fig. 7.13a) revealed that the endocardial cells in this simulation repolarized earlier than the epicardial cells. APs from the simulation with the same reduction in g_{CaL} and an additional ten-fold increase in g_{to} (fig. 7.13b) are typical of all the other simulations insofar as the cells repolarized within 5 ms of each-other.

Recall that a repolarization wavefront moving towards an electrode generates a negative deflection. The early endocardial repolarization resulted in a localized transmural repolarization from endocardium to epicardium at around 200 ms, generating a negative deflection in the proximal leads (V2 etc.). The transmural difference in recovery times for all simulations are shown in fig. 7.14. The 25 % g_{CaL} simulation has by far the largest transmural dispersion in RT (12 ms).

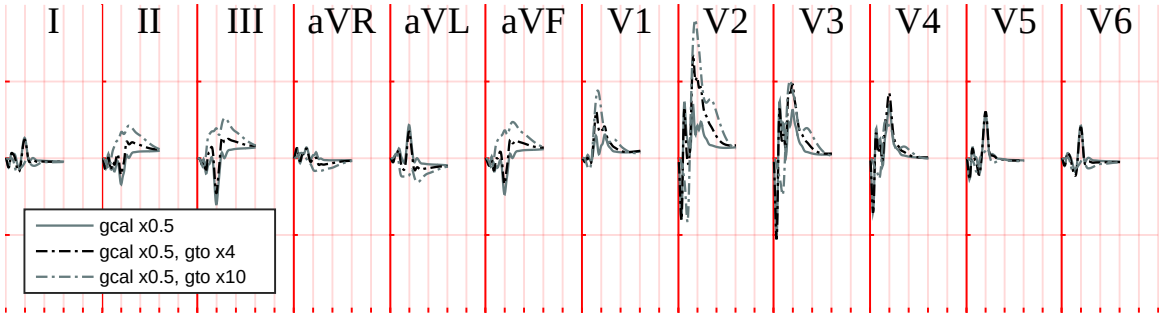

(a) 50 % of default g_{CaL}

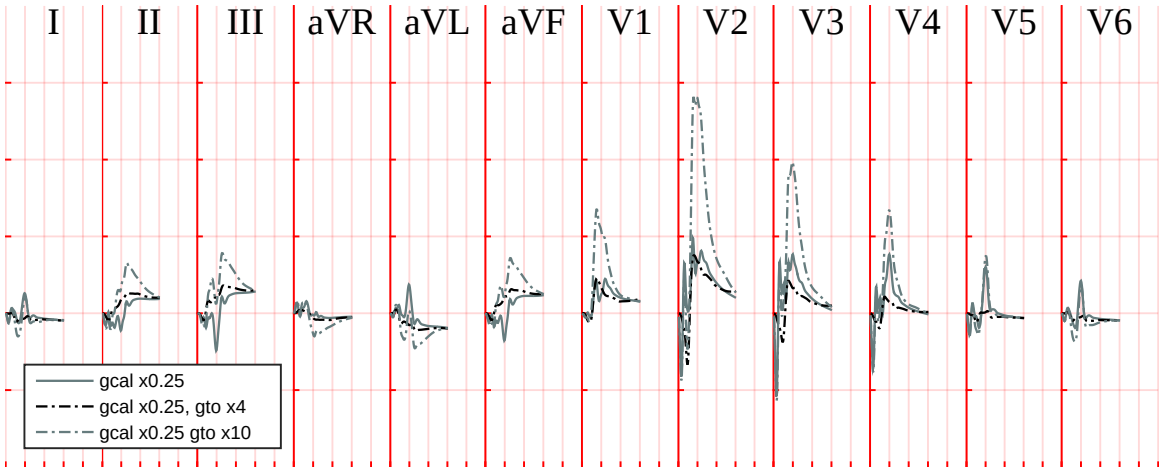

(b) 25 % of default g_{CaL}

Figure 7.8: Difference between ECGs simulated with a region of altered g_{to} and/or g_{CaL} and baseline. Only the first 120 ms are shown to highlight the QRS complex. The signals were first aligned with the lead/lag estimated using cross-correlation. Grid spacing: 40 ms, 0.025 mV.

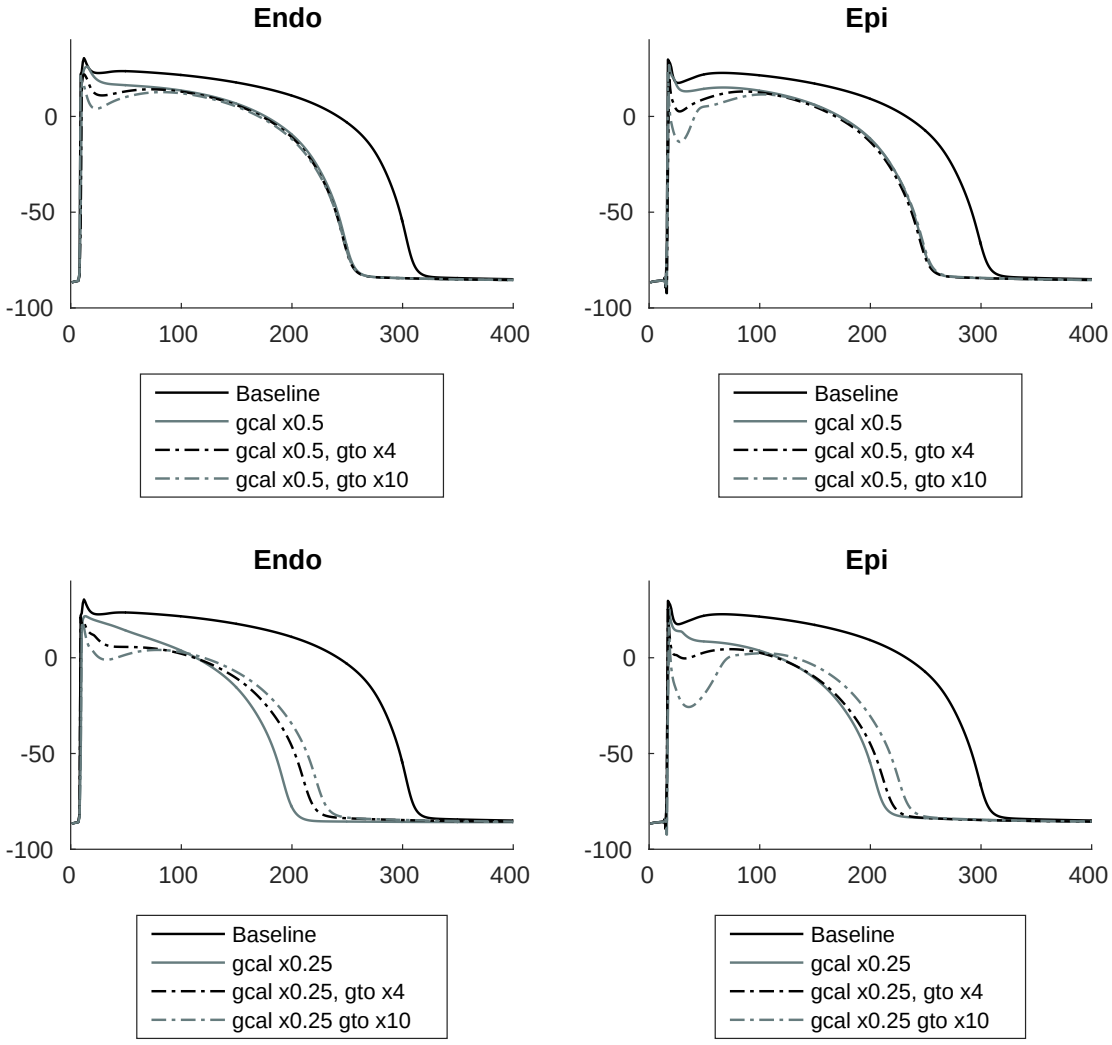

Figure 7.9: Action potentials for the baseline and six simulations with reduced g_{CaL}, on the endocardium (left) and epicardium (right). See fig. 7.5 for the sampling locations. Horizontal axes represent time in ms, vertical axes transmembrane voltage in mV.

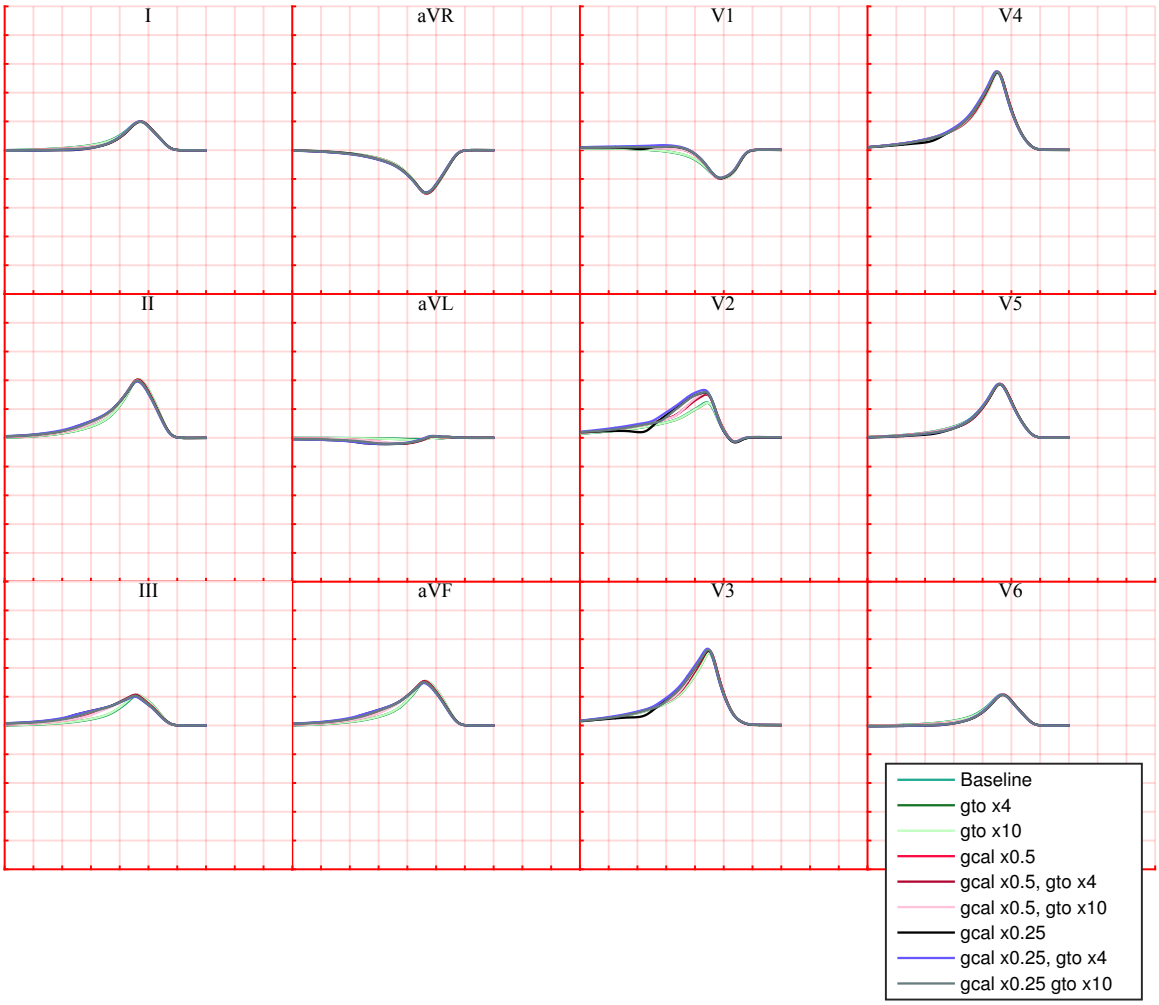

Figure 7.10: T waves for a baseline simulation and nine including a region of altered electrophysiology in g_{to} and g_{CaL}.

