Bayesian inference and Machine learning methods for the advanced analysis of electrocardiographic signals

Julià Camps

St Edmund Hall

University of Oxford

A thesis submitted for the degree of

Doctor of Philosophy

Trinity 2021

To my wife Núria
and
my parents Arlene and Julià

Acknowledgements

Personal

Prof Blanca Rodríguez, you moulded my perception of reality to guide me from engineer to scientist. My only wish is that it had been a longer journey since I still have too many things to learn from you.

Prof Kevin Burrage and Dr Brodie Lawson, that month working together at Queensland University of Technology was the spark that started this research.

Thank you, Núria Sánchez, for supporting me in this incredible adventure; you deserve half of the credit.

Dr 王芝诺 (Zhinuo Jenny Wang), thank you for our invaluable scientific discussions.

Institutional

I would like to thank my supervisors, Prof Blanca Rodríguez and Dr Ana Mincholé, to the University, especially to the Computational Cardiovascular Science (CCS) group, the Department of Computer Science, and my college, Teddy Hall (St Edmund Hall), and to my funders, Scatcherd Scholarships and the EPSRC. Thank you all for the unique experience offered in this town.

I enjoyed every day as no non-Oxford PhD student can imagine.

Abstract

The realisation of precision cardiology requires novel techniques for the non-invasive characterisation of individual patients' cardiac function to inform therapeutic and diagnostic decision-making. Both electrocardiography and imaging are used for the clinical diagnosis of cardiac disease. The integration of multimodal datasets through advanced computational methods could enable the development of the cardiac 'digital twin'. A cardiac digital twin is a comprehensive virtual tool that mechanistically reveals a patient's heart condition from clinical data and simulates treatment outcomes. The adoption of cardiac digital twins requires the non-invasive efficient personalisation of the electrophysiological properties in cardiac models.

This thesis presents novel computational techniques to exploit the synergy between Artificial Intelligence and multimodal datasets towards efficiently producing cardiac digital twins from non-invasive routinely acquired clinical data modalities, namely, ECG and cardiac magnetic resonance (CMR) imaging. I present and evaluate

- an efficient Bayesian-based inference method for estimating the ventricular activation properties from electrocardiographic data through CMR-based modelling and simulation,
- and a machine learning-based ECG delineation methodology capable of generalising to new databases to enable the translation of my inference method into real-world clinical applications.

These two methodologies, combined with CMR-based cardiac anatomical reconstruction, compose the foundations for an end-to-end cardiac digital twin generation pipeline from non-invasive standard clinical data modalities.

Research outputs

Publications

- **Camps, J.**, Lawson, B., Drovandi, C., Minchole, A., Wang, Z. J., Grau, V., Burrage, K. & Rodriguez, B. (2021). Inference of ventricular activation properties from non-invasive electrocardiography. Medical Image Analysis, 102143.
- Minchole, A., **Camps, J.**, Lyon, A., & Rodríguez, B. (2019). Machine learning in the electrocardiogram. Journal of electrocardiology, 57, S61-S64.
- Banerjee, A., **Camps, J.**, Zacur, E., Choudhury, R. P., Rodriguez, B. & Grau, V. (2021). A Completely Automated Pipeline for 3D Reconstruction of Human Heart from 2D Cine MR Slices. Accepted for publication in Philosophical Transactions of the Royal Society A.

Conference publications

- **Camps, J.**, Rodríguez, B., & Mincholé, A. (2018, September). Deep learning based QRS multilead delineator in electrocardiogram signals. In 2018 Computing in Cardiology Conference (CinC) (Vol. 45, pp. 1-4). IEEE.
- **Camps, J.**, McCarthy, A., Rodrıguez, B., & Mincholé, A. (2018). ConvNet based QRS Multilead Delineator in Electrocardiogram signals. International Conference on Machine Learning (ICML) workshop.

Awards

- EPSRC Doctoral Prize to increase the impact of the work described in this thesis. This prize included nine months of salary to stay as a postdoctoral researcher at the University of Oxford.
- The Rosanna Degani Young Investigator Award semi-finalist at the international conference Computing in Cardiology 2018.

Software

- Bayesian inference method for the estimation of human ventricular activation properties from electrocardiographic data. https://github.com/juliacamps/Inference-of-healthy-ventricular-activation-properties
- Multi-lead electrocardiogram delineation methodology to automatically extract QRS complexes from clinical electrocardiogram recordings, trained and evaluated on the QT database. https://github.com/juliacamps/ECG-QRSdelineation

Contents

Research outputs **ix**

List of Figures **xv**

List of Tables **xvii**

Glossary **xix**

1 Introduction **1**
1.1 Motivation 1
1.2 Goals 2
1.3 Outline 3

2 Background **5**
2.1 The heart 5
2.1.1 The healthy human cardiac cycle 6
2.2 Cardiac electrophysiology 9
2.2.1 Cellular cardiac electrophysiology 10
2.2.2 The cardiac electrical conduction system 12
2.2.3 Cardiac electrical propagation 13
2.3 The electrocardiogram 16
2.3.1 The 12-lead ECG test 18
2.4 ECG analysis 22
2.4.1 ECG noise-removal 23
2.4.2 ECG delineation 24
2.5 Non-invasive electrocardiographic imaging 27
2.6 Cardiac multiscale modelling and simulation 28
2.6.1 Bidomain model 28
2.6.2 Monodomain model 30
2.6.3 Eikonal model 30
2.7 Bayesian inference 31
2.7.1 Sequential Monte Carlo approximate Bayesian computation 33
2.8 Convolutional neural networks 34

2.8.1 Activation functions in deep neural networks 35
2.8.2 Regularisation for deep learning algorithms 36
2.8.3 Training deep learning algorithms 38
2.8.4 Optimisation algorithms for deep learning 40
2.8.5 Deep learning strategies for ECG delineation 41

3 Literature review **43**
3.1 Inference of the ventricular activation properties 43
3.2 ECG delineation 47
3.3 Machine learning for ECG 53
3.3.1 Deep learning for ECG 54

4 Methods **57**
4.1 Introduction 57
4.2 Overview 58
4.3 Virtual subjects for the 'target data' 61
4.4 Simulation protocols 62
4.4.1 Eikonal implementation 64
4.4.2 Electrocardiogram computation 66
4.5 Discrepancy calculation 69
4.5.1 Parameter space exploration 70
4.6 Parameter inference with SMC-ABC method 71
4.6.1 Error metrics and post-processing 74
4.6.2 Hyperparameter calibration 75

5 Inference of ventricular activation properties from activation time maps **77**
5.1 Introduction 77
5.2 Methods 78
5.3 Results 79
5.3.1 Reproduction of noise-contaminated epicardial activation maps 79
5.3.2 Root node inference from noise-contaminated epicardial activation maps 79
5.3.3 Speed inference from noise-contaminated epicardial activation maps 85
5.3.4 Inference from noise-free bidomain epicardial activation maps 87
5.3.5 Impact of the individual conduction speeds on the activation sequence 89
5.4 Discussion 89
5.5 Limitations and future work 92

6 Inference of ventricular activation properties from 12-lead ECG-QRS 95
6.1 Introduction . . . 95
6.2 Methods . . . 95
6.3 Results . . . 96
6.3.1 Reproduction of noise-contaminated ECGs . . . 96
6.3.2 Root node inference from noise-contaminated ECGs . . . 96
6.3.3 Speed inference from noise-contaminated ECGs . . . 102
6.3.4 Inference from noise-free bidomain ECGs . . . 104
6.4 Discussion . . . 105
6.4.1 Limitations and future work . . . 106

7 Deep learning QRS delineation in multi-lead clinical electrocardiograms 109
7.1 Introduction . . . 109
7.2 Methods . . . 110
7.2.1 Overview . . . 110
7.2.2 Clinical ECG recordings . . . 110
7.2.3 ECG pre-processing and data augmentation . . . 111
7.2.4 Deep learning architecture . . . 112
7.2.5 Training of the deep learning algorithms . . . 114
7.2.6 Delineation process and its evaluation . . . 117
7.2.7 Error metrics and post-processing . . . 117
7.3 Results and discussion . . . 118
7.3.1 Limitations and future work . . . 119

8 Conclusions and future directions 121
8.1 Contributions . . . 121
8.2 Future directions . . . 122
8.2.1 Motivation . . . 122
8.2.2 Proposed investigations . . . 123
8.3 Concluding remarks . . . 125

References 127

List of Figures

2.1 Schematic representation of the cardiovascular system. 6
2.2 Four-chamber illustration of the human heart. 7
2.3 Cardiac cycle in a healthy human heart. 8
2.4 Cardiomyocytes with gap junctions. 10
2.5 Human ventricular action potential. 11
2.6 Four-chamber illustration of the electrical system in the human heart. 12
2.7 Transmural properties of the cardiac wall. 14
2.8 Electrical activity in the myocardium. 15
2.9 Human ECG during sinus rhythm. 17
2.10 Placement of limb electrodes for the 12-lead ECG test. 19
2.11 Placement of precordial electrodes for the 12-lead ECG test. 20
2.12 Directionality of electrical activity represented in the limb leads. . . 21
2.13 Directionality in 3D of the 12-lead ECG. 22
2.14 Example of the variability in human ECG signals. 26
2.15 Typical convolutional neural network architecture. 35
2.16 Regularisation example. 37

4.1 Proposed inference method from 12-lead ECGs. 59
4.2 Comparison between pseudo-ECG and diffusion-based computed QRSs. 68

5.1 Proposed inference method from epicardial activation maps. 78
5.2 Agreement between recovered and target Eikonal activation maps. . 80
5.3 Root nodes inferred from noise-contaminated activation maps for Mesh-1. 81
5.4 Root nodes inferred from noise-contaminated activation maps for Mesh-2. 82
5.5 Root nodes inferred from noise-contaminated activation maps for Mesh-3. 83
5.6 Root nodes inferred from noise-contaminated activation maps for Mesh-4. 84
5.7 Conduction speeds inference from noise-contaminated activation maps. 86
5.8 Correlation heatmap between root node and conduction speed inference errors. 88

6.1 Agreement between recovered and target Eikonal 12-lead ECG-QRSs. 97
6.2 Root nodes inferred from noise-contaminated ECGs for Mesh-1. . . 98
6.3 Root nodes inferred from noise-contaminated ECGs for Mesh-2. . . 99
6.4 Root nodes inferred from noise-contaminated ECGs for Mesh-3. . . 100
6.5 Root nodes inferred from noise-contaminated ECGs for Mesh-4. . . 101
6.6 Conduction speeds inference from noise-contaminated ECGs. 103

7.1 Deep learning-based ECG multi-lead delineator. 110
7.2 Segmenter's deep learning architecture. 113

8.1 Extended automatic cardiac digital twin generation pipeline. 123

List of Tables

4.1 Conduction speed configurations considered for the virtual cohort. . 62

5.1 Root nodes errors from noise-contaminated Eikonal activation maps. 85
5.2 Prediction accuracy of the inference from bidomain activation maps. 87
5.3 Root node recovery errors from bidomain activation maps. 87
5.4 Conduction speed recovery errors from bidomain activation maps. . 88
5.5 Effect of conduction speed variations on the Eikonal activation map. 89

6.1 Root nodes errors from noise-contaminated Eikonal ECGs. 102
6.2 Prediction accuracy of the inference from bidomain ECGs. 104
6.3 Root node recovery errors from bidomain ECGs. 104
6.4 Conduction speed recovery errors from bidomain ECGs. 105

7.1 Comparison of delineation performance on the QT database. 119

Glossary

ECG . . . . . Electrocardiogram.

ECGI . . . . . Electrocardiographic imaging.

CMR . . . . . Cardiac magnetic resonance.

Cardiomyocyte . . . . Cardiac muscle cell.

Depolarisation . . . . . Activation of the cardiomyocytes leading to subsequent contractile function.

Repolarisation . . . . . Deactivation of the cardiomyocytes leading to subsequent relaxation of the tissue.

QRS . . . . . Complex composed by the Q, R, and S waves in the electrocardiogram signal representing the depolarisation of the ventricles.

Invasive . . . . . Requires penetrating the body.

Non-invasive . . . . . Does not require penetrating the body.

Sinus rhythm . . . . . Cardiac rhythm produced when the pacing is dictated by the sinoatrial node and no abnormal cardiac events are occurring.

Purkinje network . . . Tree-like network of tissue composing the electrical system in the ventricles.

Root nodes . . . . . Earliest endocardial sites of activation in the ventricles.

Anterior . . . . . Frontal-left (chest-left) side of the heart.

Posterior . . . . . Back-right side of the heart, namely, opposite to anterior.

Mesh . . . . . Graph-like structure (representation) with nodes and edges delimiting volumetric elements.

SMC-ABC . . . . . Sequential (Markov Chain) Monte Carlo - approximate Bayesian computation.

DTW . . . . . Dynamic time warping.

La vie n'est facile pour aucun de nous. Mais quoi, il faut avoir de la persévérance, et surtout de la confiance en soi. Il faut croire que l'on est doué pour quelque chose, et que, cette chose, il faut l'atteindre coûte que coûte.

Life is not easy for any of us. But what of that? We must have perseverance and above all confidence in ourselves. We must believe that we are gifted for something, and that this thing, at whatever cost, must be attained.

— Marie Curie (Curie, 1937)

1 Introduction

1.1 Motivation

Cardiovascular disease is the most common non-communicable disease globally, accounting for 17.8 million deaths in 2017, according to the World Health Organisation (Kaptoge et al., 2019). Cardiac disease increases the risk of sudden and premature death through alterations in cardiac electrophysiology and tissue structure, which are known to promote lethal arrhythmias and mechanical dysfunction.

Imaging and electrocardiographic datasets are commonly used for the diagnosis of cardiac disease. The electrocardiogram (ECG) is the most widely used clinical modality for cardiac disease diagnosis. Each beat in an ECG signal comprises a QRS complex and a T-wave, representing the ventricular cardiac cycle's activation and repolarisation phases. The clinical interpretation of ECG recordings relies on the delineation of these complexes to extract biomarkers (e.g. QRS and RR intervals). However, there is no standard methodology to delineate clinical ECG recordings, and existing approaches demonstrate significant variability in their solutions (Vančura et al., 2016).

Moreover, the information in ECG signals is confounded by anatomical and functional variability in the human population (Nguyên et al., 2015; Mincholé et al., 2019a). Non-invasive imaging, through ultrasound, computerised tomography

or cardiac magnetic resonance (CMR), is also used clinically to provide further information on cardiac anatomy, structure and mechanical function (Dall'Armellina et al., 2010; Tornvall et al., 2015; Dastidar et al., 2015; Biesbroek et al., 2018; Dastidar et al., 2019; Ferreira, 2019; Hausvater et al., 2019). Novel techniques to exploit the synergy obtained by combining ECG and imaging data and to robustly delineate ECG signals could improve data interpretation and clinical diagnosis (Berte et al., 2015).

Precision medicine aims to achieve customised healthcare therapeutic and diagnostic decision-making by combining multimodal clinical data to tailor the interpretation of test results to each patient's specific traits. Recent studies have shown the power of patient-specific image-based modelling and simulation for therapy guidance, arrhythmic biomarkers interpretation and patient's phenotypic variability interpretation (Lyon et al., 2018c; Bukhari et al., 2019; Niederer et al., 2019; Boyle et al., 2019; Martinez-Navarro et al., 2021). This technology has paved the way towards realising the 'digital twin' vision (Corral-Acero et al., 2020), referring to a comprehensive virtual tool that coherently integrates a patient's clinical data with mechanistic physiological knowledge and that can inform therapeutic and diagnostic decision-making through simulations. The generation of cardiac digital twins requires developing novel methodologies to analyse and estimate patient-specific ventricular activation properties from clinical tests data, such as ECG and CMR data.

1.2 Goals

This thesis aims to investigate new computational techniques for the efficient quantification of patient-specific ventricular activation properties using CMR-based modelling and simulation and non-invasive electrocardiographic data. To this end, I define the following specific goals:

1. To develop a novel method for the estimation of conduction properties from electrocardiographic signals using CMR-based modelling and simulation. I hypothesise that Bayesian inference methods combined with fast electrical

propagation models can efficiently explore the parameter space of the ventricular activation properties enabling to personalise cardiac models without high-performance computing infrastructure.

2. To evaluate the accuracy of the method in predicting conduction properties using epicardial activation time maps versus 12 lead-ECGs.

3. To develop a deep learning-based method for the delineation of the ECG. Accurate delineation is crucial for the automatic identification of specific features of the ECG signals (e.g. the QRS complex) used in the inference of conduction properties from clinical ECGs.

1.3 Outline

- Chapter 1 introduces the key concepts, motivation and goals of the thesis.
- Chapter 2 describes the thesis' physiological and computational background.
- Chapter 3 reviews the literature for the inference of ventricular activation properties, ECG delineation and machine learning strategies for ECG signals.
- Chapter 4 describes the methodologies developed for the inference studies presented in this thesis.
- Chapter 5 presents the application of the inference methodology (presented in Chapter 4) for recovering the human ventricular activation properties from epicardial activation time maps and CMR data.
- Chapter 6 (analogous to Chapter 5) presents the application of the inference methodology from 12-lead ECG-QRS signals.
- Chapter 7 presents a deep learning-based ECG-QRS multi-lead delineation method and validates it on clinical ECG recordings.
- Chapter 8 summarises the thesis scientific contributions and its conclusions and discusses possible future directions of this work.

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

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

— William Harvey (Harvey, 1628)

2 Background

2.1 The heart

The heart is a muscular organ that pumps blood through the cardiovascular system (Figure 2.1) to provide cells with the oxygen and nutrients they require and carry the output waste away to the organs responsible for disposing of it. The heart comprises four main chambers, following Figure 2.2 clockwise from the top-right: left atrium, left ventricle, right ventricle, and right atrium. Each of these parts has a distinct role in keeping blood circulating.

Oxygenated blood (following Figure 2.1 in the clockwise direction starting from the top-right) circulates from the lungs to the left atrium, then into the left ventricle, which pumps it through the aorta into the systemic circuit to provide nutrients to the rest of the body. Once depleted from nutrients, the blood returns from the systemic circuit to the heart's right atrium and enters the right ventricle, which pumps the blood back to the lungs, where it regains oxygen. Consequently, the ventricles have thicker walls than the atria (Figure 2.2) since they need to exert more force to move the blood outside the heart. Similarly, the left ventricle has thicker walls than the right ventricle as the former moves blood into the systemic circuit, whereas the latter only moves it to the pulmonary circuit (Figure 2.1).

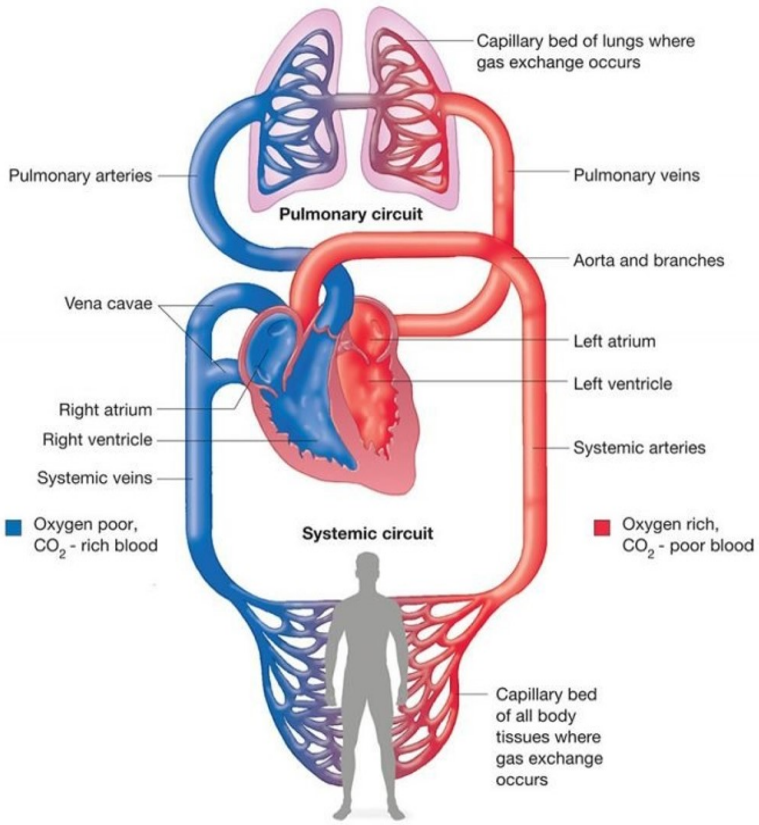

Figure 2.1: Schematic representation of the circulatory system. **Top** the pulmonary (lungs) circuit; **Mid** the four-chamber view of the heart; **Bottom** representation of the systemic circuit; **Right** (red) blood vessels carrying oxygenated blood out of the lungs into the heart, and, next, out of the heart to remaining parts of the body; **Left** (blue) blood vessels carrying blood with waste into the heart and next to the lungs to get oxygenated. From BOXYM, 2018, obtained under *fair dealing* / Cropped from original.

2.1.1 The healthy human cardiac cycle

In the initial state of the cardiac cycle, the ventricles are starting to relax (early ventricular diastole) and the atrial chambers are relaxed (Figure 2.3 – 'start'). This state causes the blood pressure outside the heart to be higher than inside the atria; thus, these get filled with blood. This would also be true for the ventricles; however, these should only be filled from the atria to ensure efficient blood circulation (Figure 2.1). Similarly, blood should never flow from the ventricles back into the atria. Thus, the ventricles are separated from the rest of the heart and circulatory system by blood pressure-gated valves (Figure 2.2 – top).

The atrioventricular valves (i.e. mitral or bicuspid valve in the left ventricle

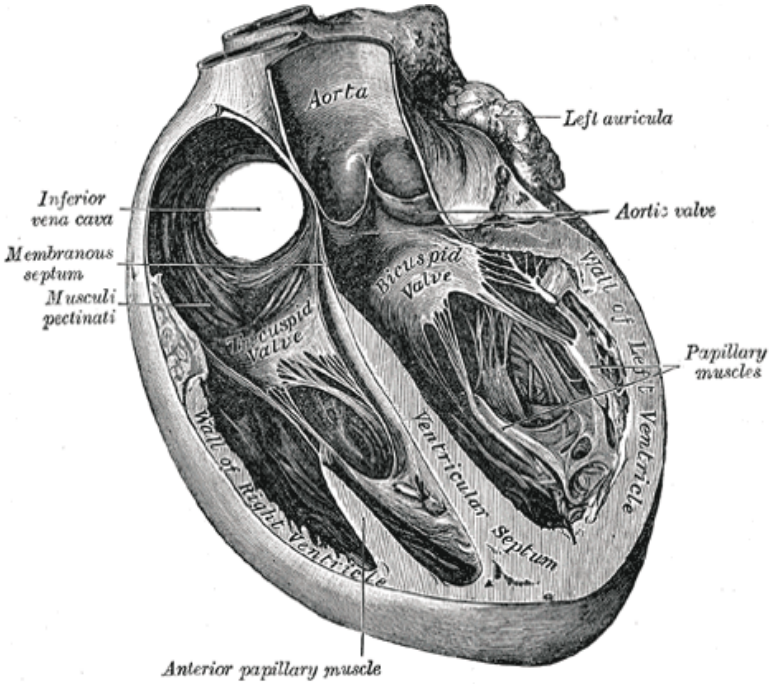

Figure 2.2: Four-chamber illustration of the human heart. Clockwise from top-right: left atrium, left ventricle, right ventricle, and right atrium. From Gray and Lewis, 1918, obtained under *fair dealing* / Cropped from original.

and tricuspid valve on the right ventricle) prevent retrograde blood flow from the ventricles back into the atria, an additional set of (semilunar) valves to prevent backflow into the ventricles from outside the organ. Consequently, during the cardiac cycle's initial state (Figure 2.3 – 'start'), despite the blood pressure outside the heart being greater than in the atria and ventricles, only the atria get filled with blood. Moreover, despite the residual pressure in the ventricles being larger than in the atria, the atrioventricular valves remain closed. This allows blood pressure to build up in the atria while the ventricles continue to relax.

Secondly (Figure 2.3 – 'ventricular filling'), the simultaneous relaxation in the ventricles and filling in the atria causes the atrioventricular pressure gradient to revert (i.e. greater blood pressure in the atria than in the ventricles), opening the atrioventricular valves. Thus, the ventricles, now relaxed (late ventricular diastole), start passively filling with blood.

Thirdly (Figure 2.3 – 'atrial contraction'), once the ventricles are almost full, the atria contracts (atrial systole) to top the ventricles up.

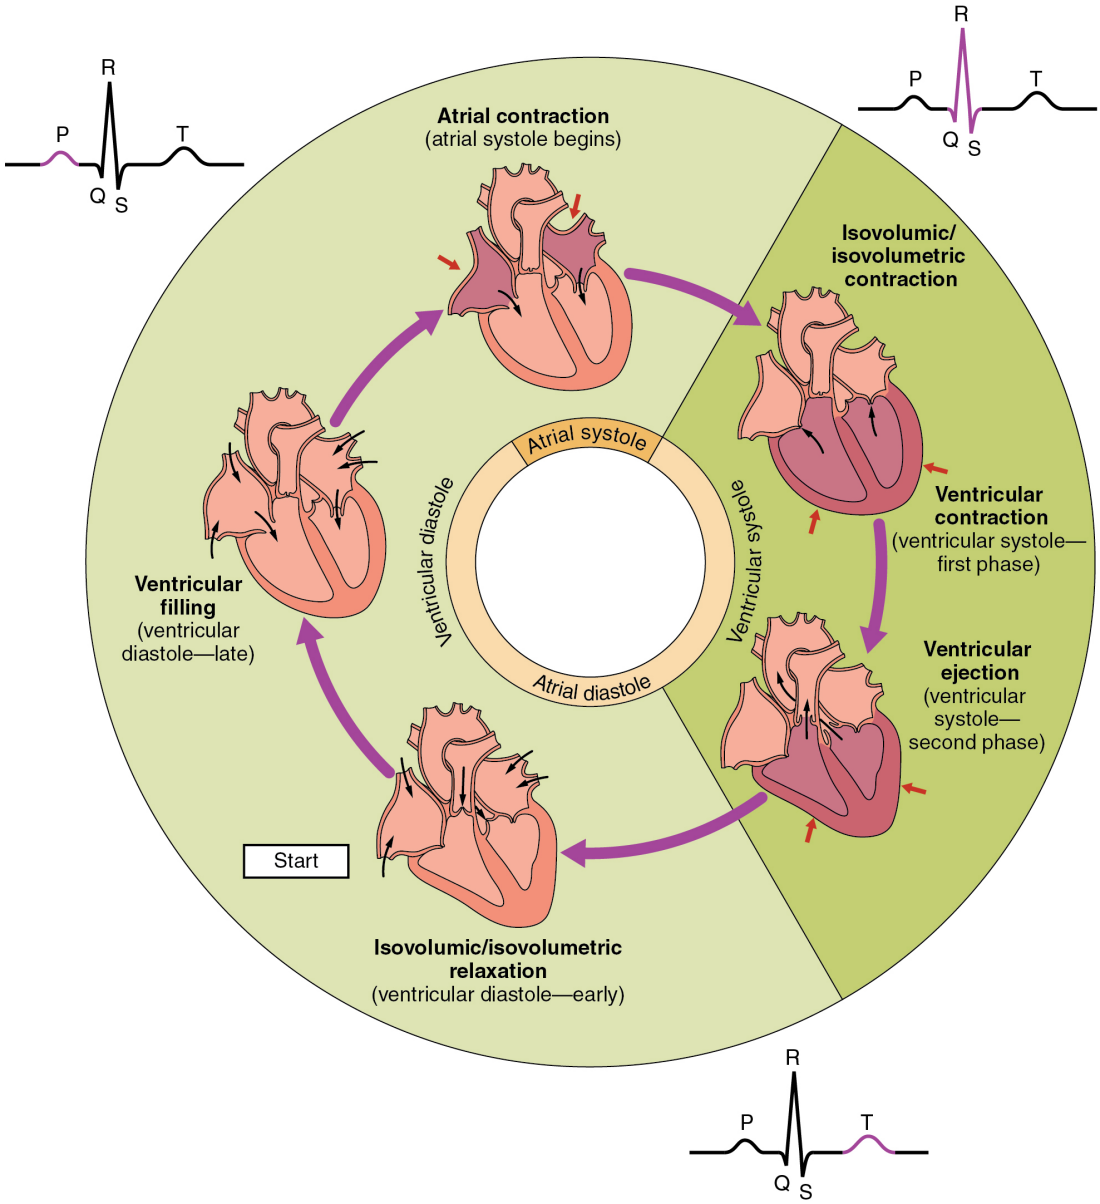

Figure 2.3: Cardiac cycle in a healthy human heart and its correspondence to the complexes in the ECG. From Commons, 2020, licensed under *Creative Commons Attribution 3.0 Unported CC BY 3.0*.

Fourthly (Figure 2.3 – 'isovolumetric contraction'), as the atrial contraction completes (end of atrial systole) and ventricular contraction commences (start of ventricular systole), the blood pressure gradient between atria and ventricles reverts (higher in the ventricles), causing the atrioventricular valves to close. The semilunar valves are still closed, and blood is incompressible; thus, the pressure builds up in the ventricles while the organ's shape and volume remain unchanged. Hence the name for this phase 'isovolumetric contraction'.

Fifthly (Figure 2.3 – 'ventricular ejection'), the contracting ventricles' blood pressure surpasses that in the aorta and pulmonary artery, causing the semilunar valves to open. This triggers the ventricular ejection phase and blood to flow through the aorta and pulmonary artery to the systemic and pulmonary circuits, respectively.