7.4 Discussion

In order to capture the main electrophysiological characteristics of ERP, a region of the right ventricle was assigned altered electrophysiological properties in the I_{to} and I_{CaL} currents. An identical simulation with no such region was performed as a baseline. The degree of dissimilarity between the simulation results and baseline were evaluated by examining the difference between the signals and calculating the RMSD, with signal alignment based on maximal cross-correlation. Action potentials recorded on the endocardial and epicardial surfaces in the ER region (figs. 7.6 and 7.9) showed a range of morphological changes in agreement with the

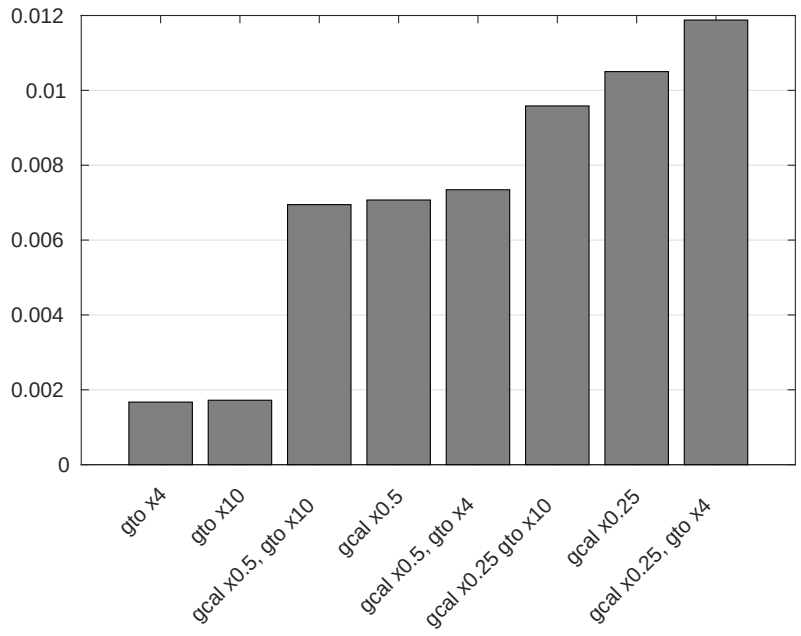

Figure 7.11: RMSDs (mV) for ST segment and T waves in each simulation and baseline. The results are grouped by the degree of change in g_{CaL}.

imposed electrophysiological abnormalities underlying the disease.

Increased g_{to} has been postulated as responsible for the J wave (Patton et al. 2016). In this study, increasing g_{to} deepened the phase 1 notch without altering the plateau or repolarization phases, and had only a minor effect on the QRS complex, slightly increasing (more positive) the QRS signal after the peak. The largest RMSDs of the QRS complexes were in simulations with a ten-fold increase in g_{to}. Nevertheless, the largest residual was only about 0.05 mV (in V2). This is less than most definitions of ERP (0.1 mV), and in any case leads V1–V3 are usually excluded (Antzelevitch et al. 2016; Macfarlane et al. 2015). Mutations in genes that result in downregulation of the calcium current have also been associated with ERS (Antzelevitch et al. 2016). Additional reduction of g_{CaL} further increased J wave elevation and QRS elevation, but still of lower magnitude than would qualify for the clinical diagnosis of ERP in the ECG.

The ST segment and T wave were mainly affected by the changes to g_{CaL} rather

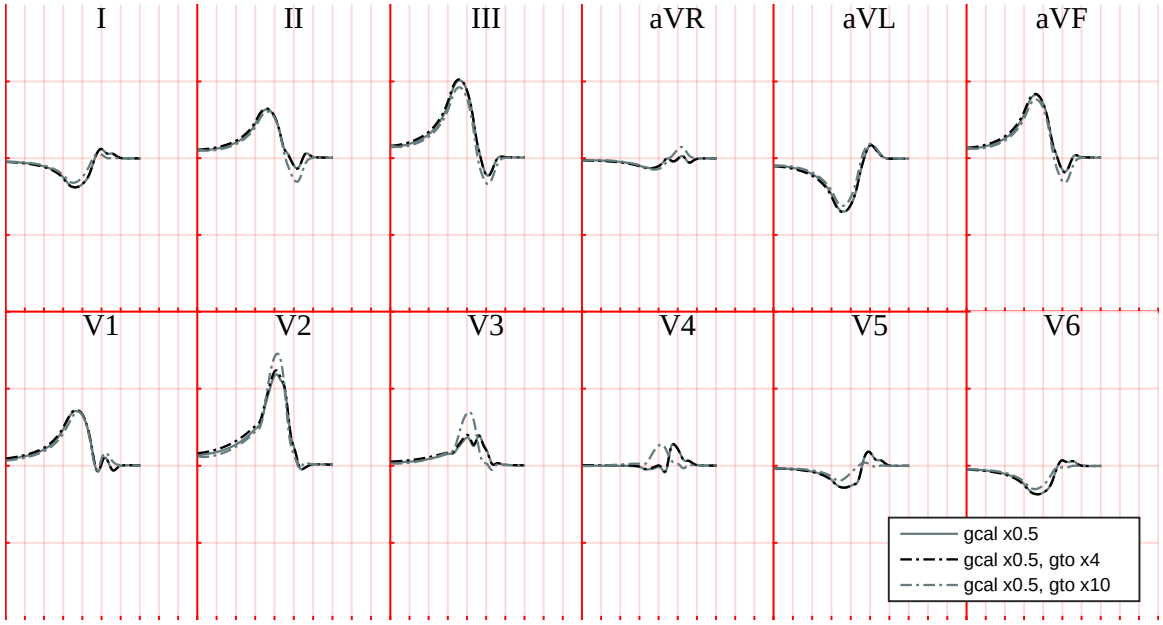

(a) 50 % of default g_{CaL}

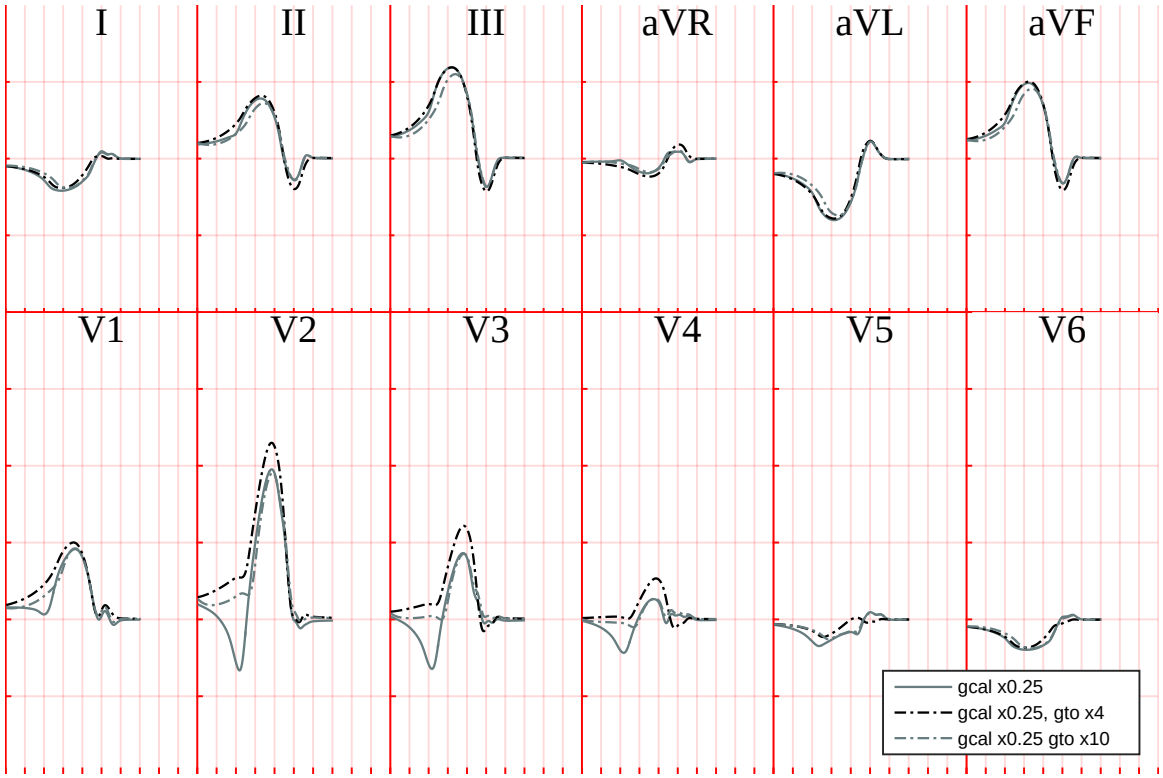

(b) 25 % of default g_{CaL}

Figure 7.12: Difference between ECGs simulated with a region of altered g_{to} and/or g_{CaL} and baseline. The first 120 ms have been removed to highlight the ST segment and T wave. Grid spacing: 40 ms, 0.025 mV.

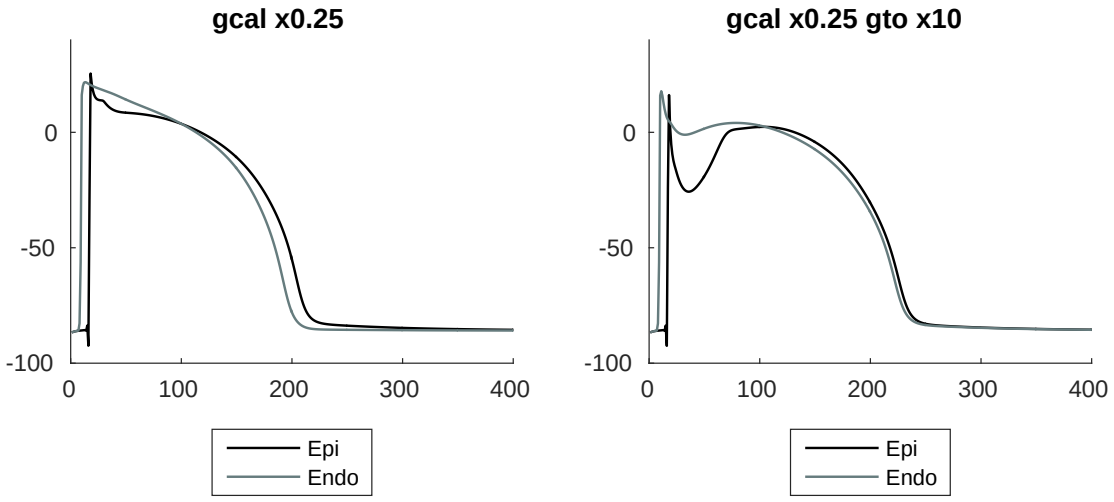

(a) 25 % of default g_{CaL} (b) 25 % of default g_{CaL} and 10×g_{to}

Figure 7.13: Action potentials from locations on the epicardial and endocardial surfaces of the ER region in two simulations. The ER region in both simulations had a 75 % reduction in g_{CaL}. (a) No change in g_{to} in the ER region. Note the slower epicardial repolarization time than in the endocardium, which was unique to this simulation. (b) Ten-fold increase in g_{to} in the ER region. Note the similar repolarization times at the epicardium and endocardium. Horizontal axes represent time in ms, vertical axes transmembrane voltage in mV. See fig. 7.5 for the sampling locations.

than g_{to}, with the largest observed differences in the leads V1–V3. This is understood since the electrodes for these leads are in relatively close proximity to the ER volume, and the T wave is a function of the plateau and repolarization phases which are substantially affected by g_{CaL} in the model. A small (< 0.05 mV) increase in most T waves was observed before the T peak without substantially changing the time to T peak. This too is intuitive because the region of reduced APD is not large enough to substantially modify the T wave, but the shorter APD provides a small, early contribution.

Other mutations associated with ERS but not investigated here are responsible for the upregulation of $I_{\mathrm{K,ATP}}$ and downregulation of I_{Na} channels, and these could be the focus of future work. A recent study by Meijborg et al. (2016) found that 'reduced sodium current in the lateral region of the heart causes inferolat-

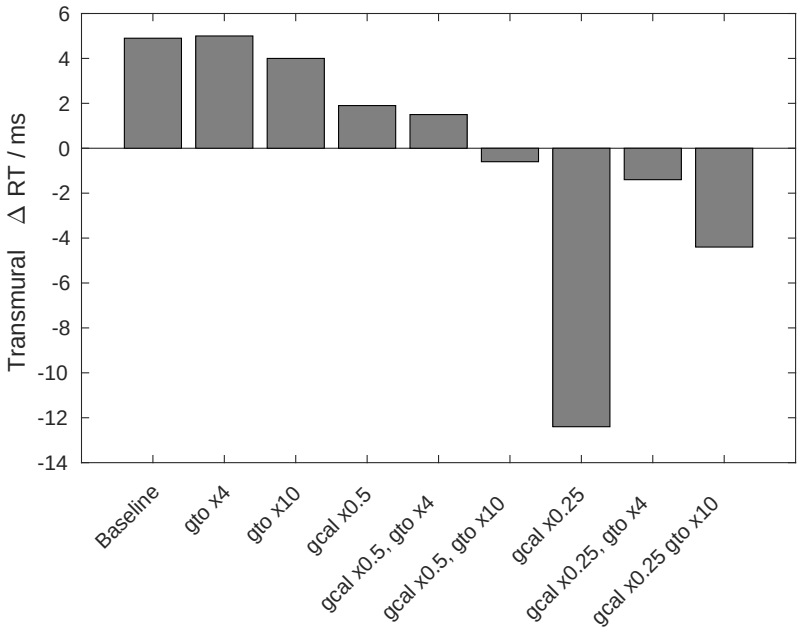

Figure 7.14: Transmural dispersion in repolarization times (RT) in the ER region for the nine simulations. Negative differences imply the endocardium repolarized first.

eral J waves on the ECG', further reinforcing the depolarization hypothesis as the mechanism underlying J point elevation in the disease. Further applications of the model might be in the investigations of arrhythmias in the disease, for example investigating the relationship between ER and the vulnerability window.

7.5 Conclusion

This investigation suggests that the changes to AP notch (g_{to}) and APD (g_{CaL}) in the demarcated region are insufficient to generate the magnitudes of features associated with the clinical diagnosis of ERP in the human ECG (QRS notching, slurring, and/or J point elevation) according to recent consensus reports (Antzelevitch et al. 2016; Macfarlane et al. 2015; Priori et al. 2013). These results suggest additional support to the depolarization over the repolarization hypothesis as the mechanism underlying J wave elevation in ERS, although further investigations are needed to

ascertain the role of other mutations in repolarization in the disease, as identified in genetic probands.

Nevertheless, the results indicate that secondary mutations in I_{CaL} might be associated with ST segment flattening in ERS which has been associated with increased risk of arrhythmic death (Tikkanen et al. 2011). These findings deserve further investigation, especially of the susceptibility of re-entry in the presence of pathological regions of altered repolarization.

Chapter 8

Conclusions and future work

8.1 Conclusions

In this thesis I have presented my contributions to multiscale mathematical models of the human ventricles and electrocardiogram using high performance computers. These include performance improvements to the ECG simulation framework, particularly in the previously neglected areas of data storage, post-processing, and data conversion. The success of these improvements made it possible to efficiently utilize HPC resources and accelerated scientific progress.

In addition to these optimization efforts I constructed and simulated models of ventricular activation and repolarization, in both normal conditions and in disease. These applications formed a natural progression due to the underlying physiology: the activation sequence is partly responsible for the repolarization sequence, and the accuracy of both are crucial to modelling their respective disorders and pathologies. These models were based on human data of multiple modalities, and the results are consistently in agreement with clinical observables.

Conclusions for each chapter are given below, followed by possible avenues for further research.

8.1.1 Improvements to simulation pipeline in HPC

Chapter 4 presents techniques and results for substantial performance improvements for large bidomain simulations in Chaste on distributed systems such as HECToR and Archer. Previously, solver performance was adequate, but other practical aspects were unacceptably slow, specifically mesh reading, data output, and post-processing. The improvements made in this these have enabled researchers access to a wider range of Chaste functionality and therefore a more complete simulation framework for cardiac simulations within tolerable resource envelopes.

The most substantial contribution to Chaste was in the way HDF5 files (the main results file format) are written to disk, including the low-level data layout. Improvements in this domain made it practically possible to utilize the complete Chaste simulation pipeline, previously impossible due to the very poor parallel scaling. Two actions were central to this result. Firstly, the data layout (the so-called chunk shape) was modified to improve efficiency when reading and writing small amounts of data, which is typical in a large parallel simulation of many processors. This resulted in a huge reduction in the time spent in post-processing simulation results and in converting results to other formats for visualization, two common requirements for researchers using Chaste. A side effect of changing the chunk shape was an increase in the time required to write results to the HDF5 file during simulation, due to the nature of the storage systems in HPC. The second action, therefore, was implementing an optional cached write mode which bundles multiple small writes into one larger one, substantially reducing the aggregate time spent writing. The overall result was no change in performance in small systems (e.g. desktop PCs), but a huge reduction in the time spent in the auxiliary stages involved in Chaste simulations.

Furthermore, bugs in the implementation of CVODE-compatible ODE solvers were identified and fixed, enabling the routine use of this solver in tissue simula-

tions, reducing total computer time by over 50 %. I implemented an adaptive time controller which improves performance in long simulations (one or more heartbeats) by fourfold by using a longer time step once the entire tissue domain has been activated, with a negligible loss of accuracy when considering the body surface potentials. Finally, two relatively small improvements (in terms of total simulation time) were made in handling meshes, as follows. Firstly, the time taken at the beginning of a simulation to read the tetrahedral mesh describing the problem domain was improved by a factor of over 23. Secondly, in mesh writing, a 'deadlock' was identified and corrected.

8.1.2 Simulating normal and pathological activation sequences and the QRS complex

Chapter 5 presents a flexible, relatively simple method for modelling the ventricular activation system and producing realistic activation sequences. In summary, endocardial nodes are stimulated with activation times proportional to the nearest 'root node', which each represent the junction of the main activation system branches with the distal Purkinje network. Multimodal data were first collected from the literature: *in vivo* and *ex vivo* ventricular activation times, body surface potential maps, and a database of control ECGs. Due to intersubject variation the data are naturally inconsistent, but the commonalities were identified and translated to the model to form a baseline, 'healthy' model. The simulated data were in broad agreement with the literature, including the simulated QRS complexes.

A sensitivity analysis determined the relationship between model parameters and ECG features. Specifically, the 'activation speed' (the constant of proportionality that determines the endocardial activation times) and intracellular myocardial conductivities impact the QRS width and amplitude (inversely with one another). The extracellular myocardial conductivities and torso conductivities mainly affect

the QRS amplitude.

The impact of the root points on the QRS complex was investigated by adjusting the locations of those on the free walls. The LV anterior branch impacted the QRS axis, manifested mainly in larger S waves in the inferior leads with a more apical location, indicating a slight shift towards the horizontal. Moving the posterior branch resulted in fractionated QRS complexes.