Lastly (Figure 2.3 – 'isovolumetric relaxation'), the ventricles' start to relax. This relaxation inverts the blood pressure gradient compared to the outside the heart, causing the semilunar valves to close. During this phase, no blood is flowing from the atria to the ventricles since the latter's pressure remains higher than in the former. Then the next cardiac cycle begins.

Maximising the heart's pumping function's efficiency requires continuously orchestrating the contraction times in all the individual cardiomyocytes throughout the atria and ventricles. This is done through perfectly timed electrical stimulation of the different regions in the heart.

2.2 Cardiac electrophysiology

The cardiac electrical clockwork is critical in the human body as it is directly linked to all other organs through the circulatory system (Figure 2.1). The cardiac muscles need to continue operating non-stop for the individual's lifespan. Therefore, the heart has evolved to be a highly efficient and sophisticated pump synchronised with perfectly timed electrical signals.

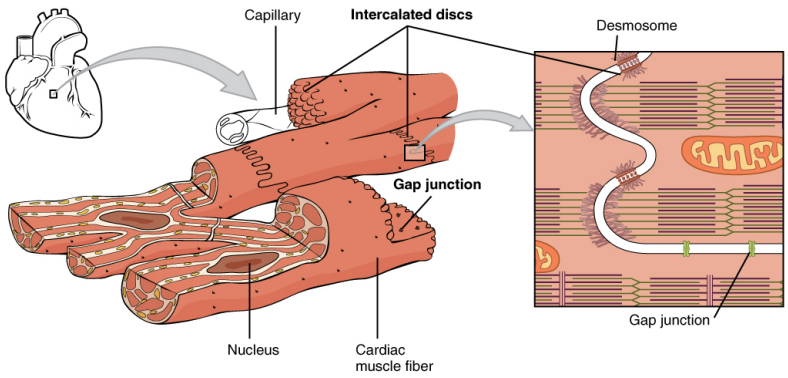

Figure 2.4: Cardiomyocytes with gap junctions. From Commons, 2020, licensed under *Creative Commons Attribution 3.0 Unported CC BY 3.0*.

2.2.1 Cellular cardiac electrophysiology

The heart's muscle tissue is composed of cardiac muscle cells (i.e. cardiomyocytes) (Figure 2.4). Cardiomyocytes are electrically and mechanically connected through intercalated disks, forming tissue fibres arranged into sheets of muscular tissue. These cells are also electrically coupled to each other via gap junctions (Figure 2.4), allowing electrical currents to propagate through the cardiac tissue. This allows cardiomyocytes to synchronise through electrical currents, allowing the heart's four chambers to operate in perfect harmony to optimise blood circulation.

At the cellular level, the heart's mechanical contraction is triggered through electrical stimulus propagated through the cardiac tissue (Figure 2.4). Cardiomyocytes can be excited by external currents and then excite other cells connected to them through gap junctions (Bers, 2002). Cardiomyocytes are characterised by having voltage-gated channels (Nerbonne and Kass, 2005). These channels allow cardiomyocytes to be excited (i.e. depolarised) and initiate action potentials (Figure 2.5), which will excite their neighbouring cells. These action potentials initiate when the cardiomyocyte's membrane potential reaches a certain threshold.

The main ions involved in the ventricular cardiomyocyte's action potential (Figure 2.5) are sodium (Na^+), potassium (K^+), and calcium (Ca^{2+}). When at rest, the cardiomyocyte has a membrane potential of about -90 millivolts (mV) (Figure 2.5 – phase 4) (Grant Augustus O., 2009). However, the cardiomyocyte has fast Na^+ channels gated at around -70 mV. Thus, when the cardiomyocyte

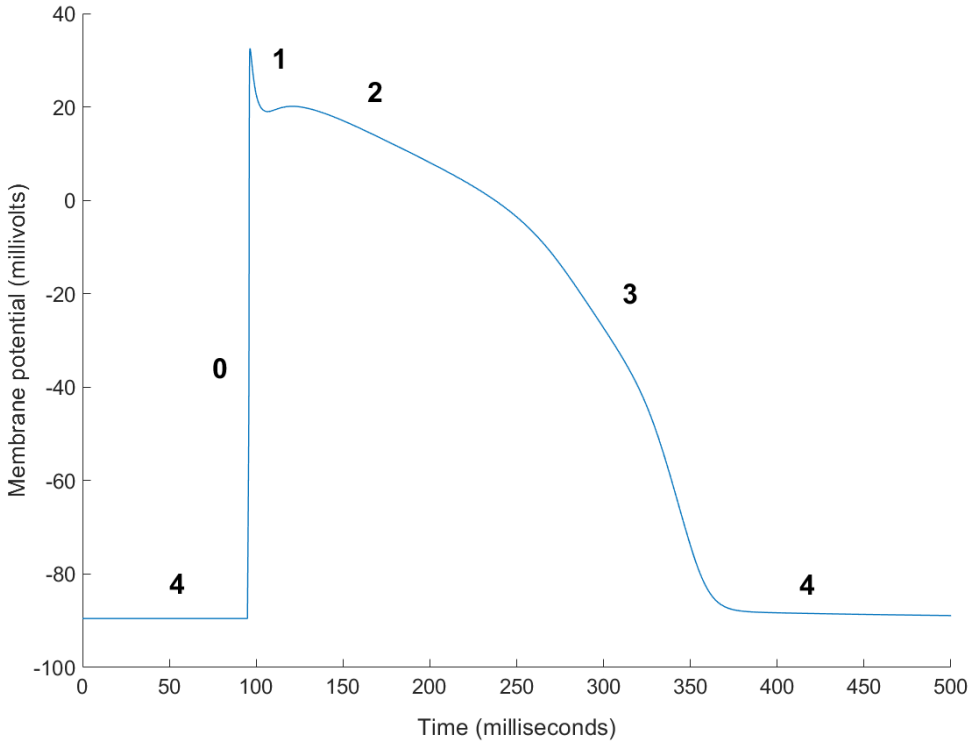

Figure 2.5: Human ventricular action potential with its phases. **4** resting phase – -90 millivolts. **0** depolarisation (upstroke) – Na^+ entering the cardiomyocyte (in) – 35 millivolts. **1** early repolarisation – K^+ leaving the cardiomyocyte (out) – 20 millivolts. **2** plateau phase – K^+ out and Ca^{2+} in. **3** repolarisation – K^+ out. Simulated using Tomek et al., 2019.

receives a stimulus from its neighbouring tissue that causes its potential to reach around -70 mV, the fast inward Na^+ channels will open, producing the action potential's upstroke (Figure 2.5 – phase 0) (i.e. depolarisation phase). During this phase, Ca^{2+} starts entering the intra-cellular space. This causes the membrane potential to increase, causing most Na^+ channels to inactivate and K^+ to start leaking out of the cell (Figure 2.5 – phase 1). However, before the cardiomyocyte's potential is able to get back to negative values, the sarcoplasmic reticulum releases Ca^{2+} into the intra-cellular space, which cancels out the changes in potential from the outward flow of K^+ (Figure 2.5 – phase 2). These Ca^{2+} influx together with the remaining open Na^+ channels determine the duration of the plateau phase. The cardiomyocyte's contraction will mostly take part during this plateau phase. Finally, the Ca^{2+} channels will inactivate, and the K^+ outflow will repolarise the cell

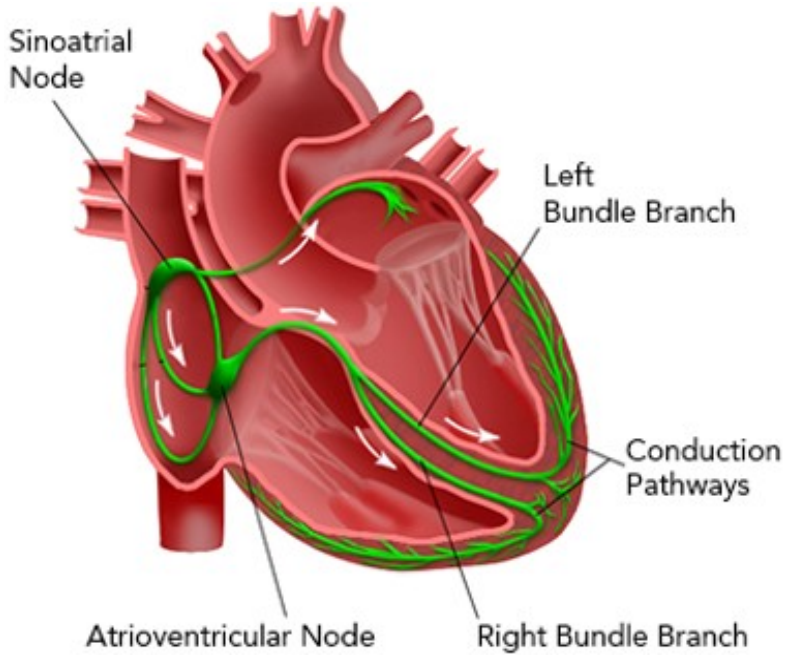

Figure 2.6: Four-chamber illustration of the electrical system in the human heart. Clockwise from top-left: right atrium with the sinoatrial node and the atrioventricular node, left atrium, left ventricle with the left bundle branch, and right ventricle with the right bundle branch. From Englewood Hospital, 2020, obtained under *fair dealing* / Cropped from original.

(Figure 2.5 – phase 3) until it gets back to its resting state (Figure 2.5 – phase 4).

Throughout the cardiac cycle (Figure 2.3), these cardiomyocytes are stimulated through gap junctions at specific timings that cause the heart to operate efficiently. The part of the heart responsible for delivering and controlling the electrical stimuli to the cardiomyocytes to orchestrate the cardiac cycle is the cardiac electrical conduction system.

2.2.2 The cardiac electrical conduction system

The cardiac electrical conduction system (Figure 2.6) is a network of fast conducting cells responsible for controlling the heart's pumping behaviour. These cells are characterised by having automaticity. The heart's normal activity, known as sinus rhythm, is orchestrated by the sinoatrial node (hence the name sinus rhythm) (Figure 2.6 – top-left). Thus, during sinus-rhythm, the rest of the cells in the cardiac electrical conduction network are stimulated through gap junctions.

The cardiac electrical conduction system (Figure 2.6) starts in the right atrium from the sinoatrial node (Figure 2.6 – top-left). Then, it branches to the left atrium (Figure 2.6 – top-right) and the atrioventricular node. From there, the conduction system enters the ventricles through the septum (i.e. his-bundle) (Figure 2.6 – centre). The ventricular part of the cardiac electrical conduction system is commonly referred to as the Purkinje network.

The Purkinje network splits into right and left bundle branches, responsible for the right and left ventricles. Further down, close to the apex (Figure 2.6 – bottom-right), the two bundles split into conduction pathways connected to the endocardial tissue (i.e. inner ventricular tissue layer) in both ventricles. The Purkinje network in humans connects only to the endocardial layer (via Purkinje-myocardial junctions) and is electrically isolated elsewhere (Durrer et al., 1970; Vigmond and Stuyvers, 2016). These Purkinje-myocardial connection points are the first ventricular locations to depolarise during sinus-rhythm activity. Thus, these locations are commonly referred to as the earliest activation sites or root nodes. Propagation of electrical excitation is faster through the Purkinje network than muscle cardiac tissue.

2.2.3 Cardiac electrical propagation

The electrical orchestration of the cardiac cycle (Section 2.1.1) is controlled by the electrical conduction system in the heart (Section 2.2.2), the electrical conduction speeds in the cardiac muscle tissue, and the dispersion in action potential duration in cardiomyocytes in different regions of the organ. The asymmetry in the electrical conduction speeds in the ventricular cardiac tissue caused by the cardiomyocytes' arrangement (Figure 2.7) has a key role in determining the activation sequence in the ventricles.

Cardiomyocytes organise themselves into fibre-like structures (Figure 2.4), which are aligned, forming muscle tissue sheets. This organised structure gives rise to differences in electrical conduction speeds (Caldwell et al., 2009). From here, we

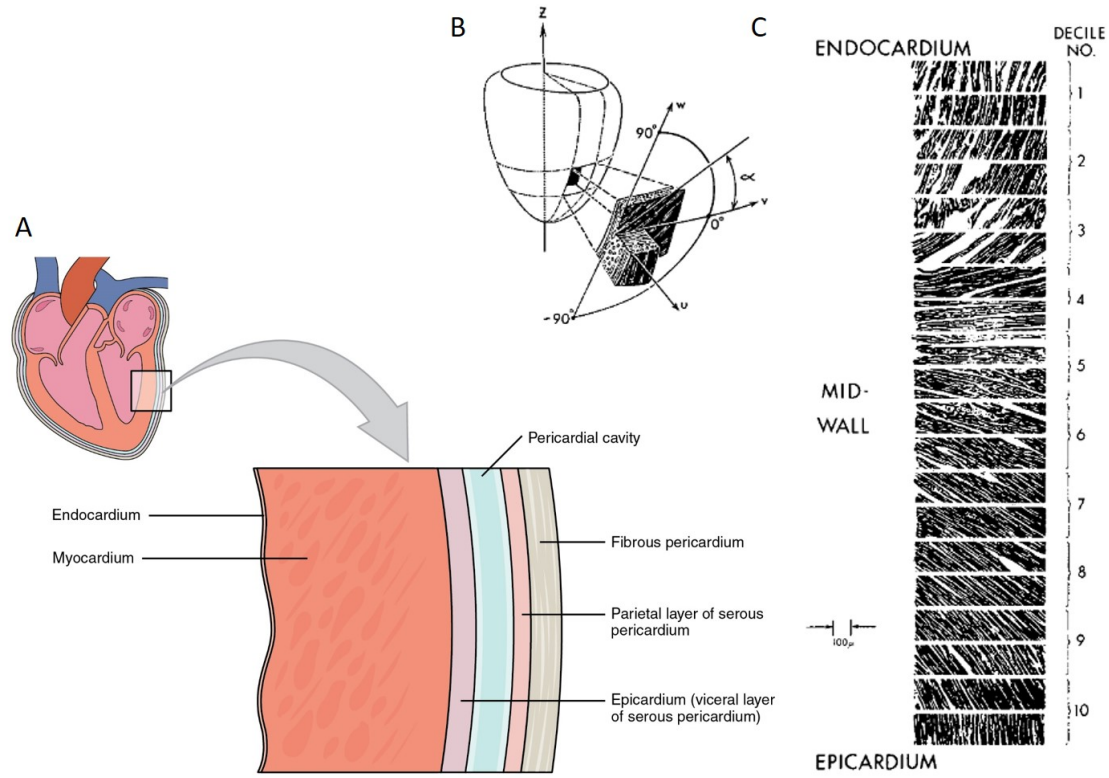

Figure 2.7: Transmural properties of the cardiac wall. **A** illustrates the transmural tissue layers in the heart, from inside outwards: endocardium, myocardium, and epicardium. **B** experimental data showing the transition in fibre angle from the endocardium to the epicardium. **C** schematic illustration showing the correspondence of the angles in **B** with the organ in space. Sources: **A** from Commons, 2020 (licensed under *Creative Commons Attribution 3.0 Unported CC BY 3.0*), **B** and **C** from Streeter et al., 1969 (obtained under *fair dealing* / Cropped from original).

distinguish three orthogonal directions relative to the orientation of the connected-cardiomyocytes (through gap junctions) (Figure 2.4), which form fibres of tissue. The electrical conduction is fastest in the direction of these fibres (i.e. fibre-directed speed). These fibres organise themselves into sheets of tissue; thus, the second-fastest direction of propagation is perpendicular to the fibre but within the tissue sheet. Lastly, the slowest conduction speed direction is the one perpendicular to the sheet of tissue. The orientation of these fibres of cardiac tissue changes transmurally in the heart (Figure 2.7.C) (Streeter et al., 1969), causing changes in electrical conduction speed. These transmural differences were measured by Streeter et al., 1969 in canine to go from 60 to -60 degrees from endocardium to epicardium (Figure 2.7.B). This asymmetry in the fibre-directed conduction speeds together with the root nodes (i.e.

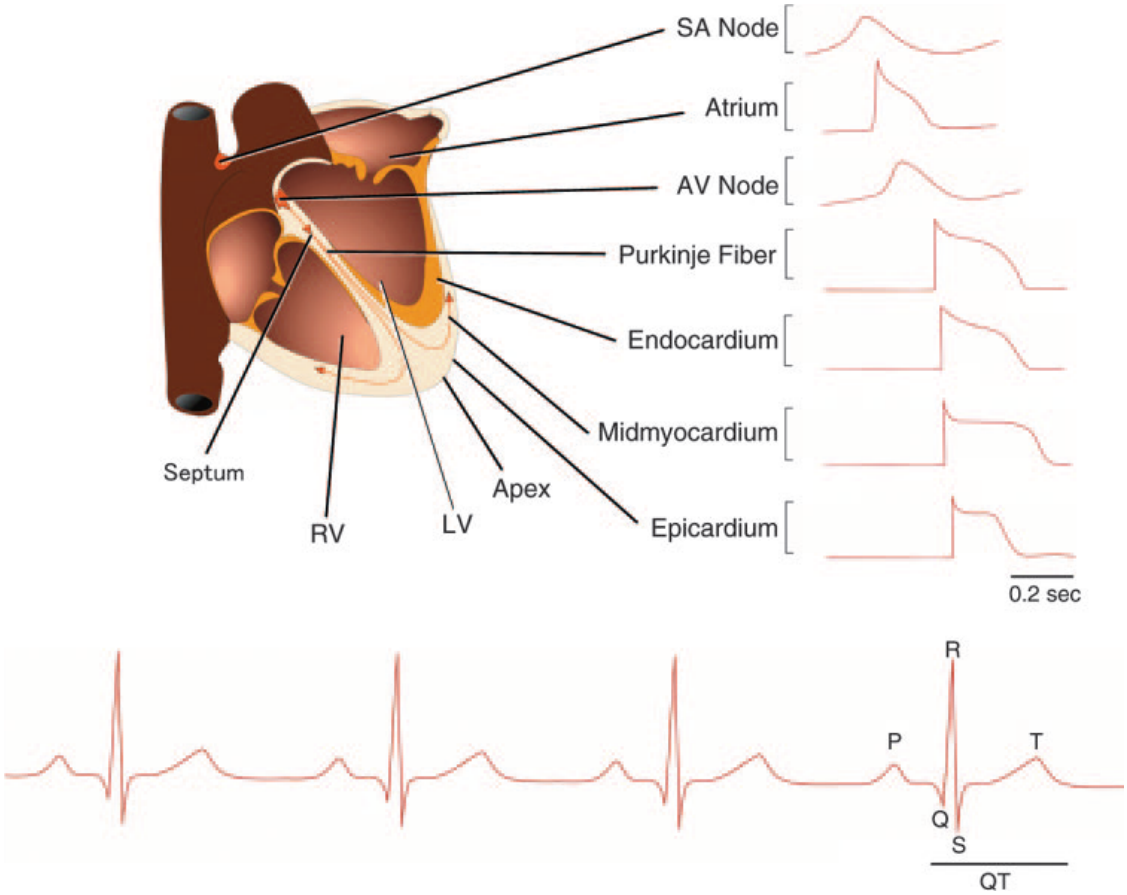

Figure 2.8: Electrical activity in the myocardium. *Top*: Schematic of a human heart with an illustration of typical action potential waveforms recorded in different regions. *Bottom*: schematic of a surface electrocardiogram; three sequential beats are displayed. From Nerbonne and Kass, 2005, Copyright © 2005 the American Physiological Society, obtained under *fair dealing*.

earliest activation sites) from the cardiac electrical conduction system synchronises the mechanical contraction of the ventricles (Figure 2.3).

The third actor in the synchronisation of the cardiac cycle is the heterogeneities in action potential duration manifested by cardiomyocytes in different parts of the heart (Figure 2.8). For example, the outer layer of tissue in the heart (i.e. epicardium) is the last part to depolarise as the activation wavefront starts from the root nodes located at the inner-most layer in the ventricles (i.e. endocardium) (Figure 2.8). Nevertheless, the epicardium needs to repolarise (relax) first to allow for efficient transmural relaxation of the organ mechanically. Similarly, the apex tissue is excited before that in the ventricles' base; however, for a uniform relaxation,

the apical cardiomyocytes have also shorter action potential duration. These heterogeneities synchronise the cardiac mechanical function during the ventricular diastole (repolarisation) phase of the cardiac cycle (Figure 2.3).

These characteristics affect the electrical wavefronts' patterns moving the heart, which are directly linked to its subsequent mechanical behaviour. Therefore, the heart's electrical activity study can explain cardiac behaviour and identify abnormalities when present.

2.3 The electrocardiogram

The electrocardiogram (ECG) is the most widely used clinical modality for cardiac disease diagnosis. Surface electrocardiography (also referred to as the ECG test) is a non-invasive test that records the heart's electrical activity as signal information. This electrical activity originates from the cardiomyocytes' excitation (i.e. action potentials) and propagates on an organ level. Then it can be recorded using electrodes attached to the patient's skin. These readings gathered from the electrodes (i.e. electrograms) can be combined with a far-field reference to compute potential differences (i.e. ECG leads). Each lead in the ECG represents the heart's electrical activity along the axis defined by the measurement and reference electrodes' spatial alignment.

Each heartbeat in the ECG comprises a P-wave, a QRS complex, and a T-wave (Figure 2.9). These complexes correspond to different phases of the cardiac cycle (Figure 2.3). As mentioned in Section 2.2.2, the cardiac cycle is initiated by the sinoatrial node and the atria's subsequent depolarisation (Figure 2.6 – top-left and Figure 2.3 – ventricular filling). The depolarisation of the atria produces a positive deflection in the ECG referred to as the P-wave. Next, the atria iso-electrically contracts to top up the ventricles, so there is a period of no recorded electrical activity given by the delay of the atrioventricular node (Figure 2.9 – PR segment). Then, the ventricles depolarise (Figure 2.3 – start of ventricular systole), which given the large myocardial-mass volume compared to the atria, produces a larger deflection giving rise to the Q (depolarisation of the septum),

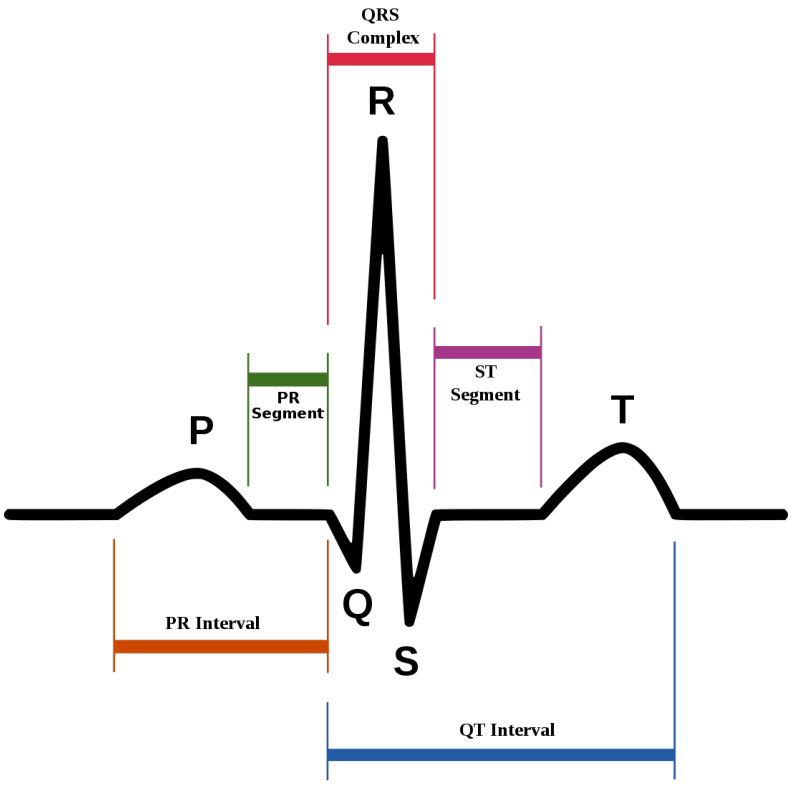

Figure 2.9: Single-beat human ECG during sinus-rhythm electrical activity with its complexes highlighted. From Commons, 2020 (*public domain*).

R (depolarisation of the ventricular free-walls), and S (depolarisation of the basal tissue) waves, giving its name to the QRS complex. During the depolarisation of the ventricles, the atria repolarise; however, the atria's repolarisation's electrical activity gets masked out by the larger amplitudes produced by the depolarisation of the ventricular cardiomyocytes. This is followed by the second period of no electrical activity (Figure 2.9 – ST segment) to allow the ventricles' muscles to contract and eject the blood out of the heart. Lastly, the ventricles repolarise (Figure 2.3 – isovolumetric relaxation), producing the T-wave. The T-wave is usually positive because, despite having an opposite sign to the repolarisation wavefront, the depolarisation wavefront also moves in the opposite direction since the epicardial myocytes are last to depolarise but first to repolarise (Figure 2.8).

2.3.1 The 12-lead ECG test

Several ECG test protocols exist, such as the standard 12-lead resting ECG test, ambulatory ECG tests (also referred to as Holter ECG tests), and stress ECG tests. These tests will vary in the number of leads and duration depending on the purpose of the application—for example, the 12-lead resting ECG test records from ten electrodes for ten seconds (Sörnmo and Laguna, 2005). Thus, the 12-lead ECG test can identify anatomical conditions, such as deformations of the heart like hypertrophy. On the other hand, ambulatory ECG tests are usually collected to detect events suspected of taking place but that were not observed during a 12-lead ECG resting test. Therefore, ambulatory ECG recordings will be between 24 and 48 hours long and devoted to identifying non-persistent abnormal events (e.g. arrhythmias). Ambulatory ECG recordings are typically composed only of two limb leads since carrying more electrodes for a whole day would significantly interfere with the patient's behaviour throughout the recording. Stress ECG is acquired from the exercise stress test. This test is meant to capture the heart's capability to cope with physical activity, aiming to reveal abnormalities hidden when resting, such as insufficient blood supply. Stress ECG recordings will usually last for a few minutes. Other types of ECG recordings include intensive care unit ECG monitoring, which takes place in a controlled environment, allowing the usage of different lead configurations depending on the patient's condition and will be as long as the life-threatening period (e.g. sometimes weeks long).

The standard 12-lead ECG uses ten strategically positioned electrodes to maximise the clinical interpretability of its recordings. Each electrode is named according to its placement. Four electrodes are placed in the limbs (Figure 2.10): right arm, left arm, left leg, and right leg. These serve to generate six limb leads (i.e. three limb and three augmented limb leads). The six remaining (precordial) electrodes are positioned in the torso's anterior left side (Figure 2.11): V1, V2, V3, V4, V5, and V6. These precordial electrodes serve to compute the six precordial leads (which share the electrodes' names).

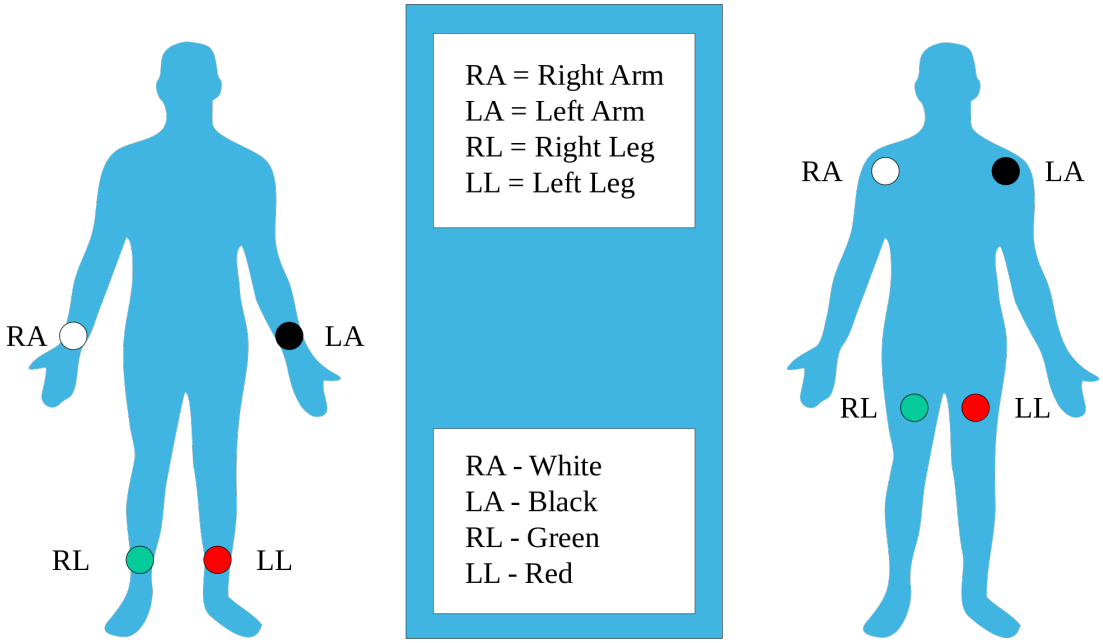

Figure 2.10: The two equivalent placement configurations (left and right) of limb electrodes for the 12-lead ECG test. From Commons, 2020 (*public domain*).