Knocking out one or more root points was successful in reproducing the typical QRS features of hemiblock or branch block seen in patients. In the hemiblocks these were rotations in the QRS axis. In the complete blocks the QRS complexes were wide and triangular, or small and fractionated, for the left bundle branch block and right bundle branch block, respectively, as seen in recordings from patients. Although computational investigations of left bundle branch block have been previously addressed in the literature (Okada et al. 2013; Potse et al. 2014; Villongco et al. 2014), the results in this thesis represent the first computational study of right bundle branch block and hemiblocks of human trifascicular system in quantitative and qualitative agreement with human ECG recordings.

8.1.3 Modelling APD heterogeneities and the T wave

Chapter 6 extends the preceding chapter by considering the T wave given a normal activation sequence. Heterogeneous APD distributions were modelled by imposing repolarization gradients based on human APD measurements, and the similarity of the resultant T waves to a database of those from controls was evaluated. The method of constructing these gradients has several strengths, including the fact that the APD dispersion is symmetrical and linear in slope, as well as the resulting combination of the gradients.

More specifically, the gradients considered were transmural, left-right (interventricular), and apico-basal. Each gradient was individual parametrized from

the human data, and combined with every other for a total of eight simulations. The majority of APD distributions produced abnormalities in at least one ECG lead when comparing to the control dataset, including T waves that were inverted, biphasic, notched, flat-topped, and 'camel humped', which can further assist in the interpretation of these abnormalities in the human ECG. The simulation with apico-basal and transmural gradients, however, generated T waves of normal appearance. This finding is in agreement with other simulation studies that used an interventricular gradient with LV APD >RV APD (Okada et al. 2011; Seemann et al. 2013), as shown in the human data.

8.1.4 Electrophysiological substrates of the ERP

Chapter 7 used the most realistic activation sequence and APD heterogeneities from the preceding two chapters to investigate the mechanisms of early repolarization pattern (ERP). Models were constructed to test the hypothesis that a localized region of increased transient outward current conductance (g_{to}) in the RV could be responsible for ERP (J point elevation, QRS notch, and/or QRS slur) by generating a transmural current due to increased depth of AP notch. The L type calcium current conductance (g_{CaL}) was also considered for its ability to shorten APD, as consistently reported as a genetic factor in patients with ERP.

Increasing g_{to} regionally had a small contribution to the latter part of the QRS complex, but not sufficiently to qualify as ERP according to recent expert consensuses, in agreement with a recent simulation study (Meijborg et al. 2016). Neither did it produce the APD shortening measured in patients, because unlike in a single cell setting, the surrounding tissue prevented the collapse of the AP dome. Reducing g_{CaL} had the expected effect in terms of shortening the APD, and had a slightly larger contribution to the QRS complexes than increasing g_{to} alone, but still was not sufficient to resemble ERP.

Finally, the impact of these changes on the T wave was also investigated in a novel simulation study of ERS and the T wave. Here the effects of reduced g_{CaL} were more substantial, including one case which had a flatter ST segment, which is a clinical risk factor of arrhythmic death in the disease.

8.2 Future work

8.2.1 Effect of intersubject variability in body and ventricle shape

The mesh used in this thesis was produced by a collaborator by combining a ventricular mesh with body surfaces segmented from a patient. This was the only mesh available due to the technical and resource challenges in geometry definition and mesh generation, and the scarcity and cost of acquiring new datasets. The effect on the ECG of various factors could therefore not be investigated, such as body shape and size, torso composition (proportions of fat and other tissue), and the position and orientation of the ventricles within the torso.

This research opportunity has been identified in the group, which is allocating efforts on developing a pipeline from scanning patients to segmentation, meshing, and finally simulation. We hope this endeavour will contribute to a deeper understanding of intersubject variability in terms of physiological differences on the ECG.

8.2.2 T waves in patients with conduction abnormalities or J wave syndromes

In chapter 5, results of simulations of conduction block were shown that successfully reproduced typical QRS complex features, and in the following chapter T waves with the normal polarity were simulated. A possible extension might be

to investigate the T waves in simulated patients with conduction abnormalities.

Similarly, in chapter 7, the effects on the T wave of modified densities of two ion channels related to genes with relatively common mutations in early repolarization syndrome were investigated. Future work could similarly evaluate the impact of rarer mutations, such as in $I_{K,ATP}$, on the T wave.

8.2.3 Simulations with alternative cardiomyocyte models

This thesis exclusively used the model of ten Tusscher and Panfilov (2006) because it offers an attractive compromise between complexity (in terms of biophysical detail, and numbers of currents and parameters) and speed of execution. Many other human ventricular models exist, however, with various strengths. In particular, some research questions require currents not explicitly represented in the ten Tusscher and Panfilov model, such as the ATP-sensitive potassium channel ($I_{K,ATP}$). Based on the modelling infrastructure developed in this thesis, work has begun in the research group to perform similar simulation studies with the recent O'Hara et al. (2011) model, currently considered the gold standard for human ventricular electrophysiology.

8.2.4 Multi-scale models of acute ischaemia and/or infarction

The activation and repolarization model presented in this thesis has great potential as a tool for investigating other pathologies. Coronary heart disease, also known as ischaemic heart disease, is a group of diseases including angina, myocardial infarction, and sudden cardiac death. It is estimated to kill 8.2 million people per year worldwide, or about 17 % of all deaths (GBD Collaborators 2015). In the UK it is responsible for 92 300 deaths per year (Finegold, Asaria, and Francis 2013).

An exciting application in this domain for the models presented in this the-

sis would be determining the contribution of acute ischaemia and/or infarction to the ECG. A computational study of this kind would have unique strengths in terms of exploring the space of possible injuries, e.g. size and location of ischaemic zone/infarct, presence and width of border zone, proportion of viable tissue, etc. Results could be validated by translating the injured region from medical images, e.g. contrast-enhanced MRI (Dall'Armellina et al. 2012) and comparing simulated ECGs to the corresponding clinical ECGs recorded at the time of image acquisition.

8.2.5 Tools for computational medicine: pharmacological screening and optimization of procedures and devices

Pharmacological therapy is commonly used for cardiac arrhythmias, but one of its main limitations is that it can result in the development of new or even worse ventricular arrhythmias (drug cardiotoxicity), particularly in diseased patients with ischaemia (Nagra, Ledley, and Kantharia 2005). Drug efficacy and safety are routinely screened using non-diseased animal models, and there is a need for new human-based methods to improve current therapies and accelerate the drug development pipeline to overcome drug cardiotoxicity. Modelling frameworks like the one developed and utilized in this thesis may be a useful tool in establishing *in silico* clinical trials to achieve these goals (Corrias et al. 2010; Zemzemi et al. 2012). Similarly, patient-specific computer models are emerging as the tools of computational medicine. Recent applications include improving the efficacy of treatment of atrial fibrillation (Bayer et al. 2016) and reducing injury caused by implantable cardioverter defibrillators (Noro et al. 2016).

Bibliography

Adler, A. et al. (2013). "What Do We Know About the "Malignant Form" of Early Repolarization?" In: *Journal of the American College of Cardiology* 62.10, pp. 863–868 (cit. on p. 131).

Akar, F. G. et al. (2002). "Unique Topographical Distribution of M Cells Underlies Reentrant Mechanism of Torsade de Pointes in the Long-QT Syndrome". In: *Circulation* 105.10, pp. 1247–1253 (cit. on p. 14).

Antzelevitch, C. (2001). "Transmural dispersion of repolarization and the T wave". In: *Cardiovascular Research* 50.3, pp. 426–431 (cit. on p. 27).

Antzelevitch, C. et al. (1991). "Heterogeneity within the ventricular wall". In: *Circ Res* 69.1, p. 427 (cit. on p. 14).

Antzelevitch, C. et al. (2016). "J-Wave syndromes expert consensus conference report: Emerging concepts and gaps in knowledge". In: *Europace*, euw235 (cit. on pp. 130, 133, 145, 148).

Anyukhovsky, E. P. et al. (1999). "The Controversial M Cell". In: *Journal of Cardiovascular Electrophysiology* 10.2, pp. 244–260 (cit. on p. 14).

Arteyeva, N. V. et al. (2015). "Action potential duration gradients in the heart ventricles and the cardiac electric field during ventricular repolarization (a model study)". In: *Journal of Electrocardiology* 48.4, pp. 678–685 (cit. on p. 127).

Balay, S. et al. (1997). "Efficient Management of Parallelism in Object Oriented Numerical Software Libraries". In: *Modern Software Tools in Scientific Computing*.

Ed. by E. Arge, A. M. Bruaset, and H. P. Langtangen. Birkhäuser Press, pp. 163–202 (cit. on p. 44).

Balay, S. et al. (2014). *PETSc Users Manual*. Tech. rep. ANL-95/11 - Revision 3.5. Argonne National Laboratory (cit. on p. 44).

Bauer, A. et al. (2002). "Effects of the IKr-blocking agent dofetilide and of the IKs-blocking agent chromanol 293b on regional disparity of left ventricular repolarization in the intact canine heart". In: *Journal of cardiovascular pharmacology* 39.3, pp. 460–467 (cit. on p. 16).

Bayer, J. D. et al. (2016). "Novel Radiofrequency Ablation Strategies for Terminating Atrial Fibrillation in the Left Atrium: A Simulation Study". In: *Frontiers in Physiology* 7 (cit. on p. 158).

Beeler, G. W. and H. Reuter (1977). "Reconstruction of the action potential of ventricular myocardial fibres". In: *The Journal of Physiology* 268.1, pp. 177–210 (cit. on p. 32).

Bernabeu, M. O., M. Wallman, and B. Rodriguez (2010). "Shock-induced arrhythmogenesis in the human heart: A computational modelling study". In: *2010 Annual International Conference of the IEEE Engineering in Medicine and Biology Society (EMBC)*, pp. 760–763 (cit. on p. 2).

Bernabeu, M. O. and D. Kay (2011). "Scalable parallel preconditioners for an open source cardiac electrophysiology simulation package". In: *Procedia Computer Science* 4, pp. 821–830 (cit. on p. 44).

Bernabeu, M. O. et al. (2014). "Chaste: a case study of parallelisation of an open source finite-element solver with applications to computational cardiac electrophysiology simulation". In: *International Journal of High Performance Computing Applications* 28.1, pp. 13–32 (cit. on pp. 41, 44, 49, 50, 79).

Bishop, M. J. et al. (2009). "Comparison of Rule-Based and DTMRI-Derived Fibre Architecture in a Whole Rat Ventricular Computational Model". In: *Functional*

Imaging and Modeling of the Heart. Ed. by N. Ayache, H. Delingette, and M. Sermesant. Lecture Notes in Computer Science 5528. Springer Berlin Heidelberg, pp. 87–96 (cit. on p. 37).

Boukens, B. J. D. et al. (2009). "Developmental Basis for Electrophysiological Heterogeneity in the Ventricular and Outflow Tract Myocardium As a Substrate for Life-Threatening Ventricular Arrhythmias". In: *Circulation Research* 104.1, pp. 19–31 (cit. on pp. 14, 27).

Boukens, B. J. and V. M. Christoffels (2012). "Electrophysiological Patterning of the Heart". In: *Pediatric Cardiology* 33.6, pp. 900–906 (cit. on p. 14).

Boukens, B. J. et al. (2015). "Transmural APD gradient synchronizes repolarization in the human left ventricular wall". In: *Cardiovascular Research* 108.1, pp. 188–196 (cit. on pp. 108, 109).

Boulakia, M. et al. (2010). "Mathematical Modeling of Electrocardiograms: A Numerical Study". In: *Annals of Biomedical Engineering* 38.3, pp. 1071–1097 (cit. on p. 104).

Bousseljot, R., D. Kreiseler, and A. Schnabel (2009). "Nutzung der EKG-Signaldatenbank CARDIODAT der PTB über das Internet". In: *Biomedizinische Technik/Biomedical Engineering* 40.s1, pp. 317–318 (cit. on pp. 26, 83, 85, 99, 103, 105, 108, 116, 125).

Brahmajothi, M. V. et al. (1997). "Regional Localization of ERG, the Channel Protein Responsible for the Rapid Component of the Delayed Rectifier, K+ Current in the Ferret Heart". In: *Circulation Research* 81.1, pp. 128–135 (cit. on p. 15).

Bryant, S. M. et al. (1998). "Regional differences in the delayed rectifier current (IKr and IKs) contribute to the differences in action potential duration in basal left ventricular myocytes in guinea-pig". In: *Cardiovascular Research* 40.2, pp. 322–331 (cit. on p. 15).

Bueno-Orovio, A. et al. (2012). "In Vivo Human Left-to-Right Ventricular Differences in Rate Adaptation Transiently Increase Pro-Arrhythmic Risk following Rate Acceleration". In: *PLoS ONE* 7.12, e52234 (cit. on pp. 16, 108, 109).

Bueno-Orovio, A. et al. (2014). "Fractional diffusion models of cardiac electrical propagation: role of structural heterogeneity in dispersion of repolarization". In: *Journal of The Royal Society Interface* 11.97, p. 20140352 (cit. on pp. 27, 114).

Burashnikov, E. et al. (2010). "Mutations in the cardiac L-type calcium channel associated with inherited J-wave syndromes and sudden cardiac death". In: *Heart Rhythm* 7.12, pp. 1872–1882 (cit. on p. 131).

Burgess, M. J. et al. (1972). "The sequence of normal ventricular recovery". In: *American Heart Journal* 84.5, pp. 660–669 (cit. on pp. 15, 16).

Caldwell, B. J. et al. (2009). "Three Distinct Directions of Intramural Activation Reveal Nonuniform Side-to-Side Electrical Coupling of Ventricular Myocytes". In: *Circulation: Arrhythmia and Electrophysiology* 2.4, pp. 433–440 (cit. on pp. 14, 38, 40).

Cardone-Noott, L. et al. (2014). "A Computational Investigation into the Effect of Infarction on Clinical Human Electrophysiology Biomarkers". In: *Computing in Cardiology*. Vol. 41, pp. 673–676 (cit. on p. 6).

Cardone-Noott, L. et al. (2016). "Human ventricular activation sequence and the simulation of the electrocardiographic QRS complex and its variability in healthy and intraventricular block conditions". In: *Europace* 18.suppl 4, pp. iv4–iv15 (cit. on p. 5).

Carro, J. et al. (2011). "A human ventricular cell model for investigation of cardiac arrhythmias under hyperkalaemic conditions". In: *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences* 369.1954, pp. 4205–4232 (cit. on p. 33).

Cheng, J. et al. (1999). "Heterogeneous distribution of the two components of delayed rectifier K+ current: a potential mechanism of the proarrhythmic effects of methanesulfonanilideclass III agents". In: *Cardiovascular Research* 43.1, pp. 135–147 (cit. on p. 16).

Clayton, R. et al. (2011). "Models of cardiac tissue electrophysiology: Progress, challenges and open questions". In: *Progress in Biophysics and Molecular Biology* 104.1–3, pp. 22–48 (cit. on p. 40).

Clerc, L. (1976). "Directional differences of impulse spread in trabecular muscle from mammalian heart." In: *The Journal of Physiology* 255.2, pp. 335–346 (cit. on pp. 37, 39).

Cohen, J. (1988). *Statistical power analysis for the behavioral sciences*. OCLC: 17877467. Hillsdale, N.J.: L. Erlbaum Associates (cit. on p. 72).

Cooper, J., R. J. Spiteri, and G. R. Mirams (2015). "Cellular cardiac electrophysiology modeling with Chaste and CellML". In: *Frontiers in Physiology* 5 (cit. on p. 44).

Corrias, A. et al. (2010). "Arrhythmic risk biomarkers for the assessment of drug cardiotoxicity: from experiments to computer simulations". In: *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences* 368.1921, pp. 3001–3025 (cit. on p. 158).

Dall'Armellina, E. et al. (2012). "Cardiovascular magnetic resonance by non contrast T1-mapping allows assessment of severity of injury in acute myocardial infarction". In: *Journal of Cardiovascular Magnetic Resonance* 14.1, p. 15 (cit. on p. 158).

De Zwaan, C., F. W. H. M. Bär, and H. J. J. Wellens (1982). "Characteristic electrocardiographic pattern indicating a critical stenosis high in left anterior descending coronary artery in patients admitted because of impending myocardial infarction". In: *American Heart Journal*. Proceedings of a symposium on

New strategies in the management of ischemic heart disease 103.4, pp. 730–736 (cit. on p. 127).

Di Diego, J. M., Z. Q. Sun, and C. Antzelevitch (1996). "I(to) and action potential notch are smaller in left vs. right canine ventricular epicardium". In: *The American journal of physiology* 271.2 Pt 2, H548–561 (cit. on p. 16).

DiFrancesco, D. and D. Noble (1985). "A Model of Cardiac Electrical Activity Incorporating Ionic Pumps and Concentration Changes". In: *Philosophical Transactions of the Royal Society of London. B, Biological Sciences* 307.1133, pp. 353–398 (cit. on p. 32).

Drouin, E. et al. (1995). "Electrophysiologic characteristics of cells spanning the left ventricular wall of human heart: Evidence for presence of M cells". In: *Journal of the American College of Cardiology* 26.1, pp. 185–192 (cit. on p. 14).

Durrer, D. et al. (1970). "Total Excitation of the Isolated Human Heart". In: *Circulation* 41.6, pp. 899–912 (cit. on pp. 83, 85, 87, 96).

Einthoven, W., G. Fahr, and A. de Waart (1950). "On the direction and manifest size of the variations of potential in the human heart and on the influence of the position of the heart on the form of the electrocardiogram". In: *American Heart Journal* 40.2, pp. 163–211 (cit. on p. 17).

Elizari, M. V., R. S. Acunzo, and M. Ferreiro (2007). "Hemiblocks Revisited". In: *Circulation* 115.9, pp. 1154–1163 (cit. on pp. 83, 95, 97, 99, 102, 105).

Finegold, J. A., P. Asaria, and D. P. Francis (2013). "Mortality from ischaemic heart disease by country, region, and age: Statistics from World Health Organisation and United Nations". In: *International Journal of Cardiology* 168.2, pp. 934–945 (cit. on p. 157).

FitzHugh, R. (1961). "Impulses and Physiological States in Theoretical Models of Nerve Membrane". In: *Biophysical Journal* 1.6, pp. 445–466 (cit. on p. 30).

Forder, J. R., M.-S. Hwang, and S. J. Blackband (2015). "Myocardial conduction network visualized by magnetic resonance microscopy/diffusion imaging and validated by histology". In: *Journal of Cardiovascular Magnetic Resonance* 17.Suppl 1, O73 (cit. on p. 3).

Foster, K. R. and H. P. Schwan (1989). "Dielectric properties of tissues and biological materials: a critical review". In: *Critical reviews in biomedical engineering* 17.1, pp. 25–104 (cit. on p. 40).

Funk, M. et al. (2013). "An Alarming Rate of Unnecessary Monitoring in the Practical Use of the Latest Standards of Electrocardiography (PULSE) Trial". In: *Journal of the American College of Cardiology* 61.10, Supplement, E1496 (cit. on p. 1).

Furukawa, T. et al. (1990). "Differences in transient outward currents of feline endocardial and epicardial myocytes". In: *Circulation research* 67.5, pp. 1287–1291 (cit. on p. 15).

Gabriel, C., S. Gabriel, and E. Corthout (1996). "The dielectric properties of biological tissues: I. Literature survey". In: *Physics in Medicine and Biology* 41.11, p. 2231 (cit. on p. 40).

Gabriel, S., R. W. Lau, and C. Gabriel (1996). "The dielectric properties of biological tissues: II. Measurements in the frequency range 10 Hz to 20 GHz". In: *Physics in Medicine and Biology* 41.11, p. 2251 (cit. on p. 40).

GBD Collaborators (2015). "Global, regional, and national age–sex specific all-cause and cause-specific mortality for 240 causes of death, 1990–2013: a systematic analysis for the Global Burden of Disease Study 2013". In: *The Lancet* 385.9963, pp. 117–171 (cit. on pp. 1, 157).

Ghosh, S. et al. (2010). "Early repolarization associated with sudden death: Insights from noninvasive electrocardiographic imaging". In: *Heart Rhythm* 7.4, pp. 534–537 (cit. on pp. 131–134).

Glukhov, A. V. et al. (2010). “Transmural Dispersion of Repolarization in Failing and Nonfailing Human Ventricle”. In: *Circulation Research* 106.5, pp. 981–991 (cit. on pp. 14, 108, 109).

Goldberger, A. L. et al. (2000). “PhysioBank, PhysioToolkit, and PhysioNet”. In: *Circulation* 101.23, e215–e220 (cit. on pp. 83, 116).