The three limb leads (Figure 2.12): lead I (one), lead II (two) and lead III (three) are calculated as follows

$$\mathrm{I} = \mathrm{LA} - \mathrm{RA} \ , \tag{2.1}$$

$$\mathrm{II} = \mathrm{LL} - \mathrm{RA} \ , \tag{2.2}$$

and

$$\mathrm{III} = \mathrm{LL} - \mathrm{LA} \ , \tag{2.3}$$

where LA, RA, LL, and LA are the limb electrograms (Figure 2.10). Note that lead III can be derived as a linear combination of leads I and II as

$$\mathrm{III} = \mathrm{II} - \mathrm{I} \ ; \tag{2.4}$$

therefore, most ECG databases omit lead III.

Wilson's central terminal, $\mathrm{V_{WCT}}$, is a virtual electrogram representing the reference at the centre of the heart computed as

$$\mathrm{V_{WCT}} = \frac{1}{3}\mathrm{RA} + \mathrm{LA} + \mathrm{LL} \ . \tag{2.5}$$

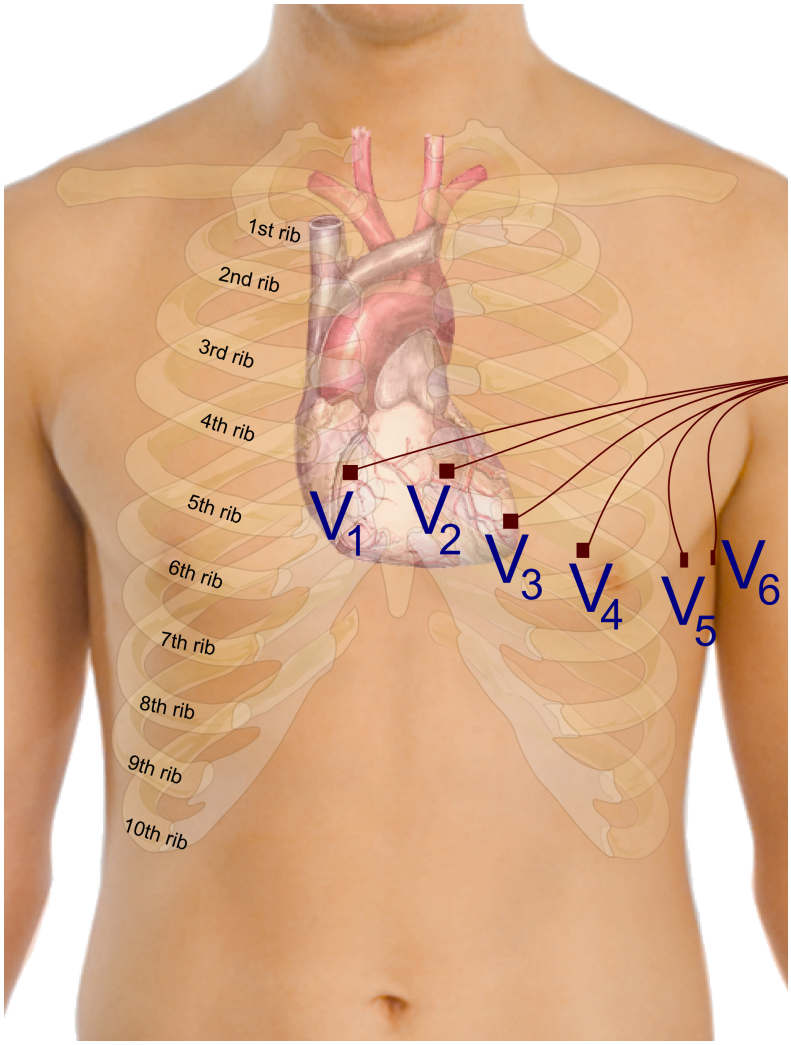

Figure 2.11: Placement of precordial electrodes for the 12-lead ECG test. From Commons, 2020, licensed under *Creative Commons CC0 1.0 Universal*.

Wilson's central terminal is combined with the limb electrograms to generate three augmented leads (Figure 2.12): augmented vector right (aVR), augmented vector left (aVL) and augmented vector foot (aVF). The augmented limb leads cover the remaining views in the 'frontal' plane of the heart. These are computed as

$$\mathrm{aVR} = \mathrm{RA} - \frac{1}{2}(\mathrm{LA} + \mathrm{LL}) = \frac{3}{2}(\mathrm{RA} + \mathrm{V_{WCT}}) \ , \tag{2.6}$$

$$\mathrm{aVL} = \mathrm{LA} - \frac{1}{2}(\mathrm{RA} + \mathrm{LL}) = \frac{3}{2}(\mathrm{LA} + \mathrm{V_{WCT}}) \ , \tag{2.7}$$

and

$$\mathrm{aVF} = \mathrm{LL} - \frac{1}{2}(\mathrm{RA} + \mathrm{LA}) = \frac{3}{2}(\mathrm{LL} + \mathrm{V_{WCT}}) \ . \tag{2.8}$$

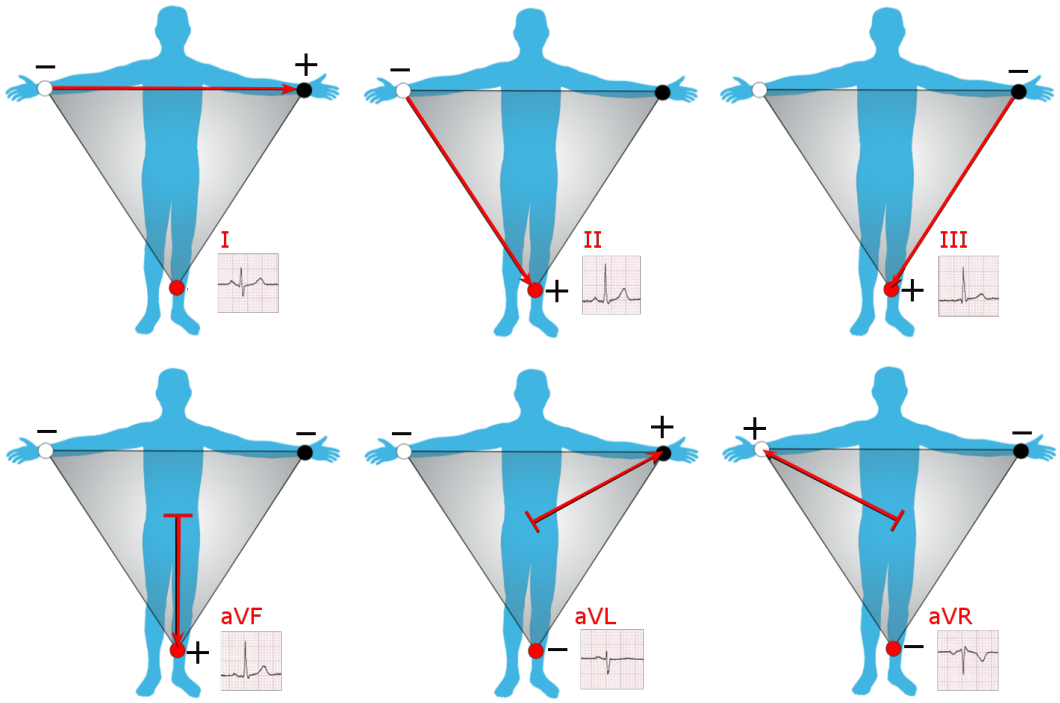

Figure 2.12: Directionality of electrical activity represented in the limb leads and augmented limb leads of the 12-lead ECG test. (Top-row) limb leads; (Bottom-row) augmented limb leads. I – lead I; II – lead II; III – lead III; aVF – augmented vector foot; aVL – augmented vector left; aVR – augmented vector right. From Commons, 2020, licensed under *Creative Commons Attribution-ShareAlike 4.0 International CC BY-SA 4.0*.

The precordial leads (Figure 2.13) are also calculated using Wilson's central terminal as reference, namely,

$$\mathrm{V}_i = \mathrm{V}'_i - \mathrm{V}_{\mathrm{WCT}} \ , \tag{2.9}$$

where V_i is the ith precordial lead (from one to six), and V'_i is the corresponding electrogram measured by the ith precordial electrode.

The leads composing the 12-lead ECG provide a reading of the heart's electrical propagation from several different perspectives. The limb and augmented limb leads represent the electrical activity across the 'frontal' plane (from arms to feet) of the heart (see Figure 2.12). Simultaneously, the precordial leads represent electrical propagation on the 'depth' plane (from chest to back) of the heart (see Figure 2.13).

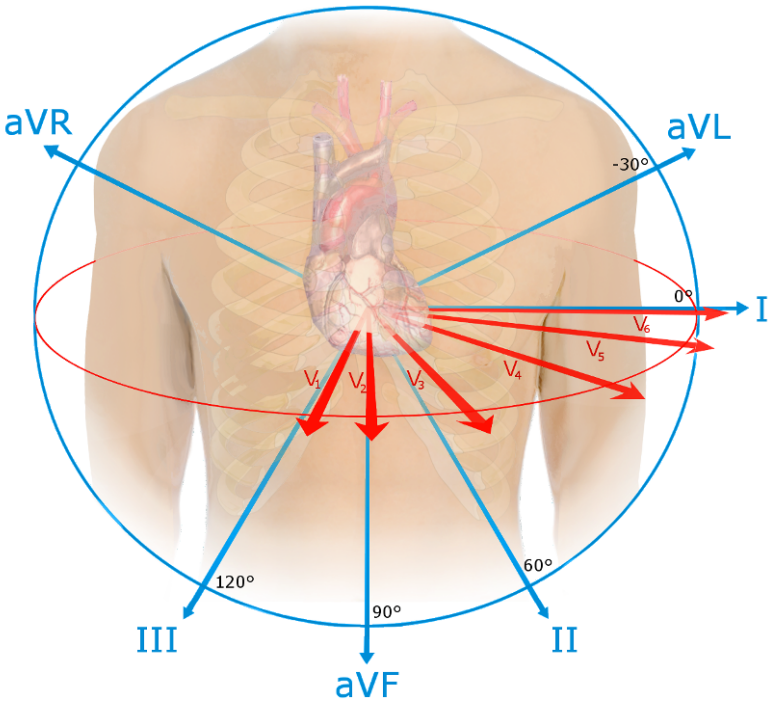

Figure 2.13: Directionality in 3D space of the 12 leads in the 12-lead ECG resting test. Limb and augmented limb leads (blue) represent the electrical activity on the arms-to-feet plane. Precordial leads (red) represent the electrical activity on the chest-to-back plane. From Commons, 2020, licensed under *Creative Commons Attribution-ShareAlike 4.0 International CC BY-SA 4.0*.

2.4 ECG analysis

In medical practice, ECG analysis is performed by computing biomarkers from the signals. These biomarkers are computed using ECG signals already delineated. For example, the RR interval is a commonly computed biomarker and requires having previously identified all the R-peaks in the signal. These characteristics summarise the heart's electrical activity behaviour and allow clinicians to cluster and classify patients to diagnose cardiac conditions (Lyon et al., 2018b). Therefore, the analysis of these signals' is typically divided into the following steps: noise-removal, delineation, biomarkers' computation, and interpretation of these biomarkers (e.g. diagnostic and therapeutic decision-making).

2.4.1 ECG noise-removal

ECG filtering removes undesired noise, which is non-relevant for the analysis. The most common types of noise found in an ECG recording are the baseline wander, powerline interference, and muscular generated electrical activity. More precisely, most ECG recordings will include baseline wander and powerline interference. Ambulatory ECG recordings will include all types of noise, depending on the patient's location and activity. Stress test recordings are characterised by having high levels of muscular noise throughout the entire record. Each of these types of noise is dealt with in a specific manner.

The baseline wander is removed to omit non-cardiac information introduced by the body movements, such as respiration or walking. Hence the frequencies associated with baseline wander will usually be lower than the ones originated by cardiac activity. Thus, this type of noise is removed with high-pass filters with constant cut-off frequencies or time-variant ones (Sörnmo and Laguna, 2005). Time-variant cut-off frequencies will allow the filter to adapt to heart rate changes and, thus, remove more significant noise levels than a filter that was fixed to the lowest heart rate possible to avoid loss of relevant information. Alternatively, one can fit a polynomial as cubic splines to the baseline wander (Sörnmo and Laguna, 2005); however, this last alternative requires having parts of the signal already delineated, which is considered a later step of the ECG analysis pipeline.

Powerlines cause electromagnetic fields which interfere with electrical recording devices when these are in range. This noise is mainly composed of frequencies of about 50 Hz (or 60 Hz, depending on the device's making) but with larger amplitudes than some of the main waveforms composing the ECG signals (e.g. the P-wave). Consequently, if not removed, powerline interference complicates the delineation of the low-amplitude waveforms in the ECG. Fortunately, the main waveforms composing the ECG signals (i.e. P-wave, QRS complex, and T-wave) are mostly composed of frequencies lower than 50 Hz (Sörnmo and Laguna, 2005). Powerline noise can be removed with a notch filter.

Muscle noise is the most challenging type of noise to remove. This type of noise is characteristic of stress tests and Holter ECG recordings when the patient is exercising. The spectral content of muscle noise overlaps with the ECG's main waveforms. Removing muscle noise is currently understood as a trade-off that must be carefully adjusted since most techniques implemented to reduce this type of noise will distort relevant heart-generated information in the ECG.

Once filtered, the ECG recordings undergo the process of delineation. This process aims to identify critical points of the waveforms composing each beat. Delineating is the most challenging step within the ECG analysis pipeline to be automated since it requires characterising all beats within the recording without prior information on the signal. The existing techniques for ECG delineation are reviewed in Section 2.4.2. The delineation marks allow computing features of interest from specific intervals (Maršánová et al., 2017). For example, these marks serve to compute the duration of the QRS and the QT and properties of the segment, such as ST-segment elevation, which has been correlated to myocardial infarction (Bliek, 2018).

2.4.2 ECG delineation

In daily clinical practice, the analysis of ECG recordings involves calculating reference biomarkers such as heart rate, QRS width, QT interval and ST segments. These biomarkers are fundamental for practitioners and automatic decision-making systems to determine if a patient has a specific cardiac condition and monitor events such as arrhythmias. Calculating these biomarkers requires positioning critical points (i.e. fiducials) within each beat, such as the QRS complex onset time (QRS onset), the QRS complex offset (QRS offset), to compute the QRS width within each beat (Figure 2.9). ECG delineation is the process of identifying these points.

Manually delineating ECG signals to calculate biomarkers is a tedious and repetitive task. Therefore, most physicians use commercial software solutions which perform the delineation as a preprocessing step to extract biomarkers of interest. Clinicians then combine their expert knowledge with the results from the

ECG analysis to make decisions. Consequently, datasets obtained from medical institutions do not include delineation marks since these are considered a preliminary computation of the ECG analysis program. Instead, ECG databases usually include ECG recordings, medical records, and, occasionally, a summary of the biomarkers calculated by the ECG analysis tool. In other words, the clinician's decision-making process will consider the summary of biomarkers together with a visual inspection of some parts of the data and their expertise on interpreting the results of the reports of the program. However, commercial ECG analysis products are known to have high variability in the calculation of the biomarkers (Vančura et al., 2016). This constitutes a problem when trying to transfer knowledge between patients analysed using different software. A standardised delineation method would enable recomputing the biomarkers without the program dependent variability and correlating the newly computed biomarkers to the clinicians' decisions. Thus, it would allow merging ECG diagnostic databases with consistent biomarkers and known diagnosis to exploit big data within the ECG context from a comprehensive biomarker-based perspective.

The state-of-the-art ECG delineation systems were designed based on available databases. Human ECG signals display significant morphological and temporal variability (Figure 2.14). However, delineation databases are small and consider selection criteria to reduce the variability in their data. For example, the QT database (Laguna et al., 1997) considered only recordings with low noise levels and slow heart rate. In other words, available ECG delineation systems require adaptations to deal with signals from different databases than those considered in the delineator's original design. Moreover, such adaptations require expertise in both the system's implementation and the incoming database to identify which delineator's parameters should be modified to account for the new database. The overhead required to get these systems to work each time has led to the omission of delineation in ECG analysis related research. Machine learning techniques, on the other hand, can automatically identify which characteristics are more relevant to explain the data, omitting the need for experts to recalibrate them to new databases.

Figure 2.14: Example of the variability in human ECG signals. These segments of different ECG recordings and leads extracted from the MIT-BIH Arrhythmia database (Moody and Mark, 2001) serve to illustrate the morphological and amplitude variability between beats and recordings in a conventional clinical database. The red grid is discretised at 20 ms on the x-axis and at 0.5 mV on the y-axis.

The final step in ECG analysis is the interpretation of its biomarkers. The variability in the biomarkers extracted from these signals can explain groups of patients associated with cardiac conditions. However, variability in inter-patient anatomical differences also gets translated into variability in their ECGs. This anatomical variability interference can cause the same biomarker to signify different information between patients or to hide disease biomarkers. Therefore, the study of this variability requires integrating domain knowledge on the heart's mechanistic aspects producing the observed ECG with its analysis. For example, combining ECG analysis with the physiological information embedded in patient-specific cardiac models will significantly increase the recordings' interpretability by ruling out unrealistic interpretations given the anatomical characteristics of the patient

being observed (Mincholé et al., 2019a).

2.5 Non-invasive electrocardiographic imaging

Electrocardiographic imaging (ECGI) is the family of methodologies that aim to estimate the source of the cardiac electrical activity from the non-invasively measured (torso's) body surface potentials and medical imaging data, such as cardiac computed tomography or CMR (Cluitmans et al., 2015). Measurements for ECGI are body surface potential maps composed of signals recorded from various electrodes distributed on the subject's torso (Ghanem et al., 2005; Giffard-Roisin et al., 2017a; *Cardioinsight Mapping Vest*). The task of estimating the source to the measured data is challenging due to its ill-posedness, its non-unique solution, and the non-linear interplay of the cardiac properties (e.g. earliest activation sites and conduction speeds) (Colli-Franzone et al., 1985). Most ECGI approaches in the literature employ regularisation-based (Tikhonov and Arsenin, 1977; Colli-Franzone et al., 1985; Wang and Rudy, 2006) strategies to narrow down the most plausible epicardial source data to the measured body surface potential maps (Cluitmans et al., 2015). Despite their limitations (Bear et al., 2019), ECGI methods continue to improve and are fast becoming an alternative to invasive cardiac mappings for therapy planning and cardiac diagnosis (Andrews Christopher M. et al., 2017; Pereira et al., 2020).

The 3D spatial representation in ECGI-based reconstructions of the epicardial potentials eases the extraction of spatiotemporal information that reveals the cardiac mechanistic behaviour. For example, ECGI can enable identifying conduction speed heterogeneities, ectopic regions, and other arrhythmic substrates (Cluitmans et al., 2015; Andrews Christopher M. et al., 2017; Kara et al., 2019). Therefore, ECGI solutions can serve for model personalisation, diagnosis, and therapy guidance (Cluitmans et al., 2018; Coll-Font et al., 2019). For example, Cuculich et al., 2017 proposed an ECGI-based guidance tool that combined with radiation-based ablation (repurpose of radiotherapy for ablation), can non-invasively guide and apply cardiac ablation in a few minutes without post-operational recovery requirements.

ECGI approaches are known to introduce artefacts and variability in their solutions through their computation and regularisation (Ghanem et al., 2005; Serinagaoglu Dogrusoz et al., 2019). With the increased cost and equipment required for conducting ECGI, these challenges have prevented these technologies from being adopted for standard screening practice to replace the 12-lead ECG. However, the steady advancements and innovations in ECGI technologies suggest that these techniques may be ready to be considered for standard practice within the next two decades. Therefore, for this thesis, I consider that epicardial activation maps can be non-invasively recovered through ECGI.

2.6 Cardiac multiscale modelling and simulation

2.6.1 Bidomain model

The bidomain model is the gold standard formulation for cardiac electrophysiology simulations. The bidomain assumes that there are two electrical subspaces in the heart, namely, the intracellular and the extracellular spaces; thus, an intracellular and an extracellular potential at every point in space, ϕ_i and ϕ_e, respectively. Therefore, transmembrane potential, V_m, is defined as $V_m = \phi_i - \phi_e$. The current densities, $\bar{J}_i$ and $\bar{J}_e$, in these two spaces, are defined as

$$\begin{aligned} \bar{J}_i &= g_{ix}\frac{\partial \phi_i}{\partial x}\bar{a}_x + g_{iy}\frac{\partial \phi_i}{\partial y}\bar{a}_y + g_{iz}\frac{\partial \phi_i}{\partial z}\bar{a}_z \\ \bar{J}_e &= g_{ex}\frac{\partial \phi_e}{\partial x}\bar{a}_x + g_{ey}\frac{\partial \phi_e}{\partial y}\bar{a}_y + g_{ez}\frac{\partial \phi_e}{\partial z}\bar{a}_z \ , \end{aligned} \tag{2.10}$$

where $\bar{a}_x$, $\bar{a}_y$, and $\bar{a}_z$ are unit vectors along the x, y, and z fibre-axes, respectively. More precisely, x represents the fibre (same direction as the fibre), y the sheet (perpendicular direction to the fibre, but along the same tissue sheet), and z the sheet-normal (normal to the tissue sheet plane) fibre-directions. Each fibre-direction has an associated conductivity, g. Thus, g_x, g_y, and g_z are the conductivities along each of the fibre-directions. Moreover, these conductivities can vary between the intracellular and extracellular domains.

The bidomain model imposes the conservation-of-current constraint such that any current leaving one subspace must enter the other one. Thus, we can write $-\nabla \cdot \bar{J}_i = \nabla \cdot \bar{J}_e = I_v$, where I_v is the transmembrane current per unit volume.

Ohm's law, $\bar{J} = \sigma \bar{E}$, defines a current density, $\bar{J}$, as the product of an electrical field, $\bar{E}$, and conductivities, σ. Furthermore, assuming that the magnetic field is not time-varying, through Faraday's law, we can redefine the electrical field in terms of the electrical potential, $\bar{E} = -\nabla\phi$. By applying both laws, we can rewrite the current densities as $\bar{J}_i = -\sigma_i \nabla\phi_i$ and $\bar{J}_e = -\sigma_e \nabla\phi_e$.

Next, by applying the conservation of currents assumption that relates the two domains through the cells' membranes, we can rewrite the density currents in terms of the transmembrane current, $\bar{I}_m$, as follows

$$\begin{aligned} \beta I_m &= \nabla \cdot (\sigma_i \nabla\phi_i) \\ \beta I_m &= -\nabla \cdot (\sigma_e \nabla\phi_e) \ , \end{aligned} \tag{2.11}$$

where β accounts for the membrane surface to volume ratio.

Moreover, by employing the capacitance formulation, the transmembrane current can be written as

$$I_m = C_m \frac{\partial V_m}{\partial t} + I_{ion}(V_m, \eta) \ , \tag{2.12}$$

where C_m is the cell's membrane capacitance, I_{ion} is the ionic current that leaks in and out of the cell due to its mechanisms, and η is the set of variables of the cell model. These ionic currents are governed by non-linear differential equations that will vary for each cell model implementation.

Finally, the bidomain model is

$$\begin{aligned} \nabla \cdot (\sigma_i \nabla\phi_i) &= \beta(C_m \frac{\partial V_m}{\partial t} + I_{ion}(V_m, \eta)) \\ \nabla \cdot (\sigma_e \nabla\phi_e) &= -\beta(C_m \frac{\partial V_m}{\partial t} + I_{ion}(V_m, \eta)) \ . \end{aligned} \tag{2.13}$$

Further information on the bidomain formulation and its boundary conditions can be found in Plonsey and Barr, 2000 and Lines et al., 2003.

2.6.2 Monodomain model

The bidomain model is implemented by a system of partial differential equations accounting for the anisotropy of the intracellular and the extracellular spaces. However, computing the bidomain for human hearts is computationally expensive, requiring hours per beat in high-performance computing environments. The monodomain model simplifies the bidomain formulation that assumes just one domain, as the name indicates. More precisely, the monodomain assumes that the intracellular and extracellular domains are anisotropic but proportional to each other in terms of conductivities, such that

$$\sigma_i = \lambda \sigma_e$$

where λ is a scalar. This assumption allows reducing the bidomain to the following monodomain formulation

$$\nabla \cdot (\sigma_m \nabla V_m) = \beta (C_m \frac{\partial V_m}{\partial t} + I_{ion}(V_m, \eta)) \tag{2.14}$$

were σ_m is the harmonic mean conductivity tensor. The monodomain model is a simplification of the bidomain that has been validated to be equivalent to the bidomain accurately in various cases (Potse et al., 2006).

2.6.3 Eikonal model

The bidomain and monodomain equations provide a reaction-diffusion-based biologically detailed definition of the electrical propagation in the human heart by assuming two electrical subspaces: the intracellular and extracellular spaces. Computation of bidomain-based simulations are costly; thus, phenomenological alternatives such as the Eikonal model have been proposed to study the mechanisms defining the cardiac activation sequence (Colli Franzone et al., 1990; Wallman et al., 2012).

The Eikonal model converts simulation of electrical wave propagation in a cardiac mesh into a shortest-path-finding problem. The eikonal computation cost is significantly lower than its counterparts reaction-diffusion-based solutions; moreover, the eikonal can simulate coarse meshes, whereas the bidomain has convergence

restrictions on the mesh resolution. These characteristics enable the eikonal to produce equivalent activation sequences at several orders of magnitude lower cost than the bidomain or monodomain models (Wallman et al., 2012).

The basic formulation of the Eikonal equation defines the propagation of a wavefront as

$$\begin{cases} \sqrt{\nabla^T d \nabla d} = 1 \ , \quad in \ \Omega \in \mathbb{R}^3 \\ \\ d(0,0,0) = 0 \end{cases} , \tag{2.15}$$

where d is a distance-vector, with the constraint that at the origin of coordinates in 3D space, the solution is zero. Note that the Euclidean distance

$$d(x,y,z) = \sqrt{x^2 + y^2 + z^2} \ , \tag{2.16}$$

is solution to the Eikonal equation.

2.7 Bayesian inference

Bayesian inference is a Bayes' theorem-based method to recover the properties underlying the data. Bayesian inference uses Bayes' rule (Equation 2.17) to update the beliefs on these properties (i.e. hypothesis) from every new data (i.e. data) becoming available.

$$P(H|E) = \frac{P(E|H) \cdot P(H)}{P(E)} \ , \tag{2.17}$$

where H is the (current) 'hypothesis', E is the (new) 'evidence', $P(H)$ is the 'prior probability' (i.e. probability of H without any new 'evidence'), $P(E)$ is the 'marginal likelihood' (i.e. probability of observing the 'evidence'), $P(E|H)$ is the 'likelihood' (i.e. the probability of observing the 'evidence' given that the 'hypothesis' is true), and $P(H|E)$ is the 'posterior probability' (i.e. probability of the 'hypothesis' given the new 'evidence'), which is what the method aims to update.

In the context of parameter estimation, Bayes' rule is usually relabelled as

$$P(\theta|x) = \frac{P(x|\theta) \cdot P(\theta)}{P(x)} \ , \tag{2.18}$$

where θ is a set of parameter values from the parameter space Q and x is a new sample of data from the data space X. Thus, the posterior $P(\theta|x)$ is a probability density in function in Q, that solves the parameter estimation. Now, the probability of observing the new data $P(x)$ can be computed as an integral over the parameter space Q

$$P(x) = \int_Q P(x|\theta)P(\theta)d\theta \ . \tag{2.19}$$

This formulation (Equation 2.19) demonstrates that the probability of observing the new data $P(x)$ is independent of the current parameter values θ being evaluated. Moreover, the prior $P(\theta)$ comes from previous knowledge (e.g. physiological ranges from human experiments). Thus, the Bayesian inference method will focus on solving the likelihood function $P(x|\theta)$.

Bayesian inference can address ill-posed problems such as estimating the ventricular properties in cardiac electrophysiological models from clinical data. In this context, we start from a clinical recording from the patient that represents their electrophysiological behaviour x (e.g. ECG). Next, we build a patient-specific cardiac model using patient imaging data (e.g. CMR). Thus, the problem to solve is: what set of parameter values θ (i.e. parameter-set) in the patient-specific cardiac model simulate the clinical data x. In other words, we want to maximise the posterior function $P(\theta|x)$. As previously mentioned, the main computation to be resolved in Bayesian inference for parameter estimation will be the likelihood $P(x|\theta)$, which, in our modelling and simulation use case, corresponds to the probability of simulating x using the parameter-set θ in our patient-specific model. Thus, this inference process would involve sampling different values for θ and simulating the cardiac model to generate some data x_{sim}. Then, compare the simulated data x_{sim} with the clinically measured data x to estimate the posterior $P(\theta|x)$, knowing that when the theoretical maximum posterior value would correspond to $x_{sim} = x$.

2.7.1 Sequential Monte Carlo approximate Bayesian computation

Sequential Monte Carlo - approximate Bayesian computation (SMC-ABC) is a Bayesian inference method that integrates the ideas behind SMC and ABC algorithms. SMC discovers regions of interest in the parameter space by gradually introducing a given sampling problem's complexity. SMC evolves a population of parameter-sets through a series of intermediary distributions. Each mutation step employs the population to inform the subsequent sampling subproblem. This process results in a high degree of exploration initially, with each intermediate distribution trading some exploration for more exploitation in promising regions of the parameter space (similarly to simulated annealing (Lew et al., 2009)).

ABC targets an approximate posterior for model parameters θ (parameter-set values), $p(\theta|\rho(x, x_{pred}) \leqslant \epsilon)$. Here $\rho(x, x_{pred})$ is a measure of discrepancy between the target data x, and x_{pred}, the predicted data from a parameter-set θ; and, ϵ is the cut-off discrepancy, with $\epsilon = 0$ recovering the true posterior (Equation 2.18). ABC defines a small discrepancy tolerance since the target data is sometimes impossible to reproduce due to noise.