Goldberger, E. (1942). “The aVl, aVr, and aVf leads: A simplification of standard lead electrocardiography”. In: *American Heart Journal* 24.3, pp. 378–396 (cit. on p. 19).

Grandi, E., F. S. Pasqualini, and D. M. Bers (2010). “A novel computational model of the human ventricular action potential and Ca transient”. In: *Journal of Molecular and Cellular Cardiology* 48.1, pp. 112–121 (cit. on pp. 32, 33).

Greenbaum, R. A. et al. (1981). “Left ventricular fibre architecture in man.” In: *British heart journal* 45.3, pp. 248–263 (cit. on p. 13).

Haïssaguerre, M. et al. (2008). “Sudden Cardiac Arrest Associated with Early Repolarization”. In: *New England Journal of Medicine* 358.19, pp. 2016–2023 (cit. on p. 130).

Haïssaguerre, M. et al. (2009). “Ventricular Fibrillation with Prominent Early Repolarization Associated with a Rare Variant of KCNJ8/KATP Channel”. In: *Journal of Cardiovascular Electrophysiology* 20.1, pp. 93–98 (cit. on p. 131).

Hanson, B. et al. (2009). “Interaction of Activation–Repolarization Coupling and Restitution Properties in Humans”. In: *Circulation: Arrhythmia and Electrophysiology* 2.2, pp. 162–170 (cit. on pp. 27, 128).

Haruta, D. et al. (2011). “Incidence and Prognostic Value of Early Repolarization Pattern in the 12-Lead ElectrocardiogramClinical Perspective”. In: *Circulation* 123.25, pp. 2931–2937 (cit. on p. 130).

Hayden, G. E. et al. (2002). "Electrocardiographic T-wave inversion: Differential diagnosis in the chest pain patient". In: *The American Journal of Emergency Medicine* 20.3, pp. 252–262 (cit. on p. 27).

Hindmarsh, A. C. et al. (2005). "SUNDIALS: Suite of nonlinear and differential/algebraic equation solvers". In: *ACM Transactions on Mathematical Software (TOMS)* 31.3, pp. 363–396 (cit. on pp. 44, 73).

Hodgkin, A. L. and A. F. Huxley (1952). "A quantitative description of membrane current and its application to conduction and excitation in nerve". In: *The Journal of Physiology* 117.4, pp. 500–544 (cit. on pp. 29, 31).

Hoekema, R., G. J. H. Uijen, and A. v. Oosterom (2001). "Geometrical aspects of the interindividual variability of multilead ECG recordings". In: *IEEE Transactions on Biomedical Engineering* 48.5, pp. 551–559 (cit. on pp. 105, 106).

Hooks, D. A. et al. (2007). "Laminar Arrangement of Ventricular Myocytes Influences Electrical Behavior of the Heart". In: *Circulation Research* 101.10, e103–e112 (cit. on p. 38).

Iyer, V., R. Mazhari, and R. L. Winslow (2004). "A Computational Model of the Human Left-Ventricular Epicardial Myocyte". In: *Biophysical Journal* 87.3, pp. 1507–1525 (cit. on p. 32).

Janse, M. J. et al. (2005). "Repolarization Gradients in the Canine Left Ventricle Before and After Induction of Short-Term Cardiac Memory". In: *Circulation* 112.12, pp. 1711–1718 (cit. on pp. 15, 16).

Janse, M. J. et al. (2012). "Repolarization gradients in the intact heart: Transmural or apico-basal?" In: *Progress in Biophysics and Molecular Biology* 109.1-2, pp. 6–15 (cit. on pp. 14, 27, 108).

Jouk, P.-S. et al. (2007). "Analysis of the fiber architecture of the heart by quantitative polarized light microscopy. Accuracy, limitations and contribution to the

study of the fiber architecture of the ventricles during fetal and neonatal life". In: *European Journal of Cardio-Thoracic Surgery* 31.5, pp. 915–921 (cit. on p. 37).

Junttila, M. J. et al. (2014). "Early repolarization as a predictor of arrhythmic and nonarrhythmic cardiac events in middle-aged subjects". In: *Heart Rhythm*. Focus Issue: Sudden Death 11.10, pp. 1701–1706 (cit. on p. 130).

Katz, A. M. (2006). *Physiology of the Heart*. Philadelphia: Lippincott Williams & Wilkins (cit. on pp. 23–25, 27).

Keller, D. U. J., R. Kalayciyan, et al. (2010). "Fast Creation of Endocardial Stimulation Profiles for the Realistic Simulation of Body Surface ECGs". In: *World Congress on Medical Physics and Biomedical Engineering, September 7 - 12, 2009, Munich, Germany*. Ed. by O. Dössel and W. C. Schlegel. IFMBE Proceedings. Springer Berlin Heidelberg, pp. 145–148 (cit. on p. 85).

Keller, D. U. J. et al. (2012). "Influence of IKs Heterogeneities on the Genesis of the T-wave: A Computational Evaluation". In: *IEEE Transactions on Biomedical Engineering* 59.2, pp. 311–322 (cit. on pp. 108, 110).

Keller, D. U. J., F. M. Weber, et al. (2010). "Ranking the Influence of Tissue Conductivities on Forward-Calculated ECGs". In: *IEEE Transactions on Biomedical Engineering* 57.7, pp. 1568–1576 (cit. on pp. 91, 102).

Keller, D. U. J. et al. (2011). "Impact of Physiological Ventricular Deformation on the Morphology of the T-Wave: A Hybrid, Static-Dynamic Approach". In: *IEEE Transactions on Biomedical Engineering* 58.7, pp. 2109–2119 (cit. on pp. 104, 128).

Konarzewska, H., G. A. Peeters, and M. C. Sanguinetti (1995). "Repolarizing K+ Currents in Nonfailing Human Hearts". In: *Circulation* 92.5, pp. 1179–1187 (cit. on p. 111).

Koncz, I. et al. (2014). "Mechanisms underlying the development of the electrocardiographic and arrhythmic manifestations of early repolarization syndrome". In: *Journal of Molecular and Cellular Cardiology* 68, pp. 20–28 (cit. on p. 131).

Krause, H., H. Antoni, and A. Fleckenstein (1966). "An electronic model for the formation of local and transmitted stimuli on the myocardium fibers based upon variable current-voltage characteristics for potassium and sodium ions". ger. In: *Pflügers Archiv Für Die Gesamte Physiologie Des Menschen Und Der Tiere* 289.1, pp. 12–36 (cit. on p. 2).

LeGrice, I. J. et al. (1995). "Laminar structure of the heart: ventricular myocyte arrangement and connective tissue architecture in the dog". In: *American Journal of Physiology - Heart and Circulatory Physiology* 269.2, H571–H582 (cit. on p. 13).

Li, G.-R. et al. (1998). "Transmural heterogeneity of action potentials andI to1 in myocytes isolated from the human right ventricle". In: *American Journal of Physiology - Heart and Circulatory Physiology* 275.2, H369–H377 (cit. on pp. 15, 16).

Li, A.-H. et al. (2014). "Dynamic Changes in Myocardial Matrix and Relevance to Disease". In: *Circulation Research* 114.5, pp. 916–927 (cit. on p. 38).

Lines, G. T. et al. (2003). "Mathematical models and numerical methods for the forward problem in cardiac electrophysiology". In: *Computing and Visualization in Science* 5.4, pp. 215–239 (cit. on p. 37).

Litovsky, S. H. and C. Antzelevitch (1988). "Transient outward current prominent in canine ventricular epicardium but not endocardium." In: *Circulation Research* 62.1, pp. 116–126 (cit. on p. 15).

Liu, D.-W. and C. Antzelevitch (1995). "Characteristics of the Delayed Rectifier Current (IKr and IKs) in Canine Ventricular Epicardial, Midmyocardial, and Endocardial Myocytes A Weaker IKs Contributes to the Longer Action Potential of the M Cell". In: *Circulation Research* 76.3, pp. 351–365 (cit. on pp. 14, 15).

Luo, C. H. and Y. Rudy (1991). "A model of the ventricular cardiac action potential. Depolarization, repolarization, and their interaction." In: *Circulation Research* 68.6, pp. 1501–1526 (cit. on p. 32).

Macfarlane, P. W. et al. (2015). "The Early Repolarization Pattern: A Consensus Paper". In: *Journal of the American College of Cardiology* 66.4, pp. 470–477 (cit. on pp. 130, 145, 148).

Mattu, A. (2003). *ECGs for the emergency physician*. London: BMJ Books (cit. on pp. 27, 126).

Maury, P. and A. Rollin (2013). "Prevalence of early repolarisation/J wave patterns in the normal population". In: *Journal of Electrocardiology* 46.5, pp. 411–416 (cit. on p. 129).

Medeiros-Domingo, A. et al. (2010). "Gain-of-function mutation S422L in the KCNJ8-encoded cardiac KATP channel Kir6.1 as a pathogenic substrate for J-wave syndromes". In: *Heart Rhythm* 7.10, pp. 1466–1471 (cit. on p. 131).

Meijborg, V. M. F. et al. (2015). "Interventricular dispersion in repolarization causes bifid T waves in dogs with dofetilide-induced long QT syndrome". In: *Heart Rhythm* 12.6, pp. 1343–1351 (cit. on p. 126).

Meijborg, V. M. F. et al. (2016). "Reduced Sodium Current in the Lateral Ventricular Wall Induces Inferolateral J-Waves". In: *Cardiac Electrophysiology* 7, p. 365 (cit. on pp. 131, 147, 155).

Miller, W. T. and D. B. Geselowitz (1978). "Simulation studies of the electrocardiogram. II. Ischemia and infarction." In: *Circulation Research* 43.2, pp. 315–323 (cit. on p. 34).

Mincholé, A. et al. (2015). "ECG-based estimation of dispersion of APD restitution as a tool to stratify sotalol-induced arrhythmic risk". In: *Journal of Electrocardiology* 48.5, pp. 867–873 (cit. on p. 82).

Mirams, G. R. et al. (2013). "Chaste: An Open Source C++ Library for Computational Physiology and Biology". In: *PLoS Computational Biology* 9.3 (cit. on p. 44).

Muszkiewicz, A. et al. (2016). "Variability in cardiac electrophysiology: Using experimentally-calibrated populations of models to move beyond the single virtual physiolog-

ical human paradigm". In: *Progress in Biophysics and Molecular Biology*. Recent Developments in Biophysics & Molecular Biology of Heart Rhythm 120.1–3, pp. 115–127 (cit. on pp. 91, 104).

Näbauer, M. et al. (1996). "Regional Differences in Current Density and Rate-Dependent Properties of the Transient Outward Current in Subepicardial and Subendocardial Myocytes of Human Left Ventricle". In: *Circulation* 93.1, pp. 168–177 (cit. on pp. 15, 109).

Nagra, B. S., G. S. Ledley, and B. K. Kantharia (2005). "Marked QT Prolongation and Torsades de Pointes Secondary to Acute Ischemia in an Elderly Man Taking Dofetilide for Atrial Fibrillation: A Cautionary Tale". In: *Journal of Cardiovascular Pharmacology and Therapeutics* 10.3, pp. 191–195 (cit. on p. 158).

Nagumo, J., S. Arimoto, and S. Yoshizawa (1962). "An Active Pulse Transmission Line Simulating Nerve Axon". In: *Proceedings of the IRE* 50.10, pp. 2061–2070 (cit. on p. 30).

Nakagawa, K. et al. (2014). "Left ventricular epicardial electrogram recordings in idiopathic ventricular fibrillation with inferior and lateral early repolarization". In: *Heart Rhythm* 11.2, pp. 314–317 (cit. on p. 131).

Nguyên, U. C. et al. (2015). "An in-silico analysis of the effect of heart position and orientation on the ECG morphology and vectorcardiogram parameters in patients with heart failure and intraventricular conduction defects". In: *Journal of Electrocardiology* 48.4, pp. 617–625 (cit. on pp. 108, 128).

Niederer, S. A. et al. (2011). "Verification of cardiac tissue electrophysiology simulators using an N-version benchmark". In: *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences* 369.1954, pp. 4331–4351 (cit. on p. 41).

Noble, D. (1962). “A modification of the Hodgkin—Huxley equations applicable to Purkinje fibre action and pacemaker potentials”. In: *The Journal of Physiology* 160.2, pp. 317–352 (cit. on pp. 2, 31).

Noble, D., A. Garny, and P. J. Noble (2012). “How the Hodgkin–Huxley equations inspired the Cardiac Physiome Project”. In: *The Journal of Physiology* 590.11, pp. 2613–2628 (cit. on pp. 2, 30).

Noro, M. et al. (2016). “Efficacy and Myocardial Injury With Subcutaneous Implantable Cardioverter Defibrillators: Computer Simulation of Defibrillation Shock Conduction”. In: *Circulation Journal: Official Journal of the Japanese Circulation Society* 80.1, pp. 85–92 (cit. on p. 158).

Obreztchikova, M. N. et al. (2006). “IKr contributes to the altered ventricular repolarization that determines long-term cardiac memory”. In: *Cardiovascular Research* 71.1, pp. 88–96 (cit. on p. 15).

O’Hara, T. et al. (2011). “Simulation of the undiseased human cardiac ventricular action potential: model formulation and experimental validation”. In: *PLoS Computational Biology* 7.5, e1002061 (cit. on pp. 15, 33, 157).

Okada, J.-I. et al. (2011). “Transmural and apicobasal gradients in repolarization contribute to T-wave genesis in human surface ECG”. In: *American Journal of Physiology - Heart and Circulatory Physiology* 301.1, H200–H208 (cit. on pp. 108, 127, 155).

Okada, J. et al. (2013). “Patient Specific Simulation of Body Surface ECG using the Finite Element Method: HUMAN ECG SIMULATION”. In: *Pacing and Clinical Electrophysiology* 36.3, pp. 309–321 (cit. on pp. 33, 104, 105, 154).

Olson, K. A. et al. (2011). “Long-term prognosis associated with J-point elevation in a large middle-aged biracial cohort: the ARIC study”. In: *European Heart Journal* 32.24, pp. 3098–3106 (cit. on p. 130).

Opthof, T., R. Coronel, and M. J. Janse (2009). "Is there a significant transmural gradient in repolarization time in the intact heart? Repolarization Gradients in the Intact Heart". In: *Circulation: Arrhythmia and Electrophysiology* 2.1, pp. 89–96 (cit. on pp. 108, 110).

Opthof, T. et al. (2007). "Dispersion of repolarization in canine ventricle and the electrocardiographic T wave: Tp-e interval does not reflect transmural dispersion". In: *Heart Rhythm* 4.3, pp. 341–348 (cit. on p. 108).

Opthof, T. et al. (2016). "Dispersion in ventricular repolarization in the human, canine and porcine heart". In: *Progress in Biophysics and Molecular Biology*. Recent Developments in Biophysics & Molecular Biology of Heart Rhythm 120.1–3, pp. 222–235 (cit. on pp. 14, 108).

Orini, M. et al. (2013). "Does the inverse activation time/action potential duration relationship apply to the in vivo human heart?" In: *Journal of Electrocardiology* 46.4, e35 (cit. on p. 128).

Patel, C. et al. (2009). "Is there a significant transmural gradient in repolarization time in the intact heart? Cellular Basis of the T Wave: A Century of Controversy". In: *Circulation: Arrhythmia and Electrophysiology* 2.1, pp. 80–88 (cit. on p. 108).

Pathmanathan, P. et al. (2010). "A numerical guide to the solution of the bidomain equations of cardiac electrophysiology". In: *Progress in Biophysics and Molecular Biology* 102.2–3, pp. 136–155 (cit. on pp. 35, 36, 41, 74).

Patton, K. K. et al. (2016). "Electrocardiographic Early Repolarization". In: *Circulation* 133.15, pp. 1520–1529 (cit. on pp. 130, 131, 145).

Péréon, Y. et al. (2000). "Differential expression of KvLQT1 isoforms across the human ventricular wall". In: *American Journal of Physiology - Heart and Circulatory Physiology* 278.6, H1908–H1915 (cit. on p. 108).

Pitt-Francis, J. et al. (2009). "Chaste: A test-driven approach to software development for biological modelling". In: *Computer Physics Communications* 180.12, pp. 2452–2471 (cit. on p. 44).

Plonsey, R. and J. Malmivuo (1995). *Bioelectromagnetism: Principles and Applications of Bioelectric and Biomagnetic Fields*. New York: Oxford University Press (cit. on pp. 19, 20, 93).

Potse, M. et al. (2006). "A comparison of monodomain and bidomain reaction-diffusion models for action potential propagation in the human heart". In: *IEEE Transactions on Biomedical Engineering* 53.12, pp. 2425–2435 (cit. on p. 2).

Potse, M. et al. (2014). "Patient-specific modelling of cardiac electrophysiology in heart-failure patients". In: *Europace* 16.suppl 4, pp. iv56–iv61 (cit. on pp. 104, 105, 154).

Priori, S. G. et al. (2013). "HRS/EHRA/APHRS Expert Consensus Statement on the Diagnosis and Management of Patients with Inherited Primary Arrhythmia Syndromes: Document endorsed by HRS, EHRA, and APHRS in May 2013 and by ACCF, AHA, PACES, and AEPC in June 2013." In: *Heart Rhythm*. Focus Issue: Devices 10.12, pp. 1932–1963 (cit. on pp. 130, 148).

Provost, J. et al. (2013). "A clinical feasibility study of atrial and ventricular electromechanical wave imaging". In: *Heart Rhythm* 10.6, pp. 856–862 (cit. on pp. 83, 85, 89, 96).

Ramanathan, C. et al. (2006). "Activation and repolarization of the normal human heart under complete physiological conditions". In: *Proceedings of the National Academy of Sciences* 103.16, pp. 6309–6314 (cit. on pp. 16, 83, 85, 89, 109).

Roberts, D. E., L. T. Hersh, and A. M. Scher (1979). "Influence of cardiac fiber orientation on wavefront voltage, conduction velocity, and tissue resistivity in the dog." In: *Circulation Research* 44.5, pp. 701–712 (cit. on p. 37).

Roberts, D. E. and A. M. Scher (1982). "Effect of tissue anisotropy on extracellular potential fields in canine myocardium in situ." In: *Circulation Research* 50.3, pp. 342–351 (cit. on p. 37).

Rodriguez, B. et al. (2015). "Human-based approaches to pharmacology and cardiology: an interdisciplinary and intersectorial workshop". In: *Europace*, euv320 (cit. on pp. 6, 30).

Rollin, A. et al. (2012). "Prevalence, Prognosis, and Identification of the Malignant Form of Early Repolarization Pattern in a Population-Based Study". In: *The American Journal of Cardiology* 110.9, pp. 1302–1308 (cit. on p. 130).

Rosenbaum, M. B. et al. (1969). "Intraventricular trifascicular blocks. Review of the literature and classification". In: *American Heart Journal* 78.4, pp. 450–459 (cit. on p. 85).

Rosso, R. et al. (2008). "J-Point Elevation in Survivors of Primary Ventricular Fibrillation and Matched Control Subjects: Incidence and Clinical Significance". In: *Journal of the American College of Cardiology* 52.15, pp. 1231–1238 (cit. on p. 130).

Rush, S., J. A. Abildskov, and R. Mcfee (1963). "Resistivity of Body Tissues at Low Frequencies". In: *Circulation Research* 12.1, pp. 40–50 (cit. on p. 40).

Sahli Costabal, F., D. E. Hurtado, and E. Kuhl (2015). "Generating Purkinje networks in the human heart". In: *Journal of Biomechanics* (cit. on pp. 104, 105).

Samie, F. H. et al. (2001). "Rectification of the Background Potassium Current A Determinant of Rotor Dynamics in Ventricular Fibrillation". In: *Circulation Research* 89.12, pp. 1216–1223 (cit. on p. 16).

Sands, G. et al. (2006). "Automated Extended Volume Imaging of Tissue using Confocal and Optical Microscopy". In: *28th Annual International Conference of the IEEE Engineering in Medicine and Biology Society, 2006. EMBS '06*, pp. 133–136 (cit. on p. 13).

Sebastian, R. et al. (2008). "Assessing influence of conductivity in heart modelling with the aim of studying cardiovascular diseases". In: *Proceedings of SPIE*. Vol. 6916, pp. 691627/1–691627/10 (cit. on pp. 39, 41).

Seemann, G. et al. (2013). "Variation of Human Ventricular Iks Heterogeneities to Reconstruct Measured Multi-Channel ECG Data". In: *Biomedical Engineering / Biomedizinische Technik* (cit. on pp. 127, 155).