The SMC approach consists of a resampling step that replaces the parameter-sets with high discrepancy and then a mutation step that recovers parameter-set uniqueness. As ABC does not define a likelihood, resampling is achieved simply by splitting the current parameter-sets into a 'to keep' and a 'to replace' group according to their discrepancy value and the current cut-off discrepancy. Then the algorithm replaces the 'to replace' group with an equivalent number of parameter-sets copied from the 'to keep' group selected at random. The copied parameter-sets are then mutated by repeatedly implementing small changes that lower their discrepancies. Finally, the SMC-ABC accepts/rejects these proposed mutations according to the Metropolis-Hastings ratio, as in Markov chain Monte Carlo (Gilks, 2005).

2.8 Convolutional neural networks

Deep learning algorithms are machine learning algorithms that extend traditional neural networks by adding hidden layers strategically connected to promote learning specific pattern types. The most popular deep learning algorithm is the convolutional neural network.

Convolutional neural networks (LeCun et al., 1998) efficiently reduce the number of weights to be trained and outperform traditional fully-connected networks when the data patterns are local (e.g. image and time-series data). Traditional fully-connected neural networks use the whole input for every operation. This causes the number of weights required for each operation to increase exponentially at each additional layer in the network. Moreover, if the patterns are local (e.g. a QRS complex), the algorithm will need to learn to ignore non-local data. Additionally, if the pattern can appear in any position in the input (e.g. a QRS complex can be in any part of the ECG segment), the algorithm would also need to learn this same pattern at every location of input independently.

On the other hand, convolutional neural networks are characterised by implementing 'convolutional layers'. These layers define the local connectivity between layers and recycle all patterns, assuming that a relevant pattern can be relevant anywhere in the input. These local patterns are complemented with convolutional operations to measure the similarity between inputs and learned patterns efficiently. These properties enable convolutional neural networks to focus only on local patterns that can appear anywhere in the input (e.g. QRS morphology in an ECG recording) using fewer weights than fully connected alternatives, thus requiring fewer data to train.

Convolutional layers are formed by kernels that share weighting parameters creating local connectivity between adjacent layers' neurones to exploit spatially local correlations. Therefore, the first convolutional layers will learn a low-level representation of spatial features that summarise the input data. Then, the second convolutional layer will learn the spatial correlations of this summarised representation, and so on. Therefore, at each next convolutional layer, the features

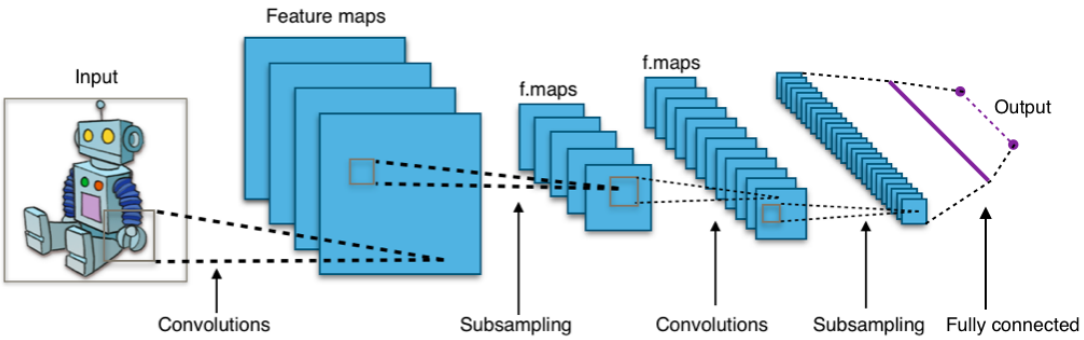

Figure 2.15: Typical convolutional neural network architecture for 2D image recognition. A convolutional neural network composed of an input layer, a first convolutional layer, a first pooling layer (subsampling), a second convolutional layer, a second pooling layer (subsampling), a fully connected layer and another fully connected layer (output). From Commons, 2020, licensed under *Creative Commons Attribution-ShareAlike 4.0 International CC BY-SA 4.0*.

learned will become more abstract. The final layers will learn high-level features, such as the representation of arrhythmic events in ECG recordings.

In summary, convolutional layers can extract features from data with underlying spatial or temporal patterns, such as images and signal data and stacking these layers allows extracting progressively more abstract patterns. Therefore, convolutional neural networks composed a powerful and efficient alternative to traditional fully-connected deep learning architectures when the underlying patterns characterising the data are spatial.

2.8.1 Activation functions in deep neural networks

Deep learning algorithms are composed of large structures of interconnected learning units (neurons). These neurons control the information flow through the network by activating when specific patterns are observed in a data sample. In convolutional neural networks, each neuron takes as input the output from applying one kernel in one location of the input space, thus, requiring only one weight to train. In contrast, in fully connected neural networks, each neuron inputs all the outputs from the previous layer. Thus, while convolutional neural networks will usually have more neurons, the number of weighting parameters will be orders of magnitude smaller than fully-connected neural networks of similar depth (number of layers).

The outputs from these neurons compose the input's feature representation at a given layer of the network. Namely, the neurons determine the expressiveness of each feature learned for a new data sample. Therefore, the training process aims to encode the relevant patterns characterising the data into the convolutional kernels while learning what patterns are relevant in different image locations and encode this information into the neurons. To this end, each neuron implements an activation function that controls the information propagated for a specific kernel in a specific input location.

Some of the popular activation functions for convolutional neural networks are the rectified linear unit (ReLU) $f(x) = max(x, 0)$, the hyperbolic tangent function (Tanh), $f(x) = \tanh(x)$, the logistic sigmoid function (Sigmoid), $f(x) = (1 + e^{-x})^{-1}$. Rectified linear units are the most popular activation function for convolutional neural networks, given their superior time-efficiency in computation cost and convergence speed (Krizhevsky et al., 2012; Szegedy et al., 2015). These advantages originate from activating the neuron only with positive values. This property significantly lowers the computation cost of the following operations and speeds up convergence compared to the Tanh and Sigmoid functions, which are more expensive to compute and more complex to train than the ReLU activation function.

2.8.2 Regularisation for deep learning algorithms

Training deep learning algorithms can cause overfitting or underfitting. Overfitting implies an over-specialisation of the training data. In contrast, underfitting implies that the algorithm could not learn the basic patterns describing these data. In either case, the deep learning algorithm will fail to generalise to new data. Overfitting is caused by training a deep learning algorithm with greater representation capacity than required to represent the data's relevant patterns. In contrast, the opposite scenario causes underfitting. In deep learning practice, the gold standard strategy is to adopt a deep learning algorithm with a more significant learning capacity than intuitively required for the task and apply regularisation to limit the degree of detail to which the data patterns are represented (Goodfellow et al., 2016).

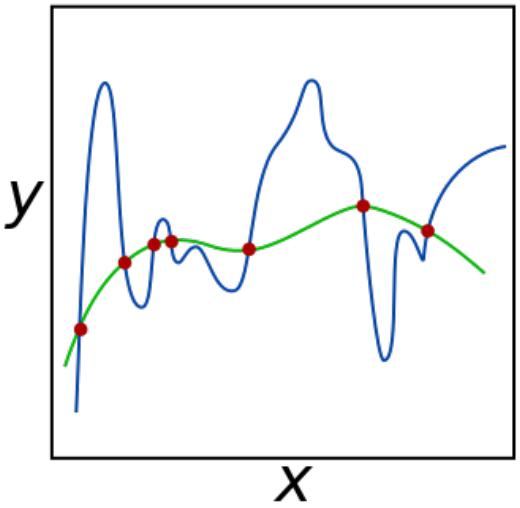

Figure 2.16: Regularisation example. The red dots represent the train data in a regression problem, the blue line is a solution without regularisation, and the green line is the solution with weight decay regularization. From Commons, 2020, licensed under *Creative Commons Attribution 4.0 International CC BY 4.0*.

Regularisation techniques aim to reduce generalisation error (Figure 2.16) while preserving training accuracy to prevent overfitting on the training data without causing underfitting (Goodfellow et al., 2016). Hereafter the regularisation strategies considered in this thesis.

Dropout

The most popular regularisation strategy for deep learning algorithms is 'dropout'. Dropout is an efficient and effective regularisation strategy Srivastava et al., 2014; Szegedy et al., 2015. Using dropout to train a deep learning algorithm provides an approximation to implementing the bagging ensemble method over numerous deep learning algorithms while preserving inexpensive computational costs. The main intuition behind dropout is to employ only a random subset of the network each time a new instance is fed to the algorithm. This approach takes only one parameter, which corresponds to the probability by which a neurone is 'dropped' (excluded).

Early stopping

Deep learning algorithms are trained iteratively to approach relevant local optima in the solutions space, which can successfully represent the target problem. These

iterations are denoted as epochs. Correctly establishing the number of epochs is important to prevent the deep learning algorithm from overfitting to the training data while avoiding useless computation by detecting if the algorithm has converged. Early stopping is a regularisation strategy that halts the training process in a deep learning algorithm to prevent over-specialising the training data.

More precisely, the number of epochs by which a deep learning algorithm may reach convergence will usually be correlated to its architecture, data, optimisation method, and internal parameters, such as the learning rate and the regularisation strategies employed. Due to its importance and complexity, deep learning algorithms are usually trained using an early stopping strategy that implements stopping criteria based on the algorithm's training performance. Defining these stopping criteria allows the training process to be halted when the training accuracy denotes overfitting. In other words, overfitting caused by too much training is characterised by increasing the training accuracy and lowering the validation accuracy.

2.8.3 Training deep learning algorithms

Cross-validation

Training deep learning algorithms is typically performed employing strategies that iterate the training process to demonstrate the algorithm's robustness before proceeding to the testing stage. The most popular training strategy is cross-validation. Cross-validation consists of splitting the training database into several random folds and at each evaluation stage during development, to train the algorithm across all the folds except the validation fold, which assesses the method's performance. This process is repeated until all folds have been used as validation. The average performance from these evaluations informs about the algorithm's expected performance to new data from the same distribution since all the data in a database is considered to belong to a single distribution. Overall, cross-validation offers an exhaustive evaluation of the algorithm while avoiding being biased to a specific choice of training and validation split of the data. Moreover, extensions

of the strategy, such as leave-one-out and nested cross-validation, make this the state-of-the-art strategy for evaluating deep learning algorithms.

Leave-one-out cross-validation extends the previous strategy by emphasising the way the data is divided into folds. In standard cross-validation, this division is performed at random, implying that the data distribution is the same in all folds. On the other hand, by dividing the data according to its distribution, for example, per-patient or per cardiac condition, it is possible to extract more robust conclusions on the algorithm's generalisation power. The best possible division would be dividing the data per database to guarantee different distributions for the same task.

Another extension to this method is nested cross-validation. Traditional cross-validation allows learning the weight values in the deep learning algorithm. However, the architecture and other hyperparameters are static to prevent overfitting. Nested cross-validation adds a validation level to the strategy to adjust these hyperparameter values without overfitting the training data. More precisely, this extension starts by dividing the data into folds and then singles out one fold as a test fold. The remaining folds are renamed as validation folds. Next, fixates the hyperparameters and performs standard cross-validation using validation folds while completely ignoring the test fold. The performance of this cross-validation is the expected performance of the current hyperparameter values. Then, this process is repeated for all possible hyperparameter values. Finally, the algorithm uses the best performing hyperparameter configuration and all the validation folds to train its weights. This trained algorithm is then evaluated on the test fold. Then the whole process is repeated by selecting the next fold as a test until all folds have been used as a test. The average test performance can then be reported. Nested cross-validation keeps the benefits of not being bound to a single split of the data while automatically tuning the hyperparameters in the deep learning algorithm. This extension of cross-validation can be combined with others, such as leave-one-out, to simultaneously exploit the benefits from various extensions.

Batch normalisation

At a lower scale, deep learning algorithms typically require large amounts of data to train. Thus, the data are usually split into randomised batches (or mini-batches) to speed up the training process by giving more evidence to the algorithm to update its weights. This well-accepted strategy allows training algorithms on multiclass data efficiently by representing all classes in each mini-batch. However, the differences in the distribution of the input data in each batch can cause an 'internal covariate shift'. Namely, the algorithm's updates behave as if chasing a moving target, demanding longer convergence times. Batch normalisation consists of standardising the inputs to a layer for each mini-batch. This standardisation accelerates the training process's convergence by enabling the algorithm to have a steadier progression through the gradients' space.

2.8.4 Optimisation algorithms for deep learning

Deep learning methods implement stochastic gradient-based algorithms to optimise the error loss after each batch is processed. These stochastic training methods perform the optimisation in mini-batches, which are partitions of the entire training-train data. Thus, deep learning algorithms compute the gradients and carry out the weight corrections per minibatch, rather than a single update per epoch. The optimal magnitude of these corrections is challenging to match; however, if these corrections are small enough, the algorithm will eventually converge. This magnitude is referred to as the learning rate and is commonly set to low values.

Recent optimisation algorithms, such as adaptive gradient algorithm (AdaGrad) (Duchi et al., 2011), implement adaptive learning rate techniques. Algorithms with adaptive learning rates can start from large learning rates to accelerate the training procedure and decrease them concerning the gradient once the algorithm is ready to converge. Other popular stochastic gradient-based optimisation algorithms are the following:

- Stochastic gradient descent (SDG) (Bottou, 1998), which, despite having a fixed learning rate, is the most widely used optimisation algorithm in deep learning applications (Goodfellow et al., 2016).
- Root mean square propagation (RMSProp) (Tieleman and Hinton, 2012) (i.e. extension of AdaGrad).
- AdaDelta (Zeiler, 2012) (i.e. extension of AdaGrad).
- Adaptive momentum (Adam) (Kingma and Ba, 2014a).

These optimisation strategies are usually paired with a weight initialisation strategy that suits them (in most cases), and then the algorithm is ready to start training.

Setting proper initial weights is key for successfully training deep learning algorithms. Finding an optimal weight initialisation strategy is challenging due to its dependence on the data, target task, deep learning architecture, regularisation strategies, and hyperparameter values. However, in practice, several initialisations will enable a deep learning algorithm to train adequately (Goodfellow et al., 2016).

2.8.5 Deep learning strategies for ECG delineation

Data augmentation for ECG signals

The best strategy for enhancing machine learning algorithms' generalisation is to increment the training data available (Goodfellow et al., 2016). Data augmentation strategies artificially augment the data available for training the deep learning algorithm. Data augmentation is a common practice in deep learning applications as it omits the need for labelling new data. The most common strategies for data augmentation consist of applying transformations or distortions to preexisting training samples. However, these transformations should be carefully designed. For example, the preservation of temporal dependencies invalidates most typical data augmentation transformations in time-series data. In other words, transformations that would affect events' order would also change the data's interpretation.

Moreover, biomedical signals, such as ECG recordings, have further rules to comply with non-biomedical time-series signals. More precisely, augmenting ECG

recordings requires that the signal's physiological soundness is preserved after the transformation. For example, a uniformly upsampled or downsampled ECG signal becomes non-physiological because heart-rate changes do not modify all waves in the signal evenly. On the other hand, strategies such as windowing (with overlap), noise-contamination, amplitude scaling, and baseline shifting would enable artificially augmenting the ECG signals to train a deep learning algorithm without physiologically corrupting them.

Pooling layers

In segmentation and delineation applications, invariance to slight changes in feature location is undesirable. Pooling layers, however, explicitly make the deep learning algorithm invariant to small changes in feature position. Thus, these layers are rarely implemented for segmentation or delineation oriented deep learning architectures. The absence of pooling layers increases the number of weights per layer in the convolutional neural network while reducing the scale of the features learnt. Thus, architectures without pooling (e.g. segmentation and delineation deep learning architectures) layers require additional convolutional layers to learn large-scale patterns from the data.

Alles Gescheite ist schon gedacht worden.
Man muss nur versuchen, es noch einmal zu denken.

All intelligent thoughts have already been thought;
what is necessary is only to try to think them again.

— Johann Wolfgang von Goethe (Goethe, 1829)

3 Literature review

3.1 Inference of the ventricular activation properties

The parameter fitting process in biophysically detailed cardiac models from non-invasive test recordings continues to be too computationally expensive for routine use in clinical applications. This motivates the need for adopting computationally efficient models. Phenomenological formulations, such as the Eikonal equations, represent a computationally efficient alternative to biophysically detailed reaction-diffusion models (e.g. bidomain and monodomain) (Wallman et al., 2012). Phenomenological models are appropriate for various clinical applications that focus on tissue-level phenomena (Chinchapatnam et al., 2008; Pernod et al., 2011).

In the literature of parameter inference for cardiac models, there are families of methods that use phenomenological electrophysiological models of propagation: traditional optimisation strategies (Dhamala et al., 2017; Molléro et al., 2018; Kunisch et al., 2019), regression strategies (Zettinig et al., 2014; Giffard-Roisin et al., 2017b), and sampling strategies (Konukoglu et al., 2011; Wallman et al., 2014; Grandits et al., 2020).

Traditional optimisation parameter inference methods can solve parts of the estimation problem for which a formulation can be defined. For example, Dhamala

et al., 2017 presented an optimisation framework for estimating the parameters of the Aliev-Panfilov model (Aliev and Panfilov, 1996). The optimisation strategy in Dhamala et al., 2017 considers a non-uniform resolution of the parameter space. Namely, heterogeneous regions of the tissue will be automatically given higher resolution compared to homogeneous ones. This spatial-adaptive resolution required them to implement the mesh as a multiscale hierarchical representation of the spatial domain. They validated the proposed methods on simulated data and recordings from seven patients diagnosed with myocardial infarction. Molléro et al., 2018 presented a 0D/3D multi-fidelity approach for personalising 3D cardiac models. Their method derives a simplified 0D version of the 3D one for performing the expensive computations. Next, they map the results of the 0D model back to the 3D one. This strategy enabled their approach to be orders of magnitude faster than just using the 3D model. Most recently, Kunisch et al., 2019 proposed an optimisation strategy for recovering the root nodes on the left endocardial surface from an epicardial activation time map using the Eikonal model. Despite showing promise in specific cases, optimisation strategies require problem specific formulations. They have not been applied to the full inverse problem considered in this thesis, namely, estimating the conduction speeds and root nodes simultaneously.

With the rise in popularity of machine learning approaches over the last few decades, some approaches that aim to learn a relationship between the data and the parameter values in cardiac models have been proposed. Zettinig et al., 2014 and Giffard-Roisin et al., 2017b proposed regression methods for the estimation of cardiac activation properties in phenomenological electrophysiological models. Zettinig et al., 2014 presented a cardiac electrical diffusivity estimation approach from ECG recordings using a phenomenological cardiac model based on a Lattice-Boltzmann formulation. They generated 9500 simulations to learn a multivariate polynomial regression (Zettinig et al., 2013). This enabled them to quantify the correlation between the uncertainty of their predictions and the ECG biomarkers. Their validation considered ECGs from nineteen dilated cardiomyopathy patients. Similarly, Giffard-Roisin et al., 2017b presented a cardiac model personalisation

study using a phenomenological model that extended the Mitchell–Schaeffer model (Mitchell and Schaeffer, 2003), with a current dipole formulation. They used a kernel ridge regression to learn the activation locations and the global ventricular conduction velocity from QRS biomarkers from body surface potential maps. Their validation considered recordings from a patient with premature ventricular contractions and five non-ischaemic implanted cardiac resynchronisation therapy patients. Regression-based inference methods are more efficient than sampling and optimisation-based inference strategies. However, regression-based methods usually lack the representation power to encapture the whole parameter space representing the human population. These methods, instead, are better suited to address specific cases where prior knowledge on the patient is available, as in Giffard-Roisin et al., 2017b . Admittedly, an ensemble of case-specific regression-based inference methods could aim to represent the whole human population similarly to sampling and optimisation-based methods. However, there is currently not enough data available to build such an ensemble-regression method. Consequently, I believe that optimisation and sampling-based strategies, despite their higher computational cost, are currently better suited to be applied to undiagnosed patients for building cardiac digital twins in the clinic.

Most sampling strategies in the literature of parameter inference for personalisation of cardiac models implemented Bayesian-based strategies combined with phenomenological cardiac models. For example, Konukoglu et al., 2011 presented a Bayesian inference-based method for the parameter estimation of the Eikonal model solved using the fast marching method (Konukoglu et al., 2007). They proposed an alternative formulation of the problem using spectral representations based on polynomial chaos expansions to work around the bottleneck of sampling with a Markov Chain Monte Carlo. Similarly, Wallman et al., 2014 presented a Bayesian inference-based method for the parameter estimation of their graph-based model (Wallman et al., 2012). Moreover, they validated the inference method using biophysically detailed simulations to know the ground-truth. They also proposed a novel strategy for selecting the number of measurements and their locations

to maximise the information gain during the parameter inference process. Their (heuristic) graph-based Eikonal model (Wallman et al., 2012) interpreted the cardiac mesh as a connected graph and computed the activation times as a shortest-path finding task using Dijkstra's algorithm (Dijkstra, 1959). The Eikonal model solved with Dijkstra's algorithm produced equivalent simulations to using the fast marching method and also than the monodomain model (Wallman et al., 2012). Giffard-Roisin et al., 2017a proposed a sparse Bayesian kernel-based regression method for estimating root nodes (pacing sites) in their phenomenological cardiac model, which had been previously presented in Giffard-Roisin et al., 2017b. They validated their approach using data from a patient artificially paced at two known ventricular locations. Nevertheless, non-Bayesian sampling-based strategies have also been considered. Grandits et al., 2020 presented a sampling-based regularisation approach to recover the conduction speeds and times of first activation in the Eikonal model using noisy epicardial activation time maps. Their evaluation considered noise-contaminated Eikonal-simulated epicardial activation maps in one biventricular geometry. Overall, sampling-based inference methodologies combined with fast models of propagation show greater flexibility than optimisation and regression-based alternatives. These methods can easily be generalised to new cases by increasing the ranges of the parameter-space (e.g. allowing the exploration of slow conductance values due to fibrosis) or including new parameters to be estimated (e.g. estimating different conduction speeds in scar regions).

No previous works aimed at recovering all human ventricular activation properties (i.e. conduction speeds and root nodes) during sinus rhythm. Most previous studies focused on disease-specific conditions affecting the Purkinje network (e.g. artificial pacemakers or premature ventricular contractions). These conditions reduce the number of root nodes to estimate from about seven (Cardone-Noott et al., 2016) to one or two (Giffard-Roisin et al., 2017a). Moreover, the number of root nodes is known in these disease cases. Reducing and making known the number of the root nodes significantly eases the inference task compared to sinus-rhythm conditions,

making them less relevant for clinical applications with no prior knowledge of each patient's Purkinje network structure.

3.2 ECG delineation

As mentioned in Subsection 2.4.2, delineation is identifying fiducials of the waveforms composing the ECG signals. This process is the core step in traditional ECG analysis since it allows the computation of waveform-related features. Due to its repetitive nature, several studies have suggested automatic delineators for ECG recordings. Automatic ECG delineators can be multilead or single-lead. Multilead approaches use multiple leads for identifying the waveform boundaries, whereas single-lead methods only process the information of one lead. Researchers have extensively compared and explored both alternatives; however, which one is better depends on the biomarkers' use and purpose.

More precisely, there is no standard on delineating the waveforms composing the ECG signals or how to combine information from several leads when delineating. Therefore, the ground-truth labels in reference databases will include criteria-dependent variability. In other words, the placement of some marks may differ depending on the criteria adopted by the expert. Moreover, this variability increases in multilead databases due to the additional criteria required to merge the information from multiple leads.

ECG delineation systems were already being implemented during the late 70's (Thomas et al., 1979). At that time, according to (Thomas et al., 1979), a review of computer-based solutions for ECG analysis, QRS delineation was performed by threshold-based methods, which implemented several rules to account for specific cases and, this way, refine the output of the algorithm. At that time, the MIT-BIH Arrhythmia database had already been collected.

The data composing the MIT-BIH Arrhythmia database were collected during the late 70's (Moody and Mark, 2001). The purpose of collecting these signals was to provide long-term ECG recordings to develop automated methods for identifying arrhythmias. In 1980, this database was made public and became

the benchmark dataset for comparing performances in several ECG analysis tasks, such as arrhythmia detection, beat classification, and R-wave's peak detection. The existing methods at that time had issues with discontinuous behaviours in some waveforms within the beat, until Nygårds and Sörnmo, 1983 proposed a delineation method that represented each QRS complex as a single positive pulse. This method was robust to noise and gradual morphology changes, whereas the main delineation strategy still was based on thresholds.

In 1985, the researchers in the European project Common Standards for Quantitative Electrocardiography (CSE) proposed a set of recommendations to update the standards in quantitative electrocardiography (The CSE Working Party, 1985) and the dataset employed to design these rules, namely, the CSE database (Willems et al., 1985). Five years later, Laguna et al., 1990 presented a rule-based QT interval delineator for ambulatory ECG recordings. Their approach followed the classic strategy in rule-based delineation systems. More precisely, firstly, they filter the signals to reduce noise. Secondly, they detect QRS complexes by searching for R-wave's peaks. Next, the method navigates through the signal using the R-wave's peaks as a reference to find the QRS onset and the T-wave's peak. Finally, this method locates the T-wave's offset using the R-wave's peak and T-wave's peak to constrain where this mark can be. Laguna et al., 1994 proposed a multilead generalisation of their previous work; moreover, they validated their multilead delineator on the CSE database. They also tested the multilead approach on records from the MIT-BIH Arrhythmia database to check the single-lead mode of the algorithm.

Cuiwei Li et al., 1995 proposed a delineator system that employed the wavelet transform to represent the signals before locating the fiducials. The validation of this wavelet transform-based approach was conducted in the MIT-BIH Arrhythmia database.

Laguna et al., 1997 published the QT database, a delineation-specific database. This database includes 105 fifteen-minute ECG recordings with two leads sampled at 250 Hz. Most of these recordings were not specifically collected for this database

but instead came from six databases, namely, the MIT-BIH Arrhythmia database, the MIT-BIH ST Change database (Albrecht, 1983), MIT-BIH Supraventricular Arrhythmia database (Greenwald and Patil, 1990), the MIT-BIH Normal Sinus Rhythm database (Goldberger et al., 2000), the European ST-T database (Taddei et al., 1992), and the MIT-BIH Long-Term ECG database (Goldberger et al., 2000). Additionally, the QT database authors collected and included ambulatory ECG data from patients who suffered sudden cardiac death while recording together with records from control patients that matched the first group in age and sex. All these databases can be found in PhysioNet, a research resource that offers access to numerous ECG datasets. The QT database includes delineation annotations for 3622 beats, a subset of which was annotated a second time independently to enable intra-annotator analysis. Moreover, all annotations were audited to correct potential inconsistencies (Laguna et al., 1997; Goldberger et al., 2000). The QT database rapidly became the standard benchmarking database for validating ECG delineation systems.

Martínez et al., 2004 proposed a wavelet transform-based delineator that extended Cuiwei Li et al., 1995, and quickly became the state-of-the-art in ECG delineation. More precisely, Martínez et al., 2004 considered the discrete wavelet transform with a quadratic spline as the prototype to represent the raw ECG data. Due to the frequential multiresolution properties of the wavelet transform, this delineation method can omit the filtering step in the typical ECG analysis pipeline. A set of carefully designed rules that correlate with how the waveforms in the beat are mapped into the wavelet transform representation compose this algorithm's decision-making engine. The algorithm has four steps: Firstly, this delineator locates the QRS complexes using thresholds on the wavelet transform scales, representing this waveform. Secondly, the delineator searches for the maximum slopes surrounding the QRS complex their algorithm to find the onset and offset of this complex, taking advantage of the derivative nature of these wavelets. Next, it uses the QRS marks to determine the time window for searching for the T-wave and delineates the T-wave using a similar strategy to the QRS complex in different wavelet transform scales.

Finally, the delineator locates and delineates the P-wave within the target beat by using a similar strategy to the T-wave. All the wave searches conducted by this algorithm use reference plausible morphologies that the waveforms may have adopted in the signal and also account for the not-found possibility. The databases employed to validate this method were the MIT-BIH Arrhythmia database, the QT database, the European ST-T database, and the CSE database. However, only the QT database and the CSE database served to validate the delineation part of the algorithm. The validation on the QT database was performed by predicting the fiducials independently for each lead. Then, for each point in time they chose the most accurate prediction. This validation strategy aimed to compensate for the fact that the database's annotators had access to all leads in the recordings when annotating the signals but only produced one prediction for each time instance and did not explain from which lead each mark was derived. Their validation reported state-of-the-art performances for QRS detection and outperformed other existing ECG delineation approaches. They obtained similar performances to Laguna et al., 1994 in on the delineation of the P-wave's peak, the P-wave's end, and both boundaries of the QRS complex (i.e. onset and offset). However, the wavelet transform-delineator (Martínez et al., 2004) achieved significantly higher performance than the previously presented multilead delineator when determining the P-wave's onset, the T-wave's peak, and the T-wave's offset. They reported the delineation error's signed mean ± its standard deviation in milliseconds (ms) for each complex: P-wave's onset — 2±14.8; P-wave's peak — 3.6±13.2; P-wave's offset — 1.9±12.8; QRS's onset — 4.6±7.7; QRS's offset — 0.8±8.7; T-wave's peak — 0.2±13.9; and T-wave's offset — -1.6±18.1. To date, this wavelet transform method continues to be considered a competitive ECG delineation method.

Hughes et al., 2004 proposed a machine learning-based ECG delineation system which focused on the analysis of the QT interval. This method implemented both hidden Markov and hidden semi-Markov models and compared their performances. However, the method's assessment is not directly comparable with other works since they employed a private dataset to validate their algorithms and the fact that they

report 5-fold cross-validation errors instead of clean test errors. Nevertheless, they can report promising results for the hidden semi-Markov model combined with a wavelet transform representation of the signals; this approach reported an absolute mean error of 3 ms and 12 ms for the Q-wave and T-offset, respectively.

Sun et al., 2005 presented a multiscale morphological derivative transform-based delineator, which combined the multiscale morphological transformation with rules to identify the characteristic points in the ECG signals, and was tested on recordings from the MIT-BIH Arrhythmia database and the QT database to assess wave-detection and wave-delineation, respectively. They reported the following delineation error's signed mean ± standard deviation (in milliseconds): P-wave's onset — 9±9.4; P-wave's offset — 12.8±13.2; QRS's onset — 3.5±6.1; QRS's offset — 2.4±10.3; T-wave's onset — 7.9±15.8; and T-wave's offset — 8.3±12.4. These performances were similar to those previously reported by Martínez et al., 2004.

Computing in Cardiology is an annual conference about computational methods for cardiovascular research. This conference hosts different challenges organised by Physionet in each edition. Interestingly, the challenge for 2006 (Moody et al., 2006) was QT interval measurement on the Physikalisch-Technische Bundesanstalt (PTB) database (Bousseljot et al., 2009). Almeida et al., 2006 presented a multilead extension of Martínez et al., 2004 that considered the vectorcardiogram (FRANK, 1956) 3D representation of the 12-lead ECG. Nevertheless, the 2006 Physionet challenge winner was a threshold-based delineator (Hayn et al., 2006) (see Schreier et al., 2003 for the description of this approach). The authors modified the strategy for combining predictions in multilead ECG recordings concerning their previous description of the method in Schreier et al., 2003 to achieve better results on the PTB database.

The lack of standardised automatic systems for delineating ECG recordings forced higher-level studies to implement their delineation algorithms. For example, Singh and Gupta, 2008 presented a time derivative and adaptive thresholding approach for ECG delineation to investigate the relevance of ECG fiducials as biometric footprints to identify different people. Consequently, more sophisticated

delineation methods continued to be proposed. Ghaffari et al., 2009 proposed a wavelet transform and adaptive thresholding-based ECG delineation strategy. Their evaluation considered the QT database. However, they replaced the fiducials in the database that corresponded to challenging beats to delineate by new marks produced from clinicians.

Vítek et al., 2009 presented a continuous wavelet transform-based delineation method, which soon after they extended in Vitek et al., 2010 to delineate the P-wave. They validated their methods on the CSE database. Similarly, Almeida et al., 2009 proposed an extension from their vectorcardiogram multilead delineator (Almeida et al., 2006), and soon after, they proposed a variant of this same method for P-wave delineation in Almeida et al., 2010.

Martínez et al., 2010 presented a delineation system based on the phasor transform (for details on the phasor transform, refer to Proakis and Manolakis, 1992). Analogously to Ghaffari et al., 2009, Martínez et al., 2010 considered that the QT database required several labelling corrections. However, when considering the original annotations in the QT database, they reported comparable performances to previously published ECG delineation methods (Martínez et al., 2004; Sun et al., 2005).

At this point in the history of ECG delineation approaches, achieving substantial improvements became quite challenging since the data available had already been intensely exploited, and no new delineation ECG databases were being collected. Therefore, new ECG delineation approaches aimed to stand out by becoming problem-specific instead of generic delineators. For example, Di Marco and Chiari, 2011 presented a wavelet transform-based delineator designed explicitly for being implemented using 32-bit integer linear algebra processing to reduce the method's computational cost. Vázquez-Seisdedos et al., 2011 presented a specific delineation method for T-wave's offset detection in noisy conditions. Homaeinezhad et al., 2014 presented a delineation system that employed pattern recognition techniques to cluster RR intervals and then delineated the P and T waves accordingly to each cluster.

Research in ECG delineation has slowed down considerably as no new public datasets exist since the QT database (Laguna et al., 1997), during the last few years, some interesting approaches have been presented in this field. Bote et al., 2018 proposed a low complexity thresholding-based delineation system specially designed to be implemented in wearable devices. Bote et al., 2018 validated their delineation approach on the MIT-BIH Arrhythmia database and the QT database. However, they adjusted the model's parameters using the QT database without keeping a test dataset to evaluate their system's generalisation. The QRS delineation method presented in this thesis was published in Camps et al., 2018. In this work I demonstrated the generalisation power of deep learning-based delineation and compared to the state-of-the-art methods. Subsequent studies in the literature were unable to demonstrate better QRS delineation results than my deep learning-based delineation methodology. For example, Jimenez-Perez et al., 2019 and Jimenez-Perez et al., 2021 adapted the U-Net architecture (Ronneberger et al., 2015) to work with ECG signals, whereas Peimankar and Puthusserypady, 2021 proposed a nearly identical architecture and training strategy to the ones used in Camps et al., 2018 .

An automatic standard ECG delineation method capable of generalising to new databases would enable extracting clinically relevant biomarkers from large unlabelled ECG datasets, such as the UK Biobank (Sudlow et al., 2015), and reducing the variability in these ECG biomarkers produced by choice of the delineation strategy.

3.3 Machine learning for ECG

Over the last decade, the increase in the availability of digital health records, including various medical data modalities for large and diverse population samples, such as the Biobank database (Sudlow et al., 2015), has enabled machine learning in medical research. Moreover, the increase of computational power and repurpose of graphics processing units for efficiently computing convolutions (Krizhevsky et al., 2012), have enabled state-of-the-art machine learning strategies, such as deep learning, to be considered in biomedical problems (Bycroft et al., 2018).

Recent studies have aimed at exploiting the synergy between the predictive power of machine learning algorithms and clinically relevant biomarkers extracted from ECG recordings. For example, Al-Zaiti et al., 2020 applied logistic regression, gradient boosting machines and artificial neural networks to predict acute coronary syndrome using biomarkers extracted from 12-lead ECG recordings. Their machine learning algorithms were able to outperform expert clinicians and commercial interpretation software on their benchmark comparison.

Deep learning is the set of machine learning strategies characterised by having a hierarchical architecture of non-linear layers. However, the property of these methods, which has raised their popularity, is the capacity of learning the representation of the data by themselves. Deep learning has recently become state-of-the-art in several biomedical problems (Mamoshina et al., 2016). Moreover, these techniques have shown great success in signal data (Faust et al., 2018).

3.3.1 Deep learning for ECG

Deep learning methods (Section 2.8) learn multi-level representations of raw data (LeCun et al., 2015). They can learn feature extractions that can easily handle multimodal data, missing information and high dimensional feature's spaces (LeCun et al., 2015; Goodfellow et al., 2016), omitting manual feature engineering. Furthermore, deep learning algorithms can outperform traditional machine learning algorithms when enough data to represent the complexity of a target problem are provided adequately. Concretely, deep learning algorithms have recently exhibited a breakthrough in several complex problems, even outperforming humans in complex tasks such as image classification (He et al., 2015) and playing games (Silver et al., 2016). Some other remarkable applications which are currently possible to implement with deep learning techniques are: face detection (Taigman et al., 2014), image generation (Oord et al., 2016a), video-frame prediction (Finn et al., 2016), speech recognition (Oord et al., 2016b) and audio generation (Oord et al., 2016b) and question answering (Weston, 2016; Graves et al., 2016).

Deep learning-based methods are the state-of-the-art ECG analysis (e.g. arrhythmia detection) (Rajpurkar et al., 2017; Mincholé et al., 2019a). Thanks to the capacity of performing automatic feature extraction from the raw signals, deep learning-based methods can avoid relying on delineation systems that have not been updated for the last decade. Furthermore, deep learning approaches have shown outstanding performances, and are more straightforward to adapt to unconsidered cases given the availability of sufficient data. On the other hand, these approaches have been criticised for being black-boxes. In other words, experts struggle to understand the reasoning process behind the decision-making engine of deep learning models; and, thus, do not trust critical decisions made by deep learning-based systems. Admittedly, critical decision-making systems should provide supportive arguments to aid the practitioner to understand what underlying evidence is being considered when performing a decision. The design of modular ECG analysis systems would allow integrating the strengths of deep learning into the feature extraction steps without losing interpretability about the biomarkers considered for the decision making ones.

Arrhythmia and heartbeat detection and classification are some of the most explored problems using deep learning models on ECG recordings (Acharya et al., 2017a; Acharya et al., 2017b; Kiranyaz et al., 2016; Majumdar and Ward, 2017; Luo et al., 2017). However, several other ECG-related problems are being approached using deep learning techniques, for example, coronary artery disease diagnosis (Tan et al., 2018), sleep apnea detection (Cheng et al., 2017), identifying shockable and non-shockable arrhythmias (Acharya et al., 2018), mortality prediction in COVID-19 patients (Bleijendaal et al., 2021a), and ECG feature discovery for specific disease types (Leur et al., 2021). Throughout, deep learning methods continued to demonstrate their superior predictive power compared to human experts in the ECG analysis domain. Bleijendaal et al., 2021b demonstrated deep learning methods outperforming expert cardiologists in the diagnosis of phospholamban *p.Arg14del* cardiomyopathy using ECG recordings.

Overall, deep learning-based methods have shown promising results for ECG analysis (Mincholé et al., 2019a; Ebrahimi et al., 2020; Somani et al., 2021). However, the scarcity of appropriately annotated ECG datasets prevents testing the generalisation capability of these methodologies. For example, (Acharya et al., 2017b; Kiranyaz et al., 2016; Majumdar and Ward, 2017; Luo et al., 2017) were validated using only the MIT-BIH Arrhythmia database (Moody and Mark, 2001) from Physionet (Goldberger et al., 2000). This scarcity of appropriately annotated data has also limited the performance obtained in deep learning-based delineation tasks to that reported in Camps et al., 2018, as observed in subsequent more sophisticated works, such as Jimenez-Perez et al., 2019.

It can scarcely be denied that the supreme goal of all theory is to make the irreducible basic elements as simple and as few as possible without having to surrender the adequate representation of a single datum of experience.

— Albert Einstein (Calaprice, 2010)

4 Methods

4.1 Introduction

This chapter investigates new computational techniques for the efficient quantification of subject-specific ventricular activation properties from non-invasive electrocardiographic data using CMR-based modelling and simulation. I present an inference method combined with fast Eikonal-based simulations and CMR-based torso-biventricular anatomical models to determine the accuracy in estimating activation properties (such as endocardial and myocardial conduction speeds and the location of the Purkinje endocardial root nodes) from activation time maps (as obtained through ECGI). I conducted the simultaneous inference of endocardial and myocardial conduction speeds and the root nodes, as these properties determine the activation sequence in the ventricles (Cardone-Noott et al., 2016). To address the challenges associated with inferring root node locations and speeds simultaneously, I propose a novel inference method based on the sequential Monte Carlo approximate Bayesian computation (SMC-ABC) algorithm (Sisson et al., 2007; Drovandi and Pettitt, 2011). This chapter describes the methodologies for this study, the cohort of virtual subjects to evaluate them, and the metrics for analysing the results from this evaluation.

4.2 Overview

Figure 4.1 presents a diagram of my inference method, including input, output, and the iterative process to infer the human ventricular activation properties from electrocardiographic recordings using CMR-based modelling and simulation. The input data (top-right – grey shaded area) include 'given data' (torso-biventricular mesh and fibre orientations) as well as the 'target data' (e.g. 12-lead ECG recordings). The iterative process (centre and bottom-left – blue shaded area) aims to find a population of models (with different parameter-sets but the same equations and subject-specific anatomical model) that yield simulations in agreement with the 'target data' (bottom-right – grey shaded area). Each model in this population implements one set of parameter values (hereafter referred to as parameter-set), namely, a value for the locations of the root nodes (earliest activation sites), the endocardial layer's isotropic speed, and the fibre, sheet (transmural), and sheet-normal directed orthogonal conduction speeds.

The inference process (Figure 4.1) starts with the Latin hypercube sampling (Iman and Conover, 1980) (top-left), generating a population of 512 parameter-sets for conduction speeds and root nodes with a uniform prior. These parameter-sets are combined with the subject's CMR-based biventricular mesh and the Eikonal formulation to produce an initial population of 512 cardiac electrophysiological models. The iterative process for the inference (centre and bottom-left – blue shaded area) follows five steps (from the top-left on the directions of arrows in Figure 4.1):

1. simulate the activation time maps from the new population of models;

2. calculate the predicted ECGs (unless when working directly on activation maps);

3. compute the discrepancies (distance) between each prediction and the target data (e.g. ECG);

4. evaluate the stopping criteria;

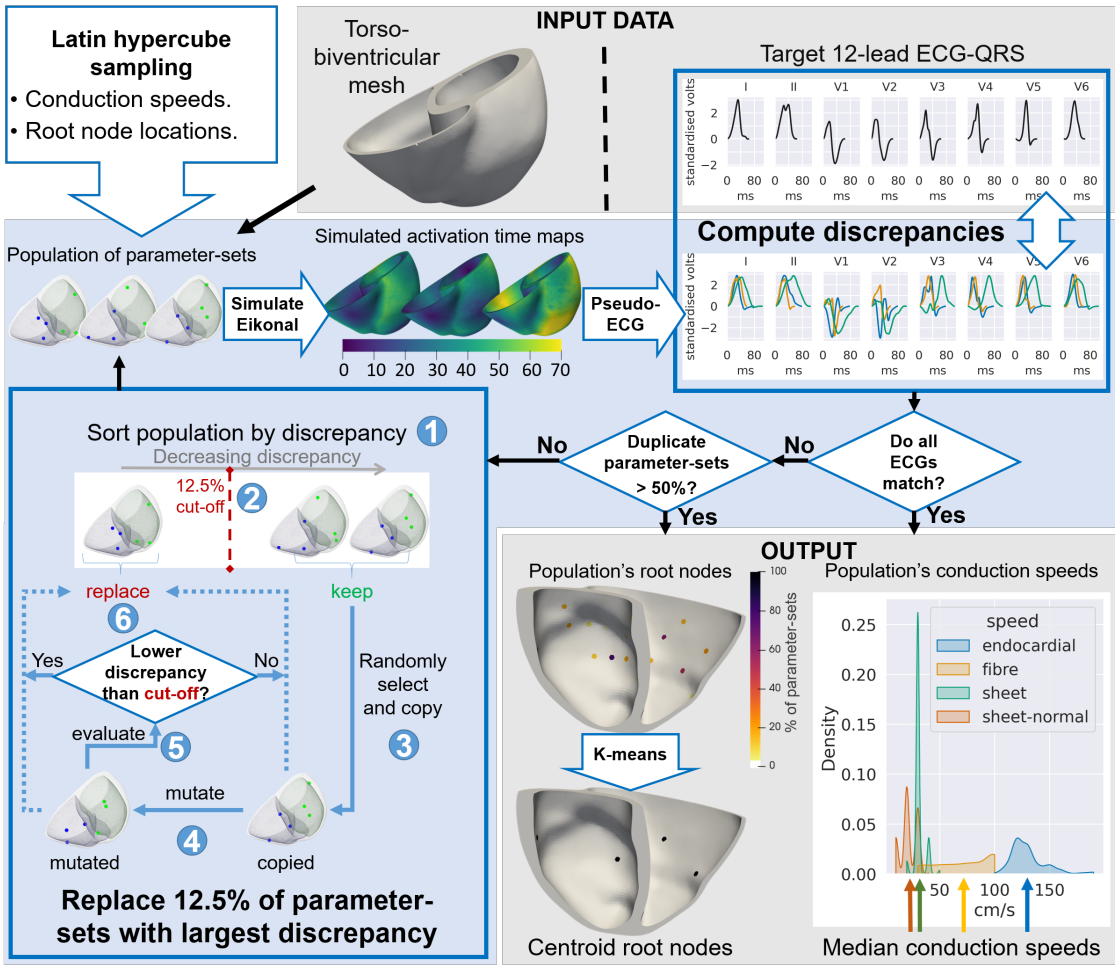

Figure 4.1: Diagram of the SMC-ABC-based inference pipeline, and its subprocesses, to recover the ventricular activation properties from CMR and electrocardiographic data. The process starts with the Latin hypercube sampling (top-left), generating a population of 512 models with different parameter values but the same propagation model equations and a subject-specific CMR-based torso-biventricular mesh. From there, the iterative process (centre and bottom-left – blue shaded area) commences in the direction of the arrows until any of the two stopping criteria (centre-right rhomboidal boxes) are fulfilled. From top-left to the right, the diagram highlights the following subprocesses: simulation of activation maps using the parameter-sets in the Eikonal model, generation of the 12-lead ECGs through the pseudo-ECG algorithm, computation of the discrepancy between each prediction and the target 12-lead ECG-QRS, stopping criteria, and, if not terminating, modification of the current population of parameter-sets (bottom-left blue box), which will replace each parameter-set with the 'high' (top 12.5%) discrepancy by either a copy of a parameter-set with low discrepancy (dotted line – right) or a mutated version of this copy (dotted line – left). The input and output data (both shaded in grey) are depicted at the top-right and bottom-right parts of the figure, respectively.

5. replace the parameter-sets corresponding to the top 12.5% highest discrepancies.

In Step-1, the Eikonal model (Section 4.4.1) simulates an activation time map from each parameter-set in the population using the subject's torso-biventricular mesh. Next, Step-2 (not needed when inferring from activation maps) computes the ECG-QRS from each activation time map using the pseudo-ECG algorithm (Section 4.4.2). Then, Step-3 computes the discrepancy (distance metric) between the predicted and target data (activation maps or ECGs) (Section 4.5). Step-4 checks if any of the following stopping criteria (Section 4.6) are fulfilled:

4.1. all predicted data 'match' the target (using a threshold);

4.2. more than 50% of the parameter-sets in the population are duplicates.

If the population does not fulfil any criteria, the iterative process continues. The following six substeps compose Step-5 (Section 4.6) (see numbers on Fig.1 – bottom-left):

5.1. sort the parameter-sets in the population according to the discrepancy values (decreasing order);

5.2. define a discrepancy cut-off to replace the parameter-sets with the top 12.5% highest discrepancies in the population;

5.3. for each parameter-set to be replaced (within the 12.5%), randomly select one to be kept and create two new copies from it;

5.4. 'mutate' only one of each pair of copies using Markov Chain Monte Carlo (Gilks, 2005);

5.5. simulate the Eikonal model for the mutated parameter-sets and compute their discrepancies;

5.6. finally, replace each parameter-set in the top 12.5% by the mutated alternative (dotted line on the left), unless its discrepancy would also have been rejected by the cut-off, in which case, use the unmodified copy for the replacement (dotted line on the right).

After Step-5, the iterative process restarts from Step-1 with simulations using the modified population of parameter-sets, and the inference process carries on until any stopping criterion in Step-4 is met.

The output area in Fig. 1 (bottom right – grey shaded area) also illustrates the post-processing of the parameter-sets in the resulting population. I combined all parameter values in the population to calculate a single solution parameter-set for each inference experiment. The strategy to combine our parameter-sets (Section 4.6) takes the median value for the conduction speeds and the k-means centroid for the root nodes throughout the population.

4.3 Virtual subjects for the 'target data'

The 'target data' were produced synthetically to know the ground-truth for the evaluation of the inference algorithm. They included noise-contaminated epicardial activation maps generated through simulating CMR-based torso-biventricular models. These twenty virtual subjects were obtained as described in Mincholé et al., 2019b through combinations of four torso-biventricular meshes constructed from the CMRs of four subjects and five conduction speed scenarios (Table 4.1) to account for the anatomical and functional variability in the healthy human population.

I considered four torso-biventricular CMR-based meshes with a variable torso cavity and biventricular myocardial-mass volumes to explore the effects of anatomical variability. These four meshes had the following torso cavity volumes: 23000 cm^3, 27000 cm^3, 54000 cm^3, and 44000 cm^3 for Mesh-1, Mesh-2, Mesh-3, and Mesh-4, respectively. The biventricular myocardial-mass volumes for these meshes were 74 cm^3, 76 cm^3, 107 cm^3, and 171 cm^3.

Target conduction speeds				
Conduction configuration	Endocardial	Fibre directed	Sheet directed	Sheet-normal directed
Normal speeds	150 cm/s	50 cm/s	32 cm/s	29 cm/s
Slow endocardium	120 cm/s	50 cm/s	32 cm/s	29 cm/s
Fast endocardium	179 cm/s	50 cm/s	32 cm/s	29 cm/s
Fast endocardium and myocardium	179 cm/s	88 cm/s	49 cm/s	45 cm/s
Slow endocardium and fast myocardium	120 cm/s	88 cm/s	49 cm/s	45 cm/s

Table 4.1: Conduction speed values selected to match the conductivities from the bidomain simulations in Mincholé et al., 2019b.

Then, as outlined in Table 4.1, I considered three fast-endocardial speeds to account for a slow (120 cm/s), normal (150 cm/s), and fast (179 cm/s) Purkinje network as Durrer et al., 1970 measured human endocardial speeds of up to 200 cm/s. I also considered two myocardial speed scenarios, namely, normal and fast myocardial conduction. The normal myocardial conduction speeds (50, 32, and 29 cm/s for the fibre, sheet, and sheet-normal speeds, respectively) were chosen to match the [extracellular, intracellular] conductivity pairs used in the bidomain model, namely [5.46, 1.5], [2.03, 0.45], and [2.03, 0.225] millisiemens/cm in the fibre, sheet and sheet-normal directions, respectively. Similarly, the fast myocardial conduction speeds (88, 49, and 45 cm/s) corresponded to conductivities [10.92, 3], [4.06, 0.9], and [4.06, 0.45] millisiemens/cm. The root node locations were set to seven homologous root node locations (three in the right and four in the left ventricle) following Cardone-Noott et al., 2016 on root node configurations that produced realistic healthy QRS complexes.

4.4 Simulation protocols

I conducted simulations using CMR-based Eikonal models solved using Dijkstra's algorithm (Dijkstra, 1959) (further described in Section 4.4.1) to generate both the 'target data' and the predictions from the inference method. All models implemented rule-based fibre orientations (Streeter et al., 1969) with fibre angle

transmurally changing from 60 to -60 degrees from endocardium to epicardium, respectively. The electrical conduction speeds were defined as orthotropic in the myocardial tissue and isotropic in the endocardia (fast endocardial layer) to emulate subendocardial the Purkinje network's effect. These modelling protocols were based on the work on generating realistic healthy human 12-lead ECG recordings with a torso-biventricular geometry by Cardone-Noott et al., 2016.

The 12-lead ECG 'target data' were computed from these Eikonal-generated activation maps using the pseudo-ECG algorithm (Gima and Rudy, 2002) (further described in Section 4.4.2), as in Mincholé et al., 2019b.

Finally, similarly to Grandits et al., 2020, I contaminated all 'target data' (i.e. epicardial activation time maps and ECGs) with white Gaussian noise to reach 20 decibels of signal-to-noise ratio as in Almeida et al., 2004 and Potyagaylo et al., 2014. The biventricular meshes used in the Eikonal models had edge-lengths ranging from 0.15 to 0.23 cm.

The same set of 'target data' were then obtained from bidomain (Section 2.6.1) simulations without noise-contamination for comparison. The bidomain simulations were conducted using the Chaste software (Pitt-Francis et al., 2009), and an extension of the O'Hara-Rudy action potential model (O'Hara et al., 2011) proposed by Dutta et al., 2017. The tetrahedral element edge lengths were set to 0.04 cm (re-meshed to ensure numerical convergence of the finite element solvers). The finer resolution of the meshes used for the bidomain simulations implied that the number of elements was 31 million compared to the 65000 in the coarser version of the mesh for the Eikonal. Moreover, the bidomain models implemented the following electrophysiological heterogeneities:

- Physiological apex to base action potential duration differences. These were modelled as a gradual decrease of the potassium ionic conductance (IKs) (from apex to base), resulting in differences of 40 ms in the action potential duration.

- Transmural heterogeneities were modelled as three regions of cellular models: the endocardial region from endocardium to 45% of the transmural width, the mid-myocardial region from 45% to 70% of the transmural width), and the epicardial region from 70% to the epicardial surface. Then, the action potential properties of the cell models in these regions were modelled as suggested in (O'Hara et al., 2011). These gradients resulted in interventricular action potential duration differences of 25 ms between left and right ventricles.

Note that these heterogeneities mainly affect the cardiac cycle's repolarisation phase. Thus, the Eikonal model without the heterogeneities is still able to reproduce a similar QRS complex. The target bidomain-ECGs were also computed from the bidomain-generated activation maps using the pseudo-ECG algorithm (Section 4.4.2).

4.4.1 Eikonal implementation

I implemented the Eikonal model (Section 2.6.3) as in Wallman et al., 2012. This implementation is hereafter reviewed together with its justification.

I assume that the wavefront propagates at constant speeds v_x, v_y, and v_z for the x, y, and z orthogonal basis of the cardiac space, respectively. Then I apply $\nabla d = v \cdot \nabla t$, where v is velocity and t is the travelling time of the wavefront $t(x, y, z)$ passing through a point (x, y, z) such that

$$\left.\begin{aligned} d_x = v_x t &\Rightarrow \partial_x d_x = v_x \partial_x t \\ d_y = v_y t &\Rightarrow \partial_y d_y = v_y \partial_y t \\ d_z = v_z t &\Rightarrow \partial_z d_z = v_z \partial_z t \end{aligned}\right\} \Longrightarrow \nabla d = \begin{pmatrix} v_x & 0 & 0 \\ 0 & v_y & 0 \\ 0 & 0 & v_z \end{pmatrix} \nabla t = V \nabla t \tag{4.1}$$

As previously described in Section 2.2.2, the activation sequence in human ventricles originates from the root nodes (i.e. earliest activation sites) in the endocardia. Thus, the effect of the Purkinje network during sinus-rhythm activity can be modelled using these root nodes combined with isotropic endocardial speed to emulate a densely connected network (Cardone-Noott et al., 2016; Mincholé et al., 2019b). This phenomenological approximation can reproduce realistic QRS complexes (Cardone-Noott et al., 2016).

By combining equations Equation (2.15) and Equation (4.1) and adjusting the constraint to account for the fact that the cardiac electrical wavefront originates from multiple locations (i.e. root nodes), I can rewrite the Eikonal equation as

$$\begin{cases} \sqrt{\nabla^T t V^2 \nabla t} = 1 \ , \quad in\ \Omega \in \mathbb{R}^3 \\ \qquad t(\Gamma_0) = 0 \end{cases} \ , \tag{4.2}$$

where Γ_0 is the set of locations (i.e. root nodes) in Ω from which the wavefront originates at $t = 0$.

This Eikonal formulation (Equation (4.2)) considers an anisotropic conduction speed medium with the fibres being aligned with the Cartesian axes for the electrical propagation. However, the fibres in the cardiac tissue are not necessarily aligned with the Cartesian axes (Section 2.2.3). To account for the fibre-dependant anisotropy of the conduction in the heart, I define the conduction speed to be a function of the fibre orientation $v(f, s, n)$, where f, s, and n, are orthonormal vectors that account for the fibre (longitudinal), sheet (transmural), and sheet-normal (normal to the tissue sheet plane) directions. Therefore, I define the conduction speeds using the fibre orientation basis such that now I have v_f, v_s, and v_n for the fibre, sheet, and sheet-normal directed speeds. Note that each edge in the biventricular mesh now will have a fibre orientation associated (f, s, n). Thus, I can compute the speed across that edge by transforming between the two bases systems as

$$(f, s, n) = \begin{array}{c} \begin{matrix} f & s & n \end{matrix} \\ \begin{pmatrix} x_f & x_s & x_n \\ y_f & y_s & y_n \\ z_f & z_s & z_n \end{pmatrix} \end{array} . \tag{4.3}$$

Then I can replace the V^2 term in the formulation in Equation (4.2) using

$$V^2 = \begin{pmatrix} x_f & x_s & x_n \\ y_f & y_s & y_n \\ z_f & z_s & z_n \end{pmatrix} \begin{pmatrix} v_f^2 & 0 & 0 \\ 0 & v_s^2 & 0 \\ 0 & 0 & v_n^2 \end{pmatrix} \begin{pmatrix} x_f & x_s & x_n \\ y_f & y_s & y_n \\ z_f & z_s & z_n \end{pmatrix}^T \tag{4.4}$$

to obtain a formulation of the Eikonal equation that accounts for the fibres' orientations in the myocardium.

When solving the Eikonal model, I know the fibre-oriented speeds, the Euclidean distance between two nodes a and b that share an edge (a, b), and the orientation of the fibres of the tissue represented by that edge. The Euclidean distance between two nodes sharing an edge is the edge-length $||\overrightarrow{(a,b)}||_2$ in the tetrahedral mesh. Therefore, by applying $t = \frac{d}{v}$ to Equation (4.4), I can compute the time cost $c_{a,b}$ for the wavefront to move between any two connected nodes, a and b, in the mesh as

$$c_{a,b} = \sqrt{\overrightarrow{(a,b)}^T V^{-2} \overrightarrow{(a,b)}} \ . \tag{4.5}$$

Wallman et al., 2012 demonstrated that constraining the wavefront to only be able to propagate through edges of the biventricular mesh yields equivalent activation sequences to the fast marching method (Konukoglu et al., 2007) while being significantly faster to compute. Hence, I discretise the Eikonal formulation, which enables the reinterpretation of the computation of the activation wavefront's travelling times (i.e. activation times), t, as a shortest-path problem on a non-directed graph. Therefore, I solve the Eikonal model using a multi-destination extension of Dijkstra's algorithm (Dijkstra, 1959). This shortest-path finding algorithm accounts for multiple sources (i.e. root nodes) of the electrical wavefront while computing the activation times for all the biventricular mesh nodes.

4.4.2 Electrocardiogram computation

I considered the eight independent leads in the 12-lead ECG, namely, leads I, II, and V1-6. The QRS complexes in specific ECG leads were computed from the activation time maps simulated through the Eikonal model using the pseudo-ECG equation (Gima and Rudy, 2002), as in Mincholé et al., 2019b. Alternative biophysically detailed approaches exist. For example, Potse, 2018 presents an accurate ECG computation approach for reaction-diffusion cardiac models, and Pezzuto et al., 2017 presents another ECG computation approach for fast phenomenological cardiac models, such as the Eikonal model. However, the cost of these approaches represented a computational bottleneck for the inference pipeline. Moreover, as illustrated in Figure 4.2, the simplest and fastest pseudo-ECG approach accurately

reproduces the ECGs computed from the bidomain. Therefore, I opted for the pseudo-ECG approach proposed in Gima and Rudy, 2002. This formulation combined with step functions, as in Potyagaylo et al., 2014, and a discretisation of the Eikonal formulation (Section 4.4.1) enables the fast sparse implementation of the pseudo-ECG equation that only updates the amplitude values using the elements showing changes in the gradients of their potential.

The pseudo-ECG equation for a 1D cable source with constant conductivity at a given electrode location, (x', y', z'), takes the form

$$\phi_e(x', y', z') = \frac{a^2\sigma_i}{4\sigma_e} \int -\nabla V_m \cdot \left[\nabla \frac{1}{r}\right] dx\, dy\, dz \ , \tag{4.6}$$

where V_m is the transmembrane potential, ∇V_m is its spatial gradient, r is the Euclidean distance Equation (2.16) from a given point (x, y, z) to the electrode's location, a is a constant that depends on the fibre's radius, and σ_i and σ_e are the intracellular and extracellular conductivities, respectively. An ECG signal is obtained by then considering this integral throughout the ventricular activation sequence period.

I combined the pseudo-ECG equation (Equation (4.6)) with a step function to compute the QRS complex as these accounts for the upstroke of the ventricular action potentials (Figure 2.5 – phase 0). The torso's conductivities modulate the amplitude of the QRS signals (Cardone-Noott et al., 2016). I consider an homogeneous bath between the heart and the electrodes, which can change the amplitude's scaling of my simulated QRSs compared the subject's. Thus, I standardise the amplitudes of simulated and subject's QRS complex in each lead to enable comparing their morphologies directly. Moreover, I assigned the V_m at each mesh's node to be one (depolarised) or zero (not depolarised yet) since the resulting QRSs are later starndardised. The standardisation process normalises the QRS signals to have zero mean and one standard deviation. After ignoring scaling factors in the original pseudo-ECG formulation, I can rewrite the equation as

$$\phi_e(x', y', z') = \sum_{j=1}^{N_{src}} -(\nabla V_m)_j \cdot \left[\nabla \frac{s_j}{r_j}\right] , \tag{4.7}$$

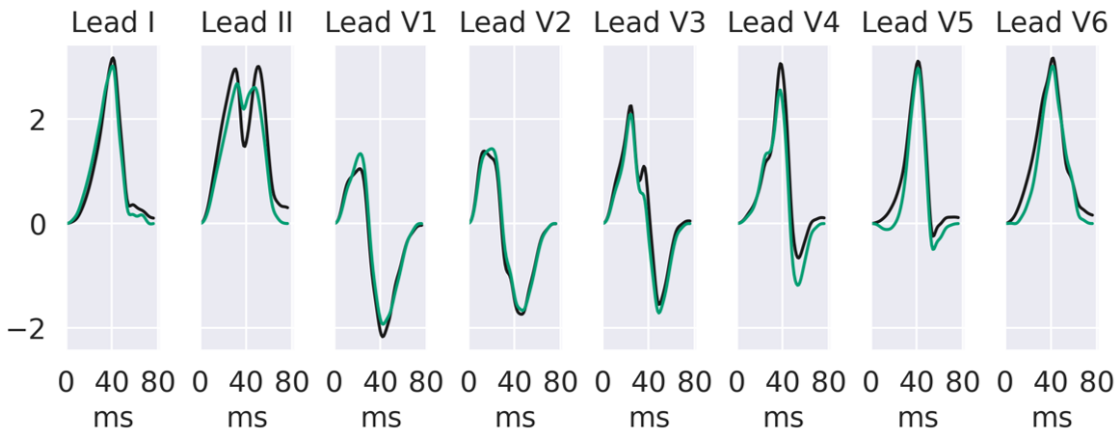

Figure 4.2: Comparison between pseudo-ECG (green line) and diffusion-based (black line) computed QRSs using the same baseline bidomain simulation. Pearson's correlation coefficient per each lead from left to right: 0.98, 0.93, 0.99, 0.99, 0.99, 0.97, 0.99, and 0.98. The amplitude of these signals is standardised to have mean zero and standard deviation one (y-axis). The time units are ms (x-axis).

where N_{src} is the number of tetrahedral source elements, $(\nabla V_m)_j$ is an estimated gradient over the jth tetrahedral element and s_j its normalised volume. Distances r_j are calculated using the centroids of each element. I estimate these gradients by assigning $V_{mi} = 1$ if the node i is activated, and $V_{mi} = 0$ if it is not. This formulation combined with the standarisation enabled generating nearly identical QRS signals to the bidomain's diffusion (Figure 4.2).

The inference process for each subject requires simulating about 100000 Eikonal models. The computational time for these Eikonal simulations using a fine mesh (31 million elements) is about 17 days using 100 CPUs. In contrast, these simulations can be done in 26 minutes using a coarse mesh (65000 elements) instead. Hence coarse meshes are necessary for the efficient generation of cardiac digital twins. However, simulating ECG signals from coarse meshes produces noise-like artefacts (Potse and Kuijpers, 2010; Tate et al., 2019; Schuler et al., 2019). Thus, I filter the ECG recordings with a lowpass filter with a cut-off frequency set to 150 Hz following the Nyquist-Shannon sampling theorem and the guidelines in frequential information in the ECG from Sörnmo and Laguna, 2005. Finally, I align the signals to start at 0 (standardised voltage units).

4.5 Discrepancy calculation

In order to calculate differences between 'predicted' and 'target' epicardial activation maps, I define their discrepancy as the average of the absolute differences between each epicardial nodes' predicted and target activation time,

$$d_{ATM} = \sum_{i=1}^{N_{nodes}} \left| \frac{pAT_i - tAT_i}{N_{nodes}} \right|, \tag{4.8}$$

where pAT_i is the predicted activation time of the ith node in the biventricular mesh, and tAT_i is the target activation time value of that node.

On the other hand, capturing differences in activation properties from differences between 12-lead QRS recordings is challenging. Dynamic time warping (DTW) is a speed-invariant dynamic programming algorithm for measuring differences between signals. However, classic DTW allows QRSs to be compared using non-physiological warping (Luzianin and Krause, 2016). I present an extension of DTW that prevents non-physiological QRS transformations by imposing a warping-slope constraint (Sakoe and Chiba, 1978), a parallelogram constraint (Itakura, 1975), and common-start and common-end warping constraints. The warping-slope constraint prevents more than two warping steps from occurring consecutively (i.e. 1 ms can warp to up to 3 ms, and vice versa). The parallelogram constraint limits the amount of warping at the signals' start and favours 'diagonal' warping such that one signal is transformed as the stretched or shrank version of the other. In other words, it rewards signals with physiologically similar morphology, namely, penalising non-monotonic warping, especially when occurring discontinuously or displaying unwarping (i.e. stretching and shrinking different parts of the same signal within the same comparison). These warping penalties within the parallelogram are more prominent at the start of the signals and decrease linearly, allowing for more flexible warping when comparing the end of the QRSs since differences in the activation sequence can have a cumulative impact on the QRS. The common-start and common-end constraints imply that both signals represent the same phenomena and should be entirely encoded into each other. Overall, my DTW method favours signals with similar morphologies, even if their source conduction speeds were scaled to fast-track

the root nodes' identification. Therefore, I implement an additional penalty based on the difference in QRS-width to identify the source model's correct speeds, which, combined with the DTW extension described, composes my DTW-based discrepancy.

4.5.1 Parameter space exploration

My approach aims to infer the root node locations and the four conduction speeds (i.e. endocardial, fibre, sheet, and sheet-normal speeds) that enable reproducing a subject's 'target data' (epicardial activation map or 12-lead ECG). I considered three sets of possible root node locations, namely, low, high and hybrid resolution of the discretisation of the root node parameter space, to expose the method's ability to represent root node locations from epicardial activation time map and ECG data modalities.

In the low resolution of the root node discretisation, the possible root node locations were preselected as the centres of each endocardial section following the American Heart Association's segmentation guidelines (Cerqueira Manuel D. et al., 2002), as in Arevalo et al., 2016, while ensuring that any point in either ventricle's endocardia had at least one root node in the same ventricle not more than 2.5 cm away. This selection strategy led to 38 candidate locations (18 in the left and 20 in the right endocardium) in the largest geometry (Mesh-4) and to 27 (11 in the left and 16 in the right endocardium) in the smallest one (Mesh-1).

A high resolution of the root node discretisation was also considered by uniformly sampling more root nodes such that the distance from each point to a root node was reduced to 1.5 cm. This high resolution produced 133 candidate root nodes (64 in the left and 69 in the right endocardium) for the largest anatomy and 98 (39 in the left and 59 in the right endocardium) for the smallest one.

Finally, I considered a hybrid resolution that combined the low resolution for the right ventricle with the high resolution for the left ventricle since the latter is presumed to have a stronger influence on the QRS. I anticipate the inference from epicardial activation maps will work better from the symmetric root node distributions (low and high) than from asymmetric ones (hybrid) as

the epicardial activation map equally represents both ventricles. On the other hand, the hybrid resolution should improve the inference from 12-lead ECGs since it emphasises the exploration of the root nodes on the left ventricle, which is known to have more influence on the ECG than the right ventricle given its larger myocardial-mass volume.

These candidate root node locations were considered a binary parameter, with the inference yielding 'in use' or 'not in use' to obtain a good match between simulated and target activation sequence. The two resolutions allowed testing the inference problem's sensitivity to the availability of root node locations, also considering the balance with feasibility and the computational expense of increasing these resolutions.

I define the range of the number of root nodes as [6, 10] since Cardone-Noott et al., 2016 demonstrated that seven root nodes were sufficient to simulate healthy QRS complexes. I considered no region or ventricle specific rules for the placement of the root nodes to allow for the straightforward extension of the method to disease conditions, such as bundle branch block.

Finally, the conduction speeds were assigned to be within predefined physiological ranges (Durrer et al., 1970). The endocardial and myocardial speeds were bounded within the ranges [100, 200] and [25, 90] cm/s, respectively. Furthermore, I constrained the fibre speed to be larger than the sheet speed, and sheet speed to be larger than the sheet-normal speed to be consistent with the literature (Frazier et al., 1988; Saffitz et al., 1994; Caldwell et al., 2009).

4.6 Parameter inference with SMC-ABC method

The combination of conduction speeds and root nodes creates an inference problem with continuous and discrete mixed-type parameter space that challenges many parameter inference algorithms. I propose the application of an SMC-ABC-based algorithm (Section 2.7.1) (Sisson et al., 2007; Drovandi and Pettitt, 2011) that I specifically designed to explore our mixed-type parameter space efficiently.

As reviewed in the description of Figure 4.1, SMC-ABC makes use of a population of models with different parameter-sets (similarly to Britton et al., 2013) that represent the parameter search space of interest. Then, at each iteration the method shrinks this parameter space of interest by replacing the parameter-sets representing regions currently associated with the highest discrepancies to the regions with currently acceptable discrepancy values. This replacing action is carried out within the resampling step (Step-4 in Figure 4.1) of the algorithm.

In generating new proposals for SMC-ABC's mutation steps (substep 4 in Step-4 from Figure 4.1), I ignore any dependencies between the discrete (root node locations) and continuous (conduction speeds) parameter spaces. For the continuous space, I use random Gaussian jumps informed by the current covariance of the parameter-sets. For the discrete root node space, however, I propose a novel strategy to generate proposals that are independent of a parameter-set's current values but informed by the whole population. Independent proposal distributions are asymmetric but can significantly accelerate SMC approaches by being more likely to suggest large jumps across the parameter space (South et al., 2019).

More precisely, I first determine the number of root node locations for a new parameter-set. To choose the number of root nodes, 80% of the time, I inherit the number of root nodes from a random parameter-set in the population. To prevent premature convergence, the remaining 20% of the time, I sample the number of root nodes from a normal distribution with a mean of seven and a standard deviation of two over the range defined for the number of root nodes (Section 4.5.1). Next, I place a Dirichlet (i.e. multivariate Beta) prior on the propensities that a candidate root node location is 'in use'. I treat the frequency of a candidate location being 'in use' in our population as a sample from a multinomial distribution. This choice of treatment allows me to calculate these propensities from the current set of parameter-sets since the Dirichlet distribution is the conjugate prior for a multinomial likelihood. Due to the properties of conjugate priors, our posterior on propensities for the candidate root node locations will also be Dirichlet distributed. Given the number of root nodes to be 'in use', I select them one at a time by first

sampling from the Dirichlet posterior implied by the 'relevant' parameter-sets in the population to obtain a set of propensities. Then I sample from the (single-trial) multinomial distribution with these propensities.

A parameter-set is 'relevant' if it has the same number of root nodes as the parameter-set being created and the same root node locations 'in use' as those selected so far during this process. The population of relevant parameter-sets that informs selecting the next candidate root node location becomes smaller after each selection. Consequently, the first few selections of root nodes will represent the existing patterns in the population of parameter-sets, while the final selections should explore new candidate root node locations. Importantly, by taking this approach and only using matching parameter-sets to update the propensities each time a candidate root node is selected, I capture interdependencies between root node locations in the population.

The SMC-ABC algorithm's iterative process considers that the current population is a solution to the inference problem when all discrepancies are smaller than a positive tolerance (Section 2.7.1). This tolerance threshold should inform the method to terminate when all the parameter-sets in the population are acceptable solutions to the inference problem. For this study, I was interested in quantifying how accurately my methodology can replicate the target data during the inference. Therefore, I defined a relatively small tolerance (on the scale of the distortion from applying white Gaussian noise-contamination with 20 dB signal-to-noise ratio). However, just using this small tolerance would force the inference to converge to the best-performing parameter-set, whereas this inverse problem is known to have non-unique solutions. Thus, I defined a second stopping criterion relative to the percentage of duplicate parameter-sets to terminate the inference before the population collapses into a single parameter-set. In other words, this second stopping criterion informs the algorithm to terminate when the Markov Chain Monte Carlo process fails to find new acceptable parameter-sets too many times (right dotted arrow in Figure 4.1 blue box – bottom left), causing 50% of the

parameter-sets in the population to be duplications. Fulfilling either stopping criterion will suffice to terminate the inference process.

4.6.1 Error metrics and post-processing

I propose a k-means based root node aggregation strategy to overcome the performance cap determined by the discretisation of the root node parameter space. More precisely, I define the resulting root nodes locations to be the centroids of the clusters found by applying k-means clustering to the values in each population of parameter-sets. I initialised the centroids in k-means as the most frequently occurring configuration for the root nodes in the population to ensure that the starting configuration was spread out and accounted for a representative number of root nodes. This aggregation strategy allows visualising the inferred root nodes directly in the shared biventricular anatomy. I also report the mean plus-minus ($\pm$) standard deviation of the (unsigned) distance between each target root node and the closest centroid from k-means and the mean $\pm$ standard deviation of the error in the number of root nodes.

I aggregate the inferred conduction speeds as the median of each speed in the population of parameter-sets. To illustrate the error, I propose a speed-normalised error metric so that I can jointly represent errors from different virtual subjects. This error was defined as

$$error = 100 * \frac{S' - S}{S} \ , \tag{4.9}$$

where S' is the inferred conduction speed value, and S is the ground-truth speed.

I adopted the mean $\pm$ standard deviation of Pearson's correlation coefficient (Bear et al., 2018; Serinagaoglu Dogrusoz et al., 2019; Schaufelberger et al., 2019) to measure disagreement between our inference predictions and the 'target data' (epicardial activation time maps or ECGs).

Additionally, I report p-values and the number of samples using the dependent t-test for paired (with the t-distribution) to support the observations drawn from the results.

4.6.2 Hyperparameter calibration

I calibrated the algorithms presented in this study using general and physiological knowledge and adjusted their values to work in one of the twenty virtual subjects (Mesh-4 with 'Normal speeds' from Table 4.1) using the low resolution for the root node discretisation. The results presented in this study were generated using the same calibration of the hyperparameters regardless of the virtual subject, the modality of the 'target data', or the resolution of the root node parameter space's discretisation. With this non-recalibration strategy, I hoped to identify which hyperparameters require subject-specific adjustments and how impactful this is.

Our society, in which there reigns an eager desire for riches and luxury, does not understand the value of science. It does not realize that science is a most precious part of its patrimony.

— Marie Curie (Curie, 1937)

5

Inference of ventricular activation properties from activation time maps

5.1 Introduction

As mentioned in Chapter 1, the goal of this thesis is to develop computational methods for the efficient generation of cardiac digital twins. Chapter 4 presented an efficient pipeline for generating cardiac digital twins from CMR and electrocardiographic data. The activation wavefront in the ventricles propagates transmurally from endocardium to epicardium. The time at which each point in the epicardium is reached by this wavefront, namely, the epicardial activation time map, can be non-invasively derived through ECGI (Section 2.5). These activation maps encapture the dynamics of the activation sequence in the ventricles. The goal of the study presented in this chapter is to investigate new computational techniques for the efficient quantification of subject-specific ventricular activation properties using CMR-based modelling and simulation and activation time maps. This chapter presents and evaluates the application of the methodology for inferring the properties underlying the dynamics of the activation sequence in the ventricles (Chapter 4), namely, the conduction speeds and root nodes (Section 2.2.3), from the epicardial activation time maps.

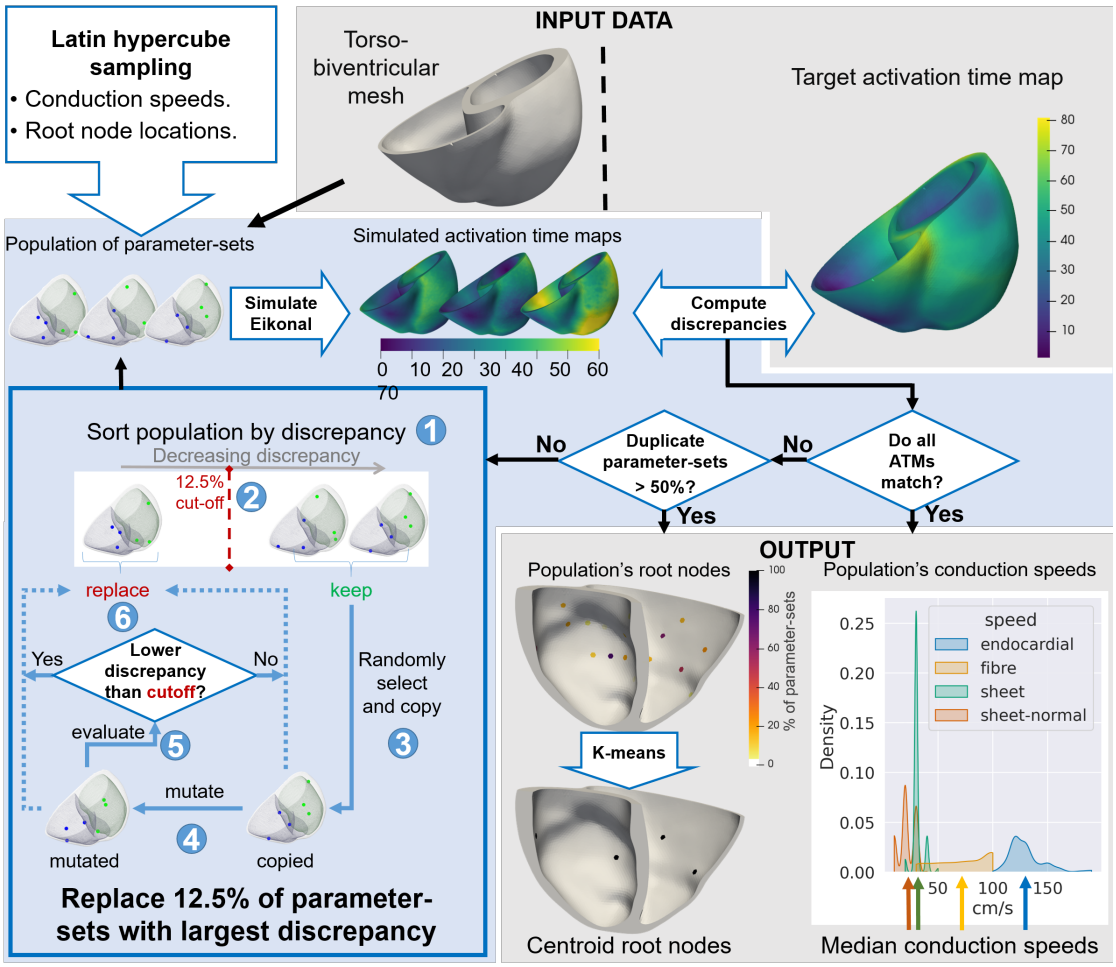

Figure 5.1: Inference pipeline to recover the ventricular activation properties from CMR and epicardial activation map data. The process starts with the Latin hypercube sampling (top-left), generating a population of 512 models with different parameter values but the same propagation model equations and a subject-specific CMR-based torso-biventricular mesh. From there, the iterative process (centre and bottom-left – blue shaded area) commences. The four steps composing the iterative process are the simulation of activation maps using the parameter-sets in the Eikonal model, computation of the discrepancy between each prediction and the target epicardial activation map, evaluation of the stopping criteria, and modification of the current population of parameter-sets.

5.2 Methods

Figure 5.1 presents the diagram of the inference method from epicardial activation maps. As mentioned in Section 4.2, the inference of the root nodes and conduction speeds (endocardial, fibre, sheet, and sheet-normal speeds) from epicardial activation maps implements the following four steps: simulate the activation maps using the Eikonal model (Section 4.4.1); compute the discrepancies between the simulated and

target epicardial activation maps (Equation 4.8); evaluate the two stopping criteria (Section 4.6); replace the parameter-sets with the highest discrepancies (Section 4.6).

The components of this inference pipeline, highlighted in Figure 5.1, are described in Chapter 4 together with the strategies to calculate the accuracy of the inference (Section 4.6.1) and to calibrate the hyperparameters (Section 4.6.2) of all these methodologies. The following sections in this Chapter present the evaluation of this inference method on the epicardial activation time maps from the twenty virtual subject cohort proposed in Section 4.3.

5.3 Results

5.3.1 Reproduction of noise-contaminated epicardial activation maps

Figure 5.2 illustrates the ability of our inference method to replicate the 'target data' from epicardial activation maps. The inference method accurately replicated the Eikonal activation maps from noise-contaminated epicardial data (Figure 5.2) showing slightly greater accuracy for the hybrid root node resolution than the low or high resolutions (Figure 5.2.B). The replication of the non-septal activation times was more accurate than the septal ones (Figure 5.2.A.vii and Figure 5.2.A.xv) except for the high-resolution case (Figure 5.2.A.xi).

5.3.2 Root node inference from noise-contaminated epicardial activation maps

This section presents the root node inference results from noise-contaminated Eikonal epicardial activation time maps as 'target data'. Figure 5.3, Figure 5.4, Figure 5.5, and Figure 5.6 illustrate the inferred root nodes locations for the virtual subjects with Mesh-1, Mesh-2, Mesh-3, and Mesh-4, respectively. The summary error metrics of this root node inference are reported in Table 5.1.

Our method inferred more accurately root nodes in non-septal areas since septal activation times are not represented in epicardial maps (Figure 5.3, Figure 5.4, Figure 5.5 and Figure 5.6). Increasing the resolution of the root node location

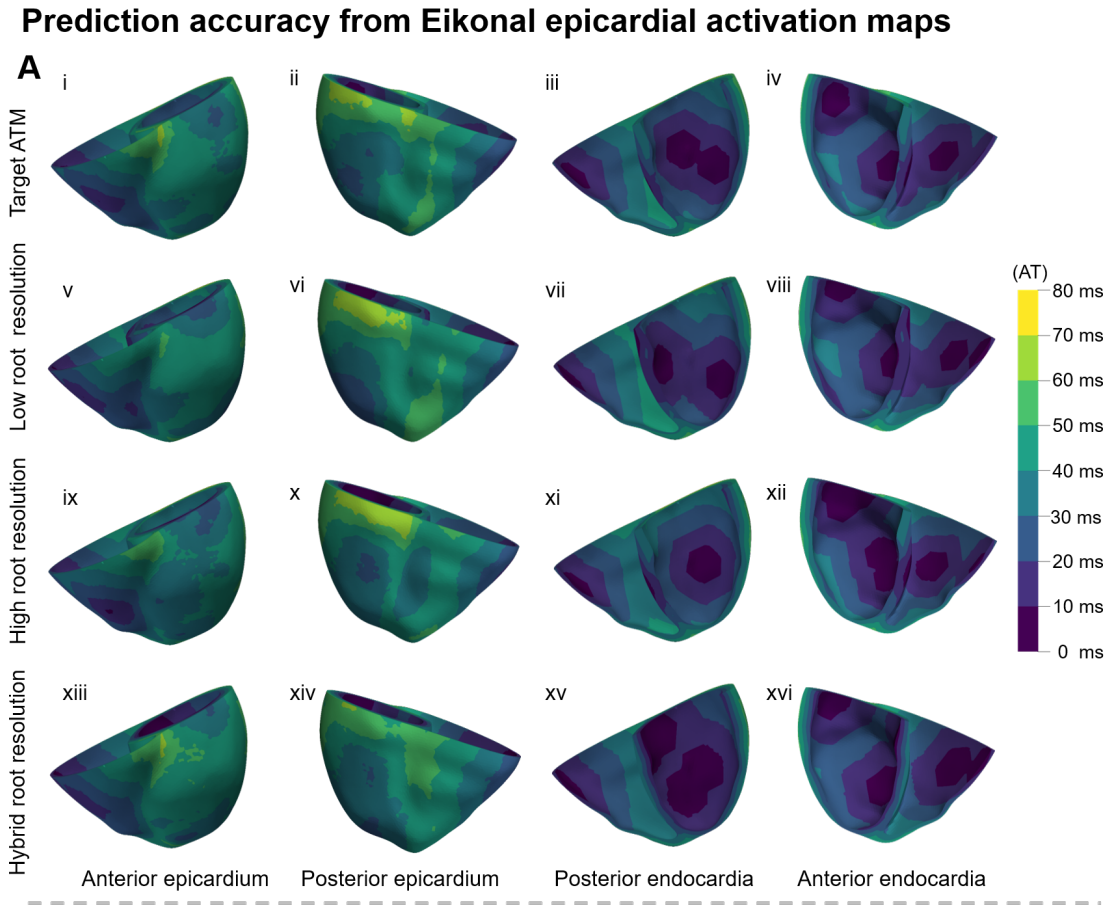

B

Anatomy	Low	High	Hybrid
Mesh-1	0.79 ± 0.06	0.78 ± 0.09	0.82 ± 0.07
Mesh-2	0.83 ± 0.07	0.79 ± 0.07	0.85 ± 0.05
Mesh-3	0.74 ± 0.07	0.84 ± 0.05	0.79 ± 0.06
Mesh-4	0.84 ± 0.06	0.86 ± 0.06	0.85 ± 0.06
All	0.80 ± 0.07	0.82 ± 0.07	**0.83 ± 0.06**

Figure 5.2: Agreement between predicted and Eikonal target activation time maps from the inference process; (**A**) on the virtual subject with Mesh-4 and 'Slow endocardial speed' properties (Table 4.1); (**B**) mean ± standard deviation of Pearson's correlation coefficients. (**A**) illustrates three activation time maps (ATM) (one per row) coloured using isochronic activation time (AT) bands of ten ms (colour bar). Top row: target; second row: prediction from the inference using the low resolution for the root node discretisation (0.93 Pearson's correlation coefficient with 'target data' – top row); third row: prediction from the inference with high root node resolution (0.89 Pearson's correlation coefficient); bottom row: prediction from the inference with hybrid root node resolution (0.87). Each column shows a different view of the same row-example. (**B**) mean ± standard deviation of Pearson's correlation coefficients between the predicted and target data. Low, High and Hybrid resolution of the discretisation of the root node parameter space.

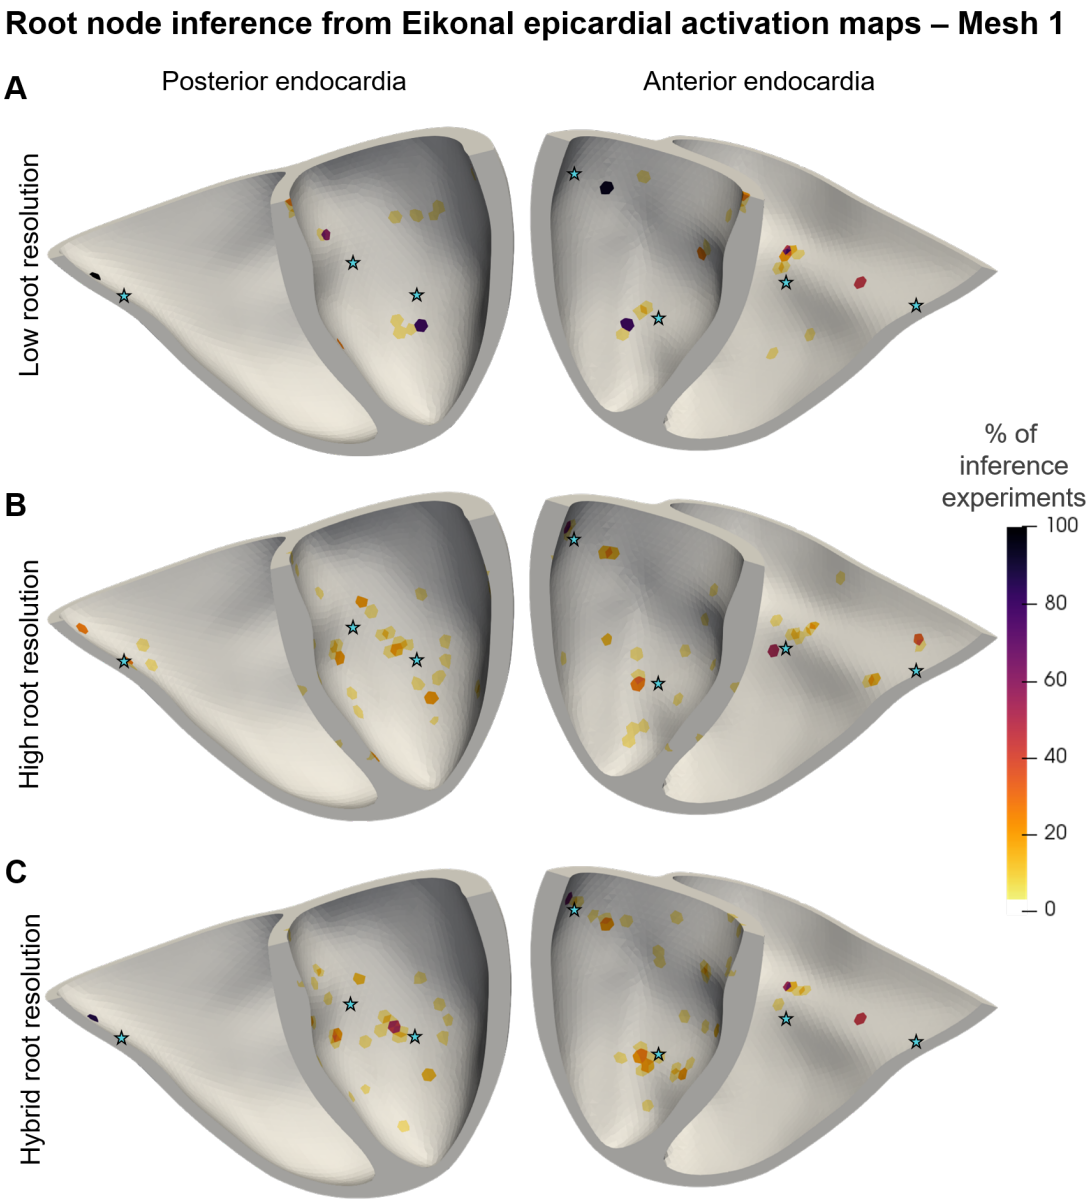

Figure 5.3: Root nodes inferred from noise-contaminated Eikonal epicardial activation maps on Mesh-1 with low (**A**), high (**B**), and hybrid (**C**) resolution of the root node space discretisation. The stars indicate the ground-truth root node locations. The endocardial surface is coloured as a heatmap showing the percentage of experiments that inferred each location as a root node.

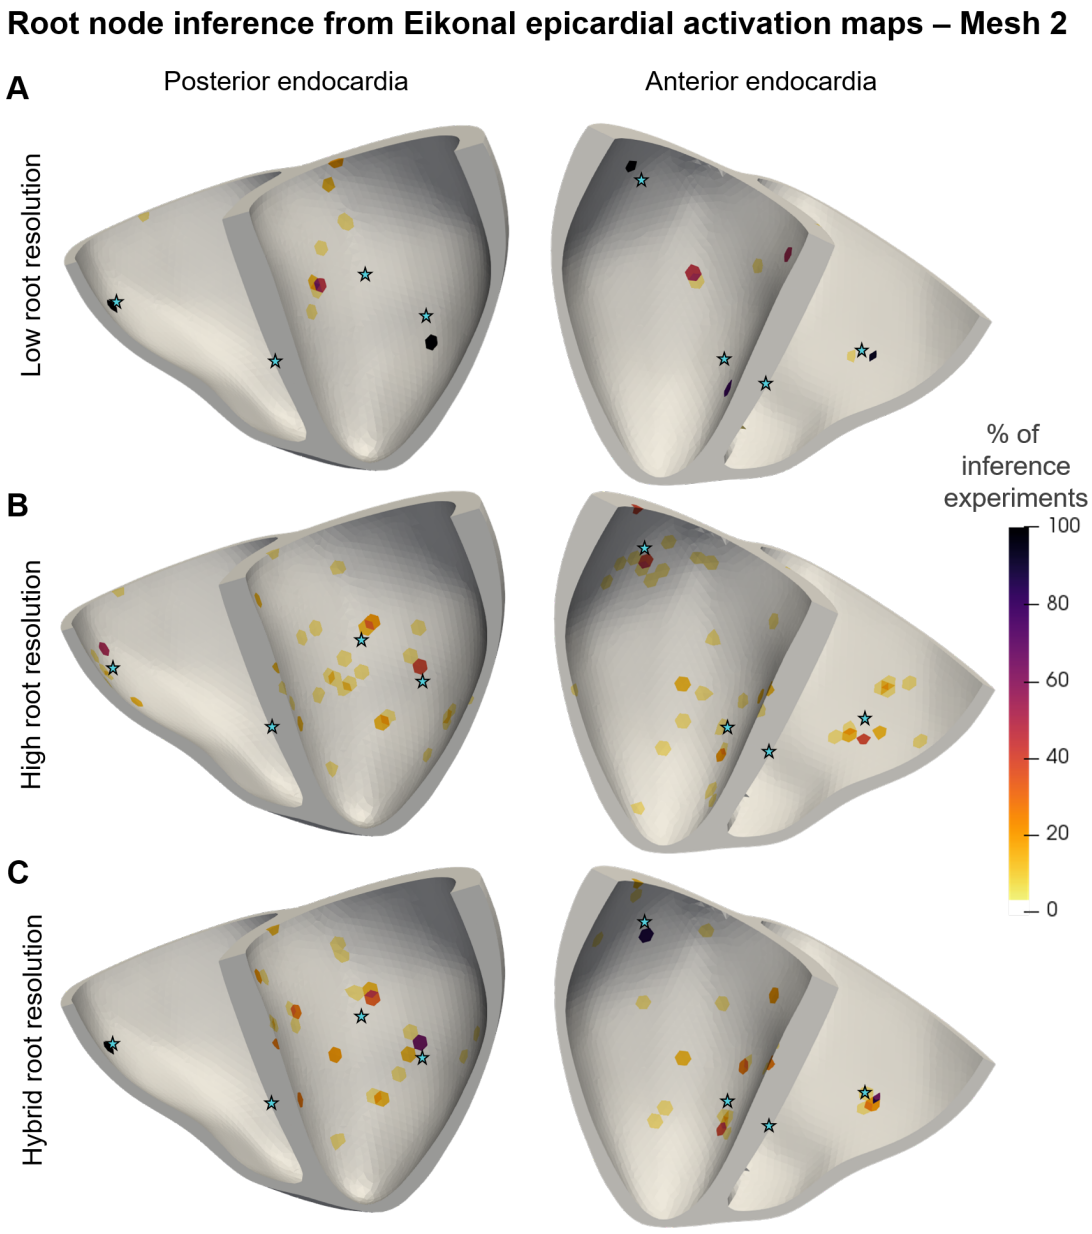

Figure 5.4: Root nodes inferred from noise-contaminated Eikonal epicardial activation maps on Mesh-2 with low (**A**), high (**B**), and hybrid (**C**) resolution of the root node space discretisation. The stars indicate the ground-truth root node locations. The endocardial surface is coloured as a heatmap showing the percentage of experiments that inferred each location as a root node.

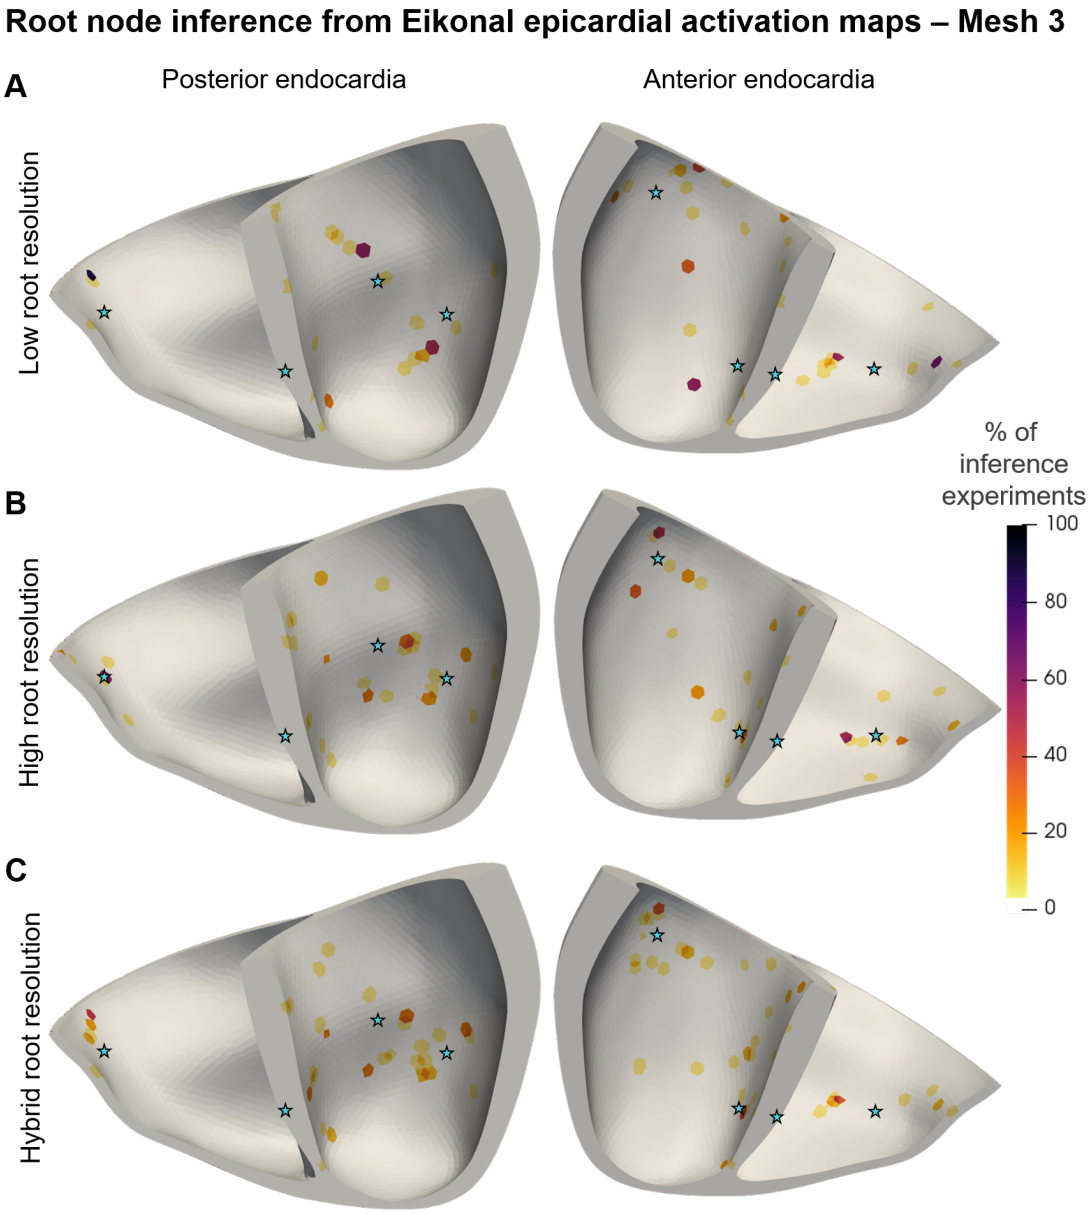

Figure 5.5: Root nodes inferred from noise-contaminated Eikonal epicardial activation maps on Mesh-3 with low (**A**), high (**B**), and hybrid (**C**) resolution of the root node space discretisation. The stars indicate the ground-truth root node locations. The endocardial surface is coloured as a heatmap showing the percentage of experiments that inferred each location as a root node.

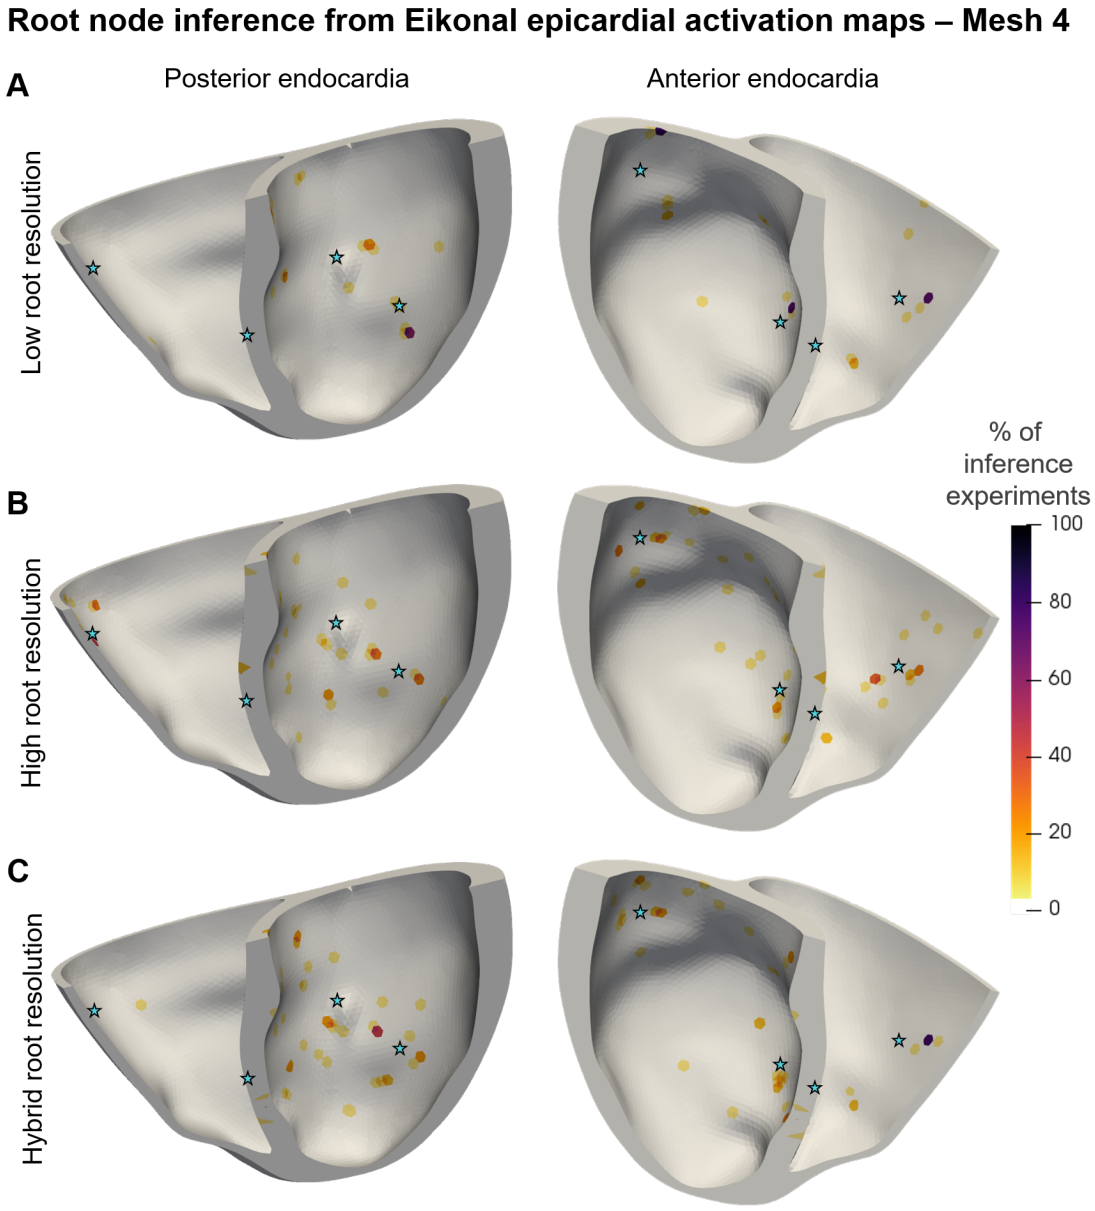

Figure 5.6: Root nodes inferred from noise-contaminated Eikonal epicardial activation maps on Mesh-4 with low (**A**), high (**B**), and hybrid (**C**) resolution of the root node space discretisation. The stars indicate the ground-truth root node locations. The endocardial surface is coloured as a heatmap showing the percentage of experiments that inferred each location as a root node.

Root node errors	**Low**	**High**	**Hybrid**
LV: Closest root node (cm)	1.1 ± 0.2	1.0 ± 0.4	0.9 ± 0.2
RV: Closest root node (cm)	1.2 ± 0.3	1.0 ± 0.4	1.4 ± 0.4
LV: Number of root nodes	1.1 ± 0.9	1.2 ± 0.8	2.4 ± 0.8
RV: Number of root nodes	0.5 ± 0.5	0.7 ± 0.7	0.6 ± 0.5

Table 5.1: Root nodes errors metrics (Section 4.6.1) from noise-contaminated Eikonal activation maps. Absolute mean ± standard deviation of the distance from the ground-truth and error in the number of root nodes inferred from all meshes. **LV** – left ventricle; **RV** – right ventricle; **Low**, **High**, and **Hybrid** resolutions of the discretisation of the root node locations.

discretisation aided in recovering root node locations closer to the ground truth from activation maps (Table 5.1). The number of root nodes was better identified from the low resolution, whereas increasing the resolution improved the root node positioning accuracy (Table 5.1). The hybrid resolution identified better the root nodes in the left ventricle than those in the right one (Table 5.1).

5.3.3 Speed inference from noise-contaminated epicardial activation maps

Figure 5.7 reports the errors for the inference of conduction speeds (Equation 4.6.1) from Eikonal epicardial activation time maps. The results are shown for low (A), high (B), and hybrid (C) resolution of root node discretisation and grouped by torso-biventricular mesh in the virtual subject (see colour legend in Figure 5.7). The right side of the figure provides a numerical summary of the left side to ease readability.

The endocardial and sheet speeds in the inference experiments implementing the low or high root node resolution were recovered more accurately than the fibre directed speed (Figure 5.7.A-B) (p-values < 0.008, 59 samples). Throughout all root node resolutions, the endocardial, sheet, and sheet-normal speeds displayed a narrower distribution of error values suggesting that these speeds were consistently recovered at the same accuracy regardless of the characteristics of the virtual subject (Figure 5.7).

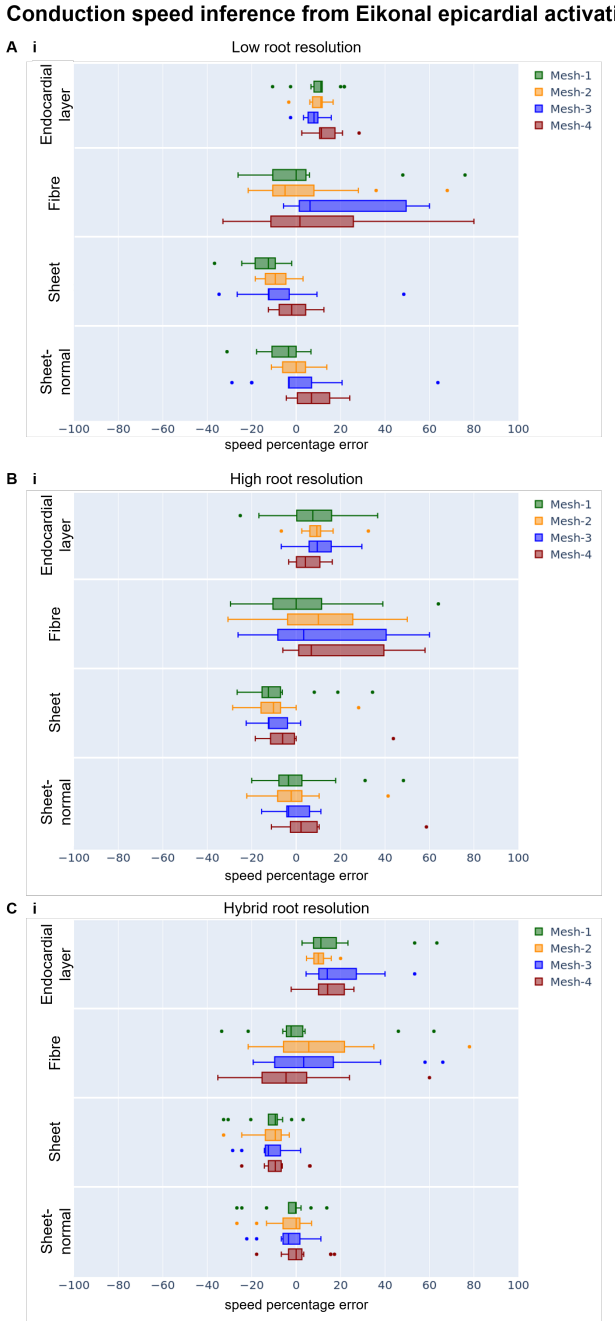

A ii

Conduction speed	Low root resolution
Endocardial	10.8 ± 5.10
Fibre	19.0 ± 21.3
Sheet	11.6 ± 9.00
Sheet-normal	9.30 ± 10.3

B ii

Conduction speed	High root resolution
Endocardial	10.5 ± 8.00
Fibre	20.0 ± 18.2
Sheet	12.4 ± 8.50
Sheet-normal	9.10 ± 11.2

C ii

Conduction speed	Hybrid root resolution
Endocardial	15.5 ± 11.9
Fibre	16.5 ± 18.1
Sheet	12.0 ± 7.20
Sheet-normal	6.60 ± 7.10

Figure 5.7: Error in the conduction speeds inference (Equation 4.6.1) from noise-contaminated Eikonal epicardial activation maps using low (**A**), high (**B**), and hybrid (**C**) resolution of the root node discretisation. The errors are computed as (Equation 4.6.1) the percentage over the ground truth conduction speeds (x-axis) represented as (**i**) box-plots grouped by anatomy (colour) and conduction speed (y-axis); (**ii**) and represented as absolute mean ± standard deviation.

Torso-biventricular anatomy	**Low**	**High**
Mesh-1	0.85 ± 0.05	0.94 ± 0.03
Mesh-2	0.86 ± 0.04	0.94 ± 0.03
Mesh-3	0.83 ± 0.03	0.95 ± 0.02
Mesh-4	0.91 ± 0.02	0.96 ± 0.01

Table 5.2: Prediction accuracy of the inference from bidomain epicardial activation time maps (mean ± standard deviation of Pearson's correlation coefficients). **Low** and **High** resolution of the discretisation of the root node parameter space.

Root node errors	**Low**	**High**
LV: Closest root node (cm)	1.1 ± 0.1	0.5 ± 0
RV: Closest root node (cm)	0.7 ± 0.1	0.4 ± 0.1
LV: Number of root nodes	0.7 ± 0.7	1.7 ± 0.5
RV: Number of root nodes	0.2 ± 0.2	0.1 ± 0.2

Table 5.3: Root node recovery error metrics (Section 4.6.1) (absolute mean ± standard deviation) from bidomain epicardial activation maps. Distance from the ground-truth and error in the number of root nodes inferred. **LV** – left ventricle; **RV** – right ventricle; **Low** and **High** resolution of the discretisation of the root node parameter space.

5.3.4 Inference from noise-free bidomain epicardial activation maps

Whereas the inference accuracy from noise-contaminated epicardial activation maps (Figure 5.2.B) was mostly unaffected by the resolution of the discretisation of the root nodes, increasing the root node resolution consistently improved the accuracy for replicating epicardial activation maps in the inference experiments targeting noise-free bidomain-simulated data (Table 5.2).

The high resolution of the root node location discretisation also helped recover root node locations closer to the ground truth from epicardial activation maps (Table 5.3). However, this increase in resolution was typically accompanied by an overestimation of the number of root node locations in the inference from bidomain epicardial activation maps.

In agreement with our previous findings from Eikonal 'target data' (Figure 5.7), the endocardial and sheet speed were also consistently better recovered from bidomain epicardial activation time maps than the fibre and sheet-normal speeds (Table 5.4). Moreover, the choice of resolution of the root node discretisation

Conduction speed	Low	High
Endocardial	12.9 ± 3.9	2.5 ± 1.8
Fibre	27.6 ± 8.2	19.0 ± 9.2
Sheet	5.1 ± 5.6	4.8 ± 2.6
Sheet-normal	7.4 ± 9.2	5.6 ± 4.7

Table 5.4: Conduction speed recovery error metrics (Equation 4.6.1) (absolute mean ± standard deviation) from bidomain epicardial activation maps. **Low** and **High** resolution of the discretisation of the root node parameter space.

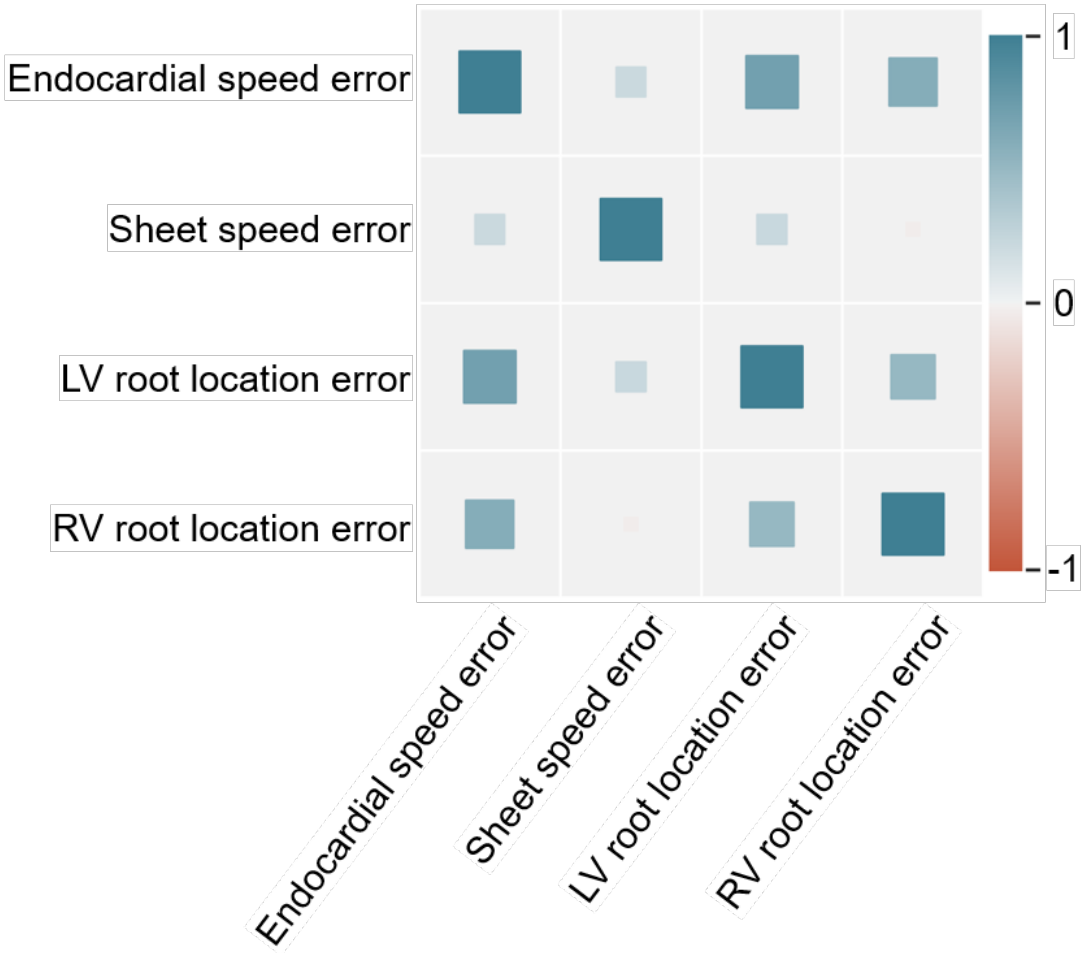

Figure 5.8: Heatmap of Pearson's correlation coefficients between the root node and the conduction speed inference errors from bidomain epicardial activation time maps.

considerably improved the estimation of the conduction speeds (Table 5.4), as it did for the root nodes (Table 5.3), suggesting a positive correlation between the identifiability of these activation properties (Figure 5.8).

Torso-biventricular anatomy	Endo x2	Endo /2	Fibre x2	Fibre /2	Sheet x2	Sheet /2	Norm x2	Norm /2
Mesh-1	0.921	0.93	0.994	0.995	0.966	0.951	0.997	0.996
Mesh-2	0.922	0.93	0.995	0.995	0.964	0.949	0.996	0.995
Mesh-3	0.93	0.921	0.996	0.996	0.959	0.953	0.996	0.996
Mesh-4	0.929	0.924	0.997	0.998	0.954	0.947	0.998	0.998

Table 5.5: Pearson's correlation coefficients between activation maps from different speeds. Correlation coefficient between the baseline 'normal speeds' (Table 4.1), namely, 150 cm/s, 50 cm/s, 32 cm/s, and 29 cm/s for the endocardial, fibre-directed, sheet-directed, and sheet-normal directed speeds, respectively, and a simulation of the same calibration with one speed changed. Each row displays the results for one of my four torso-biventricular anatomies and each column displays the correlation coefficients when modifying the indicated speed. **Endo x 2**: twice the endocardial speed (300 cm/s). **Endo / 2**: half the endocardial speed (150 cm/s). **Fibre x 2**: twice the fibre-directed speed (100 cm/s). **Fibre / 2**: half the fibre-directed speed (25 cm/s). **Sheet x 2**: twice the sheet-directed speed (64 cm/s). **Sheet / 2**: half the sheet-directed speed (16 cm/s). **Norm x 2**: twice the sheet-normal directed speed (58 cm/s). **Norm / 2**: half the sheet-normal directed speed (14.5 cm/s).

5.3.5 Impact of the individual conduction speeds on the activation sequence

Table 5.5 reports the Pearson's correlation coefficient between activation maps simulated from the same parameter-set but changing a conduction speed's value. These results quantify the identifiability of each speed individually in the Eikonal model when modelling a healthy heart using a fast isotropic endocardial layer speed. The reference parameter-set considered the 'Normal speeds' configuration from Table 4.1, namely, 150 cm/s, 50 cm/s, 32 cm/s, and 29 cm/s for the endocardial, fibre-directed, sheet-directed, and sheet-normal directed speeds, respectively.

5.4 Discussion

This study presents the application of the inference method described in Chapter 4 to estimate the root nodes and conduction speeds from epicardial activation map data. My approach aims to serve as an efficient tool for generating cardiac 'digital twins', which is of paramount importance for precision cardiology (Corral-Acero et al., 2020). I conducted the simultaneous inference of endocardial, fibre, sheet,

and sheet-normal conduction speeds, and the location and number of the root nodes in the endocardium, as these properties determine the activation sequence in the ventricles. The evaluation of the method was conducted on a cohort of twenty virtual subjects (Section 4.3) to consider the effect of functional (five conduction speed configurations) and anatomical (four anatomies with biventricular myocardial-mass volumes ranging from 74 cm^3 to 171 cm^3) variability in the healthy human population on the performance of the inference process. I considered the noise-contaminated simulations as the ground truth for the inferences. The impact of the resolution of root node discretisation on the inference errors was also tested.

In addition to the novel methodology and its ability to simultaneously work with continuous and discrete parameters, the key findings of the study are as follows. Firstly, the method enabled finding populations of models that produced activation sequences with nearly identical patterns to those observed from the subject's epicardial activation time maps (Figure 5.2 and Table 5.2) while recovering the root nodes and the conduction speeds (Figure 5.3, Figure 5.4, Figure 5.5, Figure 5.6, Table 5.1, Table 5.3, Figure 5.7, Table 5.4). Secondly, the parameters that define the activation sequence are the root nodes, the endocardial speed, and the sheet speed. Thirdly, epicardial activation maps favoured the identification of non-septal root nodes.

A key contribution of this study is the development of an inference method to estimate the two types of parameters, namely, the anisotropic conduction speeds (continuous) and the number and locations of the root nodes (discrete) in both ventricles simultaneously. Previous work considered either a known number of root nodes (in simple cases, e.g. paced hearts) or known locations. For example, Giffard-Roisin et al., 2017a estimated the location of two root nodes in the left ventricle and the conduction speeds from ECG data; whereas, Grandits et al., 2020 estimated the activation times of a known set of root nodes and the conduction speeds from epicardial activation maps.

The evaluation of the inference methodology was conducted using synthetic data to know the ground-truth values for the ventricular activation properties as

these cannot be measured in healthy subjects. This allowed exploring the effects of the variability in tissue conduction properties in the healthy human population. I mimicked clinical conditions by contaminating the 'target data' using noise to demonstrate the feasibility of my method for clinical environments. In addition, I repeated the experiments using a biophysically detailed model (bidomain) to test the capability of my phenomenological models to perform the inference on synthetic data. This ability to reproduce the synthetic and noise-contaminated (realistic) data demonstrates that my inference methodology is relevant for clinical applications such as the generation of cardiac digital twins (Corral-Acero et al., 2020) towards the realisation of precision medicine.

All my results for the inference of conduction speeds (Figure 5.7) demonstrated that the endocardial and the sheet speeds were better identified than the fibre and sheet-normal speeds in healthy ventricles in sinus rhythm. When assuming healthy conditions, defined as root nodes in both ventricles and homogeneous tissue-conduction properties, as in Durrer et al., 1970, the endocardial and the sheet-directed speeds (alongside the root nodes) dominated the activation sequence patterns. This difference in the speeds' identifiability was due to the relatively negligible impact of the fibre and sheet-normal speeds on the activation sequence (Table 5.5). These speeds mostly act on the plane parallel to the endocardial layer, governed by a faster isotropic speed (i.e. endocardial speed). Conversely, the sheet directed speed defines the time delay for the endocardial-generated activation patterns to reach the epicardial surface since its orientation is perpendicular to the other speeds. Consequently, under healthy sinus-rhythm conditions, the activation properties that can be inferred from electrocardiographic data were the root nodes, the fast endocardial speed, and the sheet directed speed. However, I hypothesise that the fibre-directed speed will play a more relevant role in pathological conditions, for example, in the presence of slow conductance regions (e.g. scars and fibrosis) or wall thickening (e.g. hypertrophy) since this speed will have time to take over the endocardial patterns.

The inference of the root nodes from epicardial activation maps (Figure 5.3, Figure 5.4, Figure 5.5, and Figure 5.6) was better for root nodes in non-septal than in the septal areas. This was because septal root nodes have the least influence on the epicardial activation time map.

The inference results using noiseless bidomain-generated epicardial activation time maps as 'target data' (Section 5.3.4) demonstrated that the baseline ability to recover the conduction speeds is conditioned by the identifiability of the root nodes (Figure 5.8). In other words, higher resolutions allowed for both more accurate root node (Table 5.3) and speed (Table 5.4) estimations, highlighting the importance of the choice of resolution since it determines the baseline inference accuracy. These results also demonstrate that my inference method generalises to a biophysically detailed source model (bidomain), indicating suitability for clinical applications.

5.5 Limitations and future work

This study only considered sinus-rhythm activity in healthy subjects. Therefore, modifications to account for Purkinje-related abnormalities, such as bundle branch block, and tissue conduction heterogeneites, such as scars, would be required for the presented method to be implemented in the clinic.

I implemented the Purkinje network as an isotropic fast endocardial layer with root node locations that emulate the hotspots of Purkinje myocardial junctions. This strategy has been extensively used in previous studies (Wallman et al., 2014; Cardone-Noott et al., 2016; Lyon et al., 2018b) and shows equivalent behaviour in the activation maps from that observed from human data reported by Durrer et al., 1970. However, Durrer et al., 1970 also reported differences of about ten ms in the activation between root nodes located two cm apart from each other that my modelling strategy has not taken into account. A possible extension could include a realistic representation of the cardiac conduction system (Sebastian et al., 2013; Grandits et al., 2020), although this would add additional complexity and uncertainty to the inference problem. Note that the dominance of the endocardial and sheet speeds would prevail despite the inclusion of a Purkinje network, as in

Grandits et al., 2020, since the cause of this dominance is the parallel influence of a fast speed in the endocardial plane (e.g. fast endocardial speed or realistic Purkinje network) and the fibre and sheet-normal speeds. Another possible improvement could be to test the implications of using a different fibre orientation model (Bayer et al., 2012; Rossi et al., 2014; Doste et al., 2019).

The rule-based fibre orientation method of choice for this study was designed from canine experimental data (Streeter et al., 1969). Canine myocardial fibre structures hold a strong resemblance to human fibre orientations (Bayer et al., 2012). Nevertheless, considering a human-based rule-based fibre orientation (Rossi et al., 2014; Doste et al., 2019) would increase the clinical relevance of the methodology. Moreover, the choice of transmural angle change should be updated to match the values reviewed by Doste et al., 2019. However, the conclusions drawn in this study would remain unchanged regardless of the adopted strategy for the fibre orientations since most state-of-the-art methodologies, as reviewed in Piersanti et al., 2021, would produce a similar transmural direction for the sheet-directed speed, which together with the (fibre orientation invariant) endocardial speed, dominates the activation sequence in my results.

Epicardial activation time maps are not a standard clinical recording; thus, requiring an intermediate processing step, such as ECGI, for generating the input data to solve the inference problem. This shortcoming could delay the translation of the method into clinical applications until ECGI tools become a standard themselves.

Humanity needs dreamers, for whom the disinterested development of an enterprise is so captivating that it becomes impossible for them to devote their care to their own material profit.

— Marie Curie (Curie, 1937)

6

Inference of ventricular activation properties from 12-lead ECG-QRS

6.1 Introduction

The ECG is the most widely used clinical modality for cardiac disease diagnosis. Extending the inference of ventricular activation properties, presented in Chapter 5, to infer from standard 12-lead ECG data instead of epicardial activation maps would accelerate the wider adoption of the methodologies developed in this thesis. The goal of the study presented in this chapter is to investigate new computational techniques for the efficient quantification of subject-specific ventricular activation properties using CMR-based modelling and simulation and 12-lead ECG-QRS. This chapter presents and evaluates a 12-lead ECG-QRS extension of the inference methodology (Chapter 4). This extended inference composes the foundations of an efficient cardiac digital twin generation pipeline that considers two widely available clinical data modalities, namely, CMR and ECG.

6.2 Methods

As mentioned in Section 4.2, the inference from 12-lead ECG-QRS (Figure 4.1) implements the following five steps: simulate the activation maps using the Eikonal

model (Section 4.4.1); calculate the simulated QRS complexes using the pseudo-ECG algorithm (Section 4.4.2); compute the discrepancies between the simulated and target QRS complexes (Section 4.5); evaluate the two stopping criteria (Section 4.6); replace the parameter-sets with the highest discrepancies (Section 4.6).

Analogously to the previous chapter, the components of this inference pipeline, highlighted in Figure 4.1, are described in Chapter 4 together with the strategies to calculate the accuracy of the inference (Section 4.6.1) and to calibrate the hyperparameters (Section 4.6.2) of all these methodologies. The following sections in this Chapter present the evaluation of this extended inference method on the 12-lead ECG-QRS recordings from the twenty virtual subject cohort proposed in Section 4.3.

6.3 Results

6.3.1 Reproduction of noise-contaminated ECGs

Figure 6.1 illustrates the ability of my inference method to replicate the target ECGs. The replication of the Eikonal 12-lead ECGs was most accurate when considering the hybrid (Section 4.5.1) root node resolution (Figure 6.1.B). Moreover, the DTW-based discrepancy successfully allowed the inference method to replicate all leads with similar accuracies (Figure 6.1.A).

6.3.2 Root node inference from noise-contaminated ECGs

Figure 6.2, Figure 6.3, Figure 6.4, and Figure 6.5 illustrate the inferred root nodes locations from noise-contaminated 12-lead ECG-QRS for the virtual subjects with Mesh-1, Mesh-2, Mesh-3, and Mesh-4, respectively. Table 6.1 reports the summary error metrics for all virtual subjects.

The inference of root nodes was more accurate on the anterior side of the heart than on the posterior (Figure 6.2, Figure 6.3, Figure 6.4 and Figure 6.5) due to the left-anterior positioning of the precordial leads in the 12-lead ECG test. Increasing the resolution of the root node discretisation had a little visual effect on the distribution of root nodes identified from ECG (Figure 6.2, Figure 6.3, Figure 6.4 and Figure 6.5) compared to the inference from epicardial activation

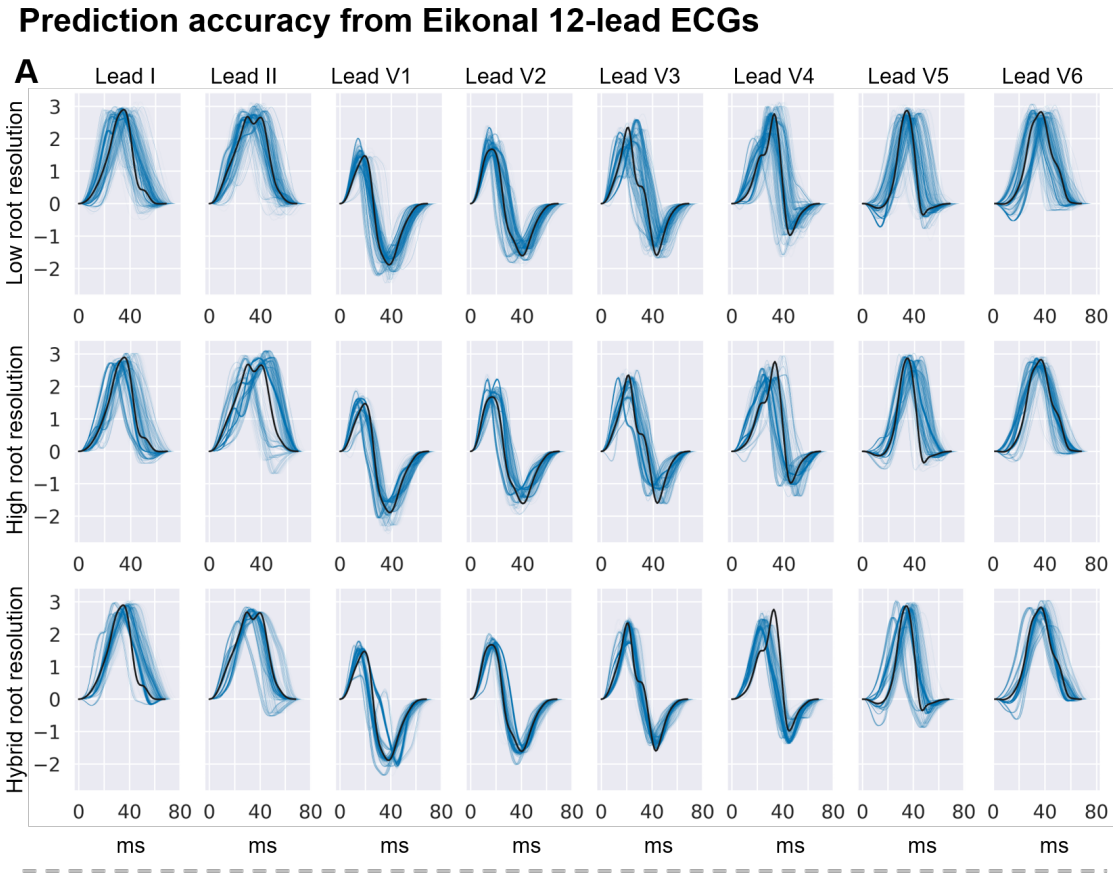

B

Anatomy	Low	High	Hybrid
Mesh-1	0.81 ± 0.21	0.78 ± 0.21	0.78 ± 0.21
Mesh-2	0.81 ± 0.20	0.76 ± 0.22	0.81 ± 0.19
Mesh-3	0.87 ± 0.14	0.86 ± 0.14	0.90 ± 0.11
Mesh-4	0.85 ± 0.16	0.83 ± 0.17	0.87 ± 0.15
All	0.84 ± 0.18	0.81 ± 0.19	**0.84 ± 0.17**

Figure 6.1: Agreement between predicted and Eikonal target ECGs from the inference process; (**A**) on the virtual subject with Mesh-4 and 'Slow endocardial speed' properties (Table 4.1); (**B**) mean ± standard deviation of Pearson's correlation coefficients. (**A**) shows the 'target' (black) and 'predicted' (blue) standardised QRSs for the inference guided by 12-lead ECG recordings. Each row accounts for different root node resolutions, from top-to-bottom: low (0.86 average Pearson's correlation coefficient), high (0.86), and hybrid (0.88) root node resolutions. Each plot includes the predictions from repeating the same experiment three times (512 QRSs each) to demonstrate the method's robustness. The amplitude of standardised QRS signals is unitless. (**B**) mean ± standard deviation of Pearson's correlation coefficients between the predicted and target ECGs. Low, High and Hybrid resolution of the discretisation of the root node parameter space.

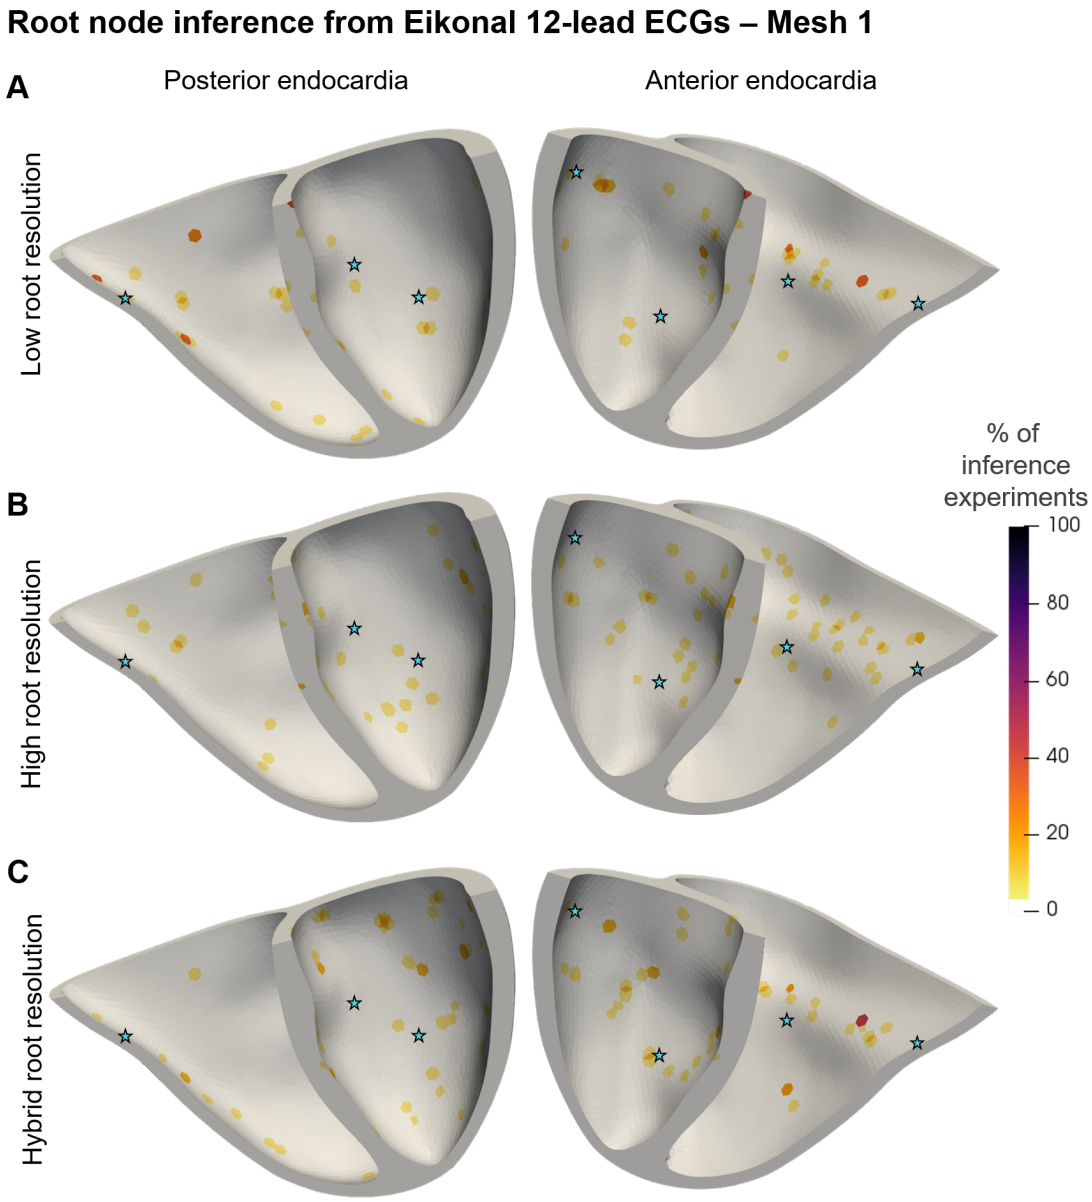

Figure 6.2: Root nodes inferred from noise-contaminated Eikonal 12-lead ECG-QRS on Mesh-1 with low (**A**), high (**B**), and hybrid (**C**) resolution of the root node space discretisation. The stars indicate the ground-truth root node locations. The endocardial surface is coloured as a heatmap showing the percentage of experiments that inferred each location as a root node.

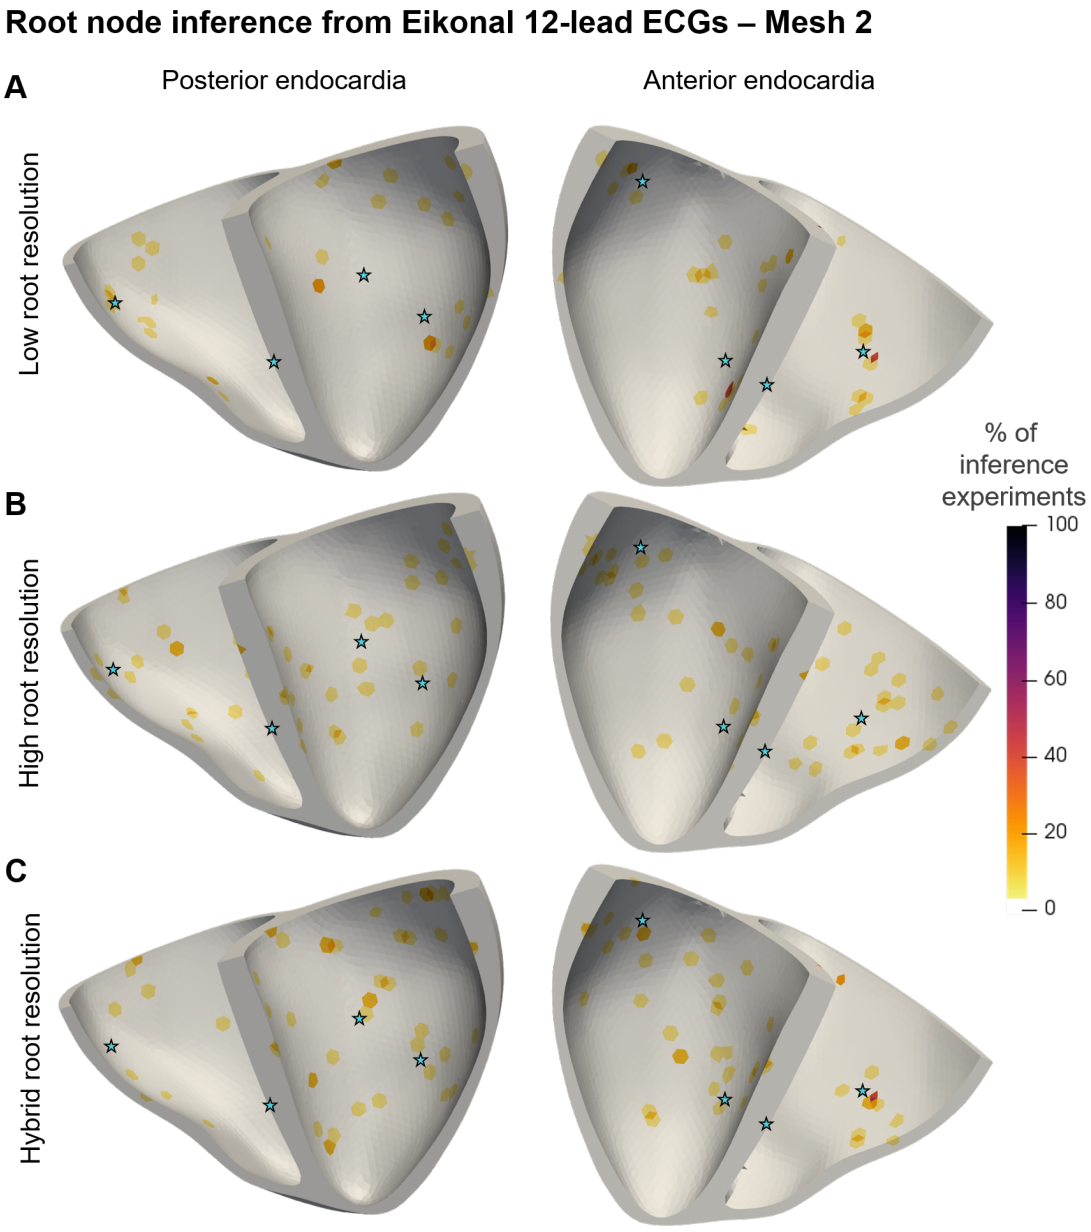

Figure 6.3: Root nodes inferred from noise-contaminated Eikonal 12-lead ECG-QRS on Mesh-2 with low (**A**), high (**B**), and hybrid (**C**) resolution of the root node space discretisation. The stars indicate the ground-truth root node locations. The endocardial surface is coloured as a heatmap showing the percentage of experiments that inferred each location as a root node.

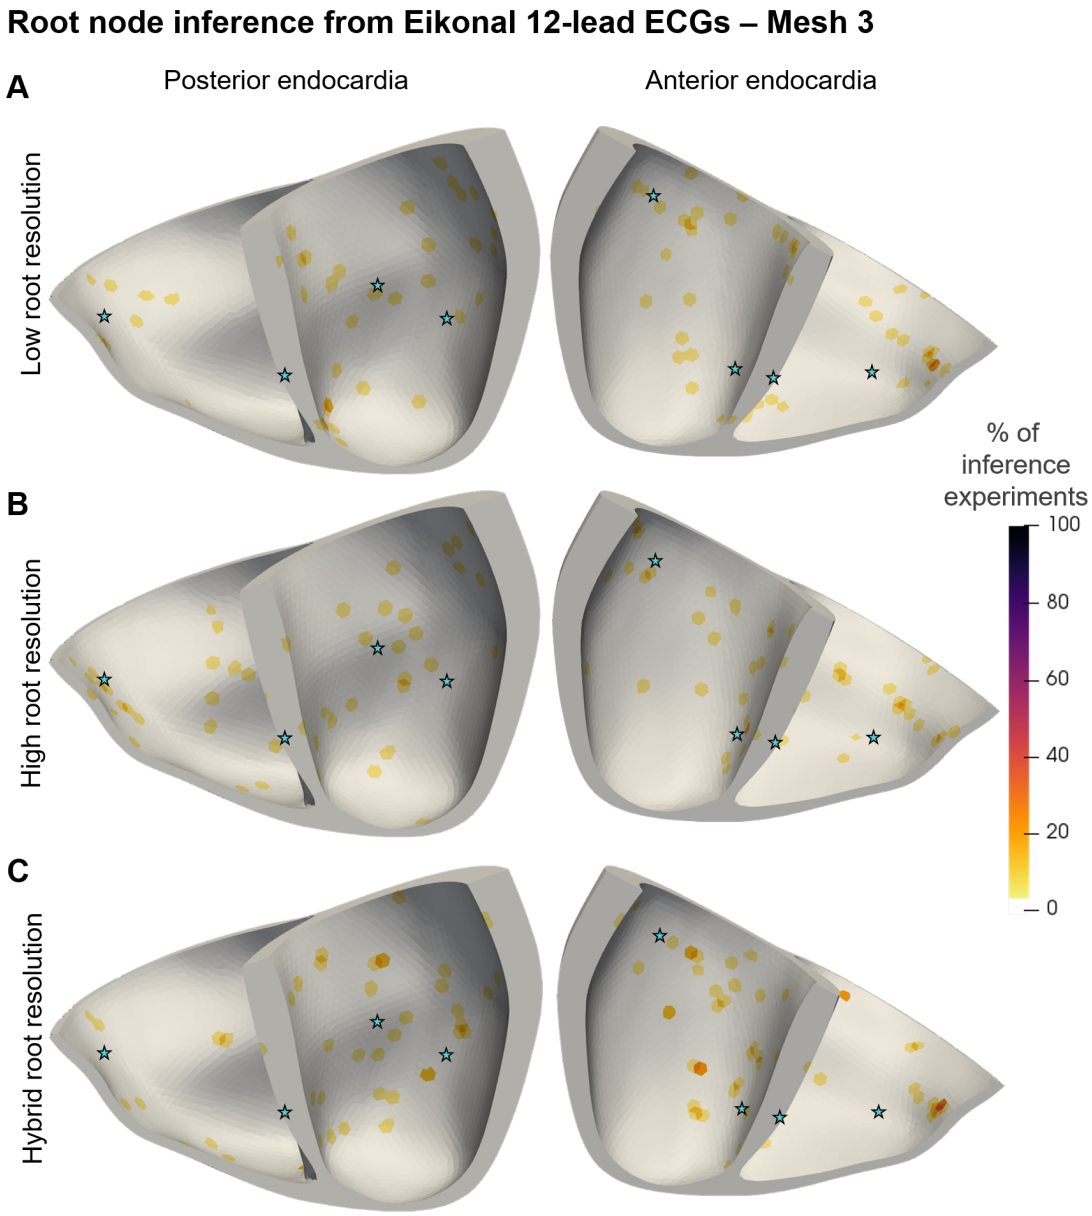

Figure 6.4: Root nodes inferred from noise-contaminated