Sekiya, S. et al. (1984). "Distribution of action potential durations in the canine left ventricle". In: *Japanese heart journal* 25.2, pp. 181–194 (cit. on pp. 15, 16).

Shannon, T. R. et al. (2004). "A Mathematical Treatment of Integrated Ca Dynamics within the Ventricular Myocyte". In: *Biophysical Journal* 87.5, pp. 3351–3371 (cit. on p. 32).

Shipley, R. A. and W. R. Hallaran (1936). "The four-lead electrocardiogram in two hundred normal men and women". In: *American Heart Journal* 11.3, pp. 325–345 (cit. on p. 129).

Sicouri, S. and C. Antzelevitch (1991). "A subpopulation of cells with unique electrophysiological properties in the deep subepicardium of the canine ventricle. The M cell." In: *Circulation Research* 68.6, pp. 1729–1741 (cit. on p. 14).

Silverman, M. E., D. Grove, and C. B. Upshaw (2006). "Why Does the Heart Beat? The Discovery of the Electrical System of the Heart". In: *Circulation* 113.23, pp. 2775–2781 (cit. on p. 3).

Sinner, M. F. et al. (2010). "Association of Early Repolarization Pattern on ECG with Risk of Cardiac and All-Cause Mortality: A Population-Based Prospective Cohort Study (MONICA/KORA)". In: *PLOS Med* 7.7, e1000314 (cit. on p. 130).

Smith, N. P., M. L. Buist, and A. J. Pullan (2003). "Altered T Wave Dynamics in a Contracting Cardiac Model". In: *Journal of Cardiovascular Electrophysiology* 14, S203–S209 (cit. on p. 128).

Society for Cardiological Science and Technology (2010). *Recording a standard 12-lead electrocardiogram An Approved Methodology* (cit. on pp. 17, 42, 43).

Southern, J. A. et al. (2009). "Solving the Coupled System Improves Computational Efficiency of the Bidomain Equations". In: *IEEE Transactions on Biomedical Engineering* 56.10, pp. 2404–2412 (cit. on p. 37).

Streeter, D. D. et al. (1969). "Fiber Orientation in the Canine Left Ventricle during Diastole and Systole". In: *Circulation Research* 24.3, pp. 339–347 (cit. on pp. 37, 44).

Sundnes, J. et al. (2009). "A Second-Order Algorithm for Solving Dynamic Cell Membrane Equations". In: *IEEE Transactions on Biomedical Engineering* 56.10, pp. 2546–2548 (cit. on p. 44).

Szabó, G. et al. (2005). "Asymmetrical distribution of ion channels in canine and human left-ventricular wall: epicardium versus midmyocardium". In: *Pflügers Archiv - European Journal of Physiology* 450.5, pp. 307–316 (cit. on pp. 15, 108).

Szentadrassy, N. et al. (2005). "Apico–basal inhomogeneity in distribution of ion channels in canine and human ventricular myocardium". In: *Cardiovascular Research* 65.4, pp. 851–860 (cit. on pp. 16, 108).

Taccardi, B. (1963). "Distribution of Heart Potentials on the Thoracic Surface of Normal Human Subjects". In: *Circulation Research* 12.4, pp. 341–352 (cit. on pp. 83, 85, 90).

Taggart, P. et al. (2001). "Transmural repolarisation in the left ventricle in humans during normoxia and ischaemia". In: *Cardiovascular Research* 50.3, pp. 454–462 (cit. on pp. 15, 27, 110).

Taggart, P. et al. (2003). "Electrotonic cancellation of transmural electrical gradients in the left ventricle in man". In: *Progress in Biophysics and Molecular Biology*. Mechano-Electric Feedback and Cardiac Arrhythmias 82.1–3, pp. 243–254 (cit. on p. 108).

Ten Tusscher, K. H. W. J. and A. V. Panfilov (2006). "Alternans and spiral breakup in a human ventricular tissue model". In: *American Journal of Physiology - Heart and Circulatory Physiology* 291.3, H1088–H1100 (cit. on pp. 32–34, 39, 109, 132, 157).

Ten Tusscher, K. H. W. J. et al. (2004). "A model for human ventricular tissue". In: *American Journal of Physiology - Heart and Circulatory Physiology* 286.4, H1573–H1589 (cit. on pp. 32, 33).

Tikkanen, J. T. et al. (2009). "Long-Term Outcome Associated with Early Repolarization on Electrocardiography". In: *New England Journal of Medicine* 361.26, pp. 2529–2537 (cit. on p. 130).

Tikkanen, J. T. et al. (2011). "Early Repolarization: electrocardiographic phenotypes associated with favorable long-term outcome". In: *Circulation* 123.23, pp. 2666–2673 (cit. on pp. 130, 131, 149).

Trayanova, N. A. (2011). "Whole-Heart Modeling Applications to Cardiac Electrophysiology and Electromechanics". In: *Circulation Research* 108.1, pp. 113–128 (cit. on p. 2).

Tung, L. (1978). "A bi-domain model for describing ischemic myocardial d-c potentials." Thesis. Dept. of Electrical Engineering and Computer Science, Massachusetts Institute of Technology (cit. on p. 34).

Vigmond, E. J. and B. D. Stuyvers (2016). "Modeling our understanding of the His-Purkinje system". In: *Progress in Biophysics and Molecular Biology*. Recent Developments in Biophysics & Molecular Biology of Heart Rhythm 120.1–3, pp. 179–188 (cit. on p. 106).

Villongco, C. T. et al. (2014). "Patient-specific modeling of ventricular activation pattern using surface ECG-derived vectorcardiogram in bundle branch block". In: *Progress in Biophysics and Molecular Biology* 115.2-3, pp. 305–313 (cit. on pp. 33, 154).

Viswanathan, P. C., R. M. Shaw, and Y. Rudy (1999). "Effects of IKr and IKs Heterogeneity on Action Potential Duration and Its Rate Dependence". In: *Circulation* 99.18, pp. 2466–2474 (cit. on p. 110).

Volders, P. G. A. et al. (1999). "Repolarizing K+ Currents ITO1 and IKs Are Larger in Right Than Left Canine Ventricular Midmyocardium". In: *Circulation* 99.2, pp. 206–210 (cit. on p. 16).

Vos, M. A. et al. (1998). "Enhanced Susceptibility for Acquired Torsade de Pointes Arrhythmias in the Dog With Chronic, Complete AV Block Is Related to Cardiac Hypertrophy and Electrical Remodeling". In: *Circulation* 98.11, pp. 1125–1135 (cit. on p. 16).

Washio, T., J. Okada, and T. Hisada (2010). "A Parallel Multilevel Technique for Solving the Bidomain Equation on a Human Heart with Purkinje Fibers and a Torso Model". In: *SIAM Review* 52.4, pp. 717–743 (cit. on p. 2).

Watanabe, H. et al. (2011). "Electrocardiographic Characteristics and SCN5A Mutations in Idiopathic Ventricular Fibrillation Associated With Early Repolarization". In: *Circulation: Arrhythmia and Electrophysiology* 4.6, pp. 874–881 (cit. on p. 131).

Watanabe, T., P. M. Rautaharju, and T. F. McDonald (1985). "Ventricular action potentials, ventricular extracellular potentials, and the ECG of guinea pig." In: *Circulation Research* 57.3, pp. 362–373 (cit. on pp. 15, 16).

Watanabe, T. et al. (1983). "Heterogeneity of the action potential in isolated rat ventricular myocytes and tissue." In: *Circulation Research* 52.3, pp. 280–290 (cit. on p. 16).

Wellens, H. J. (2008). "Early Repolarization Revisited". In: *New England Journal of Medicine* 358.19, pp. 2063–2065 (cit. on p. 131).

Wettwer, E. et al. (1994). "Transient outward current in human ventricular myocytes of subepicardial and subendocardial origin." In: *Circulation Research* 75.3, pp. 473–482 (cit. on pp. 15, 16).

Wilhelms, M., O. Dossel, and G. Seemann (2011). "In Silico Investigation of Electrically Silent Acute Cardiac Ischemia in the Human Ventricles". In: *IEEE Transactions on Biomedical Engineering* 58.10, pp. 2961–2964 (cit. on p. 82).

Wilson, F. N. et al. (1934). "Electrocardiograms that represent the potential variations of a single electrode". In: *American Heart Journal* 9.4, pp. 447–458 (cit. on p. 19).

Wilson, F. N. et al. (1944). "The precordial electrocardiogram". In: *American Heart Journal* 27.1, pp. 19–85 (cit. on p. 20).

Yan, G.-X. and C. Antzelevitch (1996). "Cellular Basis for the Electrocardiographic J Wave". In: *Circulation* 93.2, pp. 372–379 (cit. on p. 131).

Young et al. (1998). "Extended confocal microscopy of myocardial laminae and collagen network". In: *Journal of Microscopy* 192.2, pp. 139–150 (cit. on p. 13).

Zemzemi, N. and B. Rodriguez (2015). "Effects of L-type calcium channel and human ether-a-go-go related gene blockers on the electrical activity of the human heart: a simulation study". In: *Europace* 17.2, pp. 326–333 (cit. on p. 82).

Zemzemi, N. et al. (2012). "Computational assessment of drug-induced effects on the electrocardiogram: from ion channel to body surface potentials". In: *British Journal of Pharmacology* (cit. on pp. 37, 158).

Zheng, Y. et al. (2016). "Transmural, interventricular, apicobasal and anteroposterior action potential duration gradients are all essential to the genesis of the concordant and realistic T wave: A whole-heart model study". In: *Journal of Electrocardiology* 49.4, pp. 569–578 (cit. on pp. 108, 114, 127).

Zicha, S. et al. (2004). “Transmural expression of transient outward potassium current subunits in normal and failing canine and human hearts”. In: *The Journal of Physiology* 561.3, pp. 735–748 (cit. on p. 108).

Zygmunt, A. C., R. J. Goodrow, and C. Antzelevitch (2000). “I(NaCa) contributes to electrical heterogeneity within the canine ventricle”. In: *American Journal of Physiology - Heart and Circulatory Physiology* 278.5, H1671–H1678 (cit. on pp. 15, 108).

Zygmunt, A. C. et al. (2001). “Larger late sodium conductance in M cells contributes to electrical heterogeneity in canine ventricle”. In: *American Journal of Physiology - Heart and Circulatory Physiology* 281.2, H689–H697 (cit. on pp. 15, 108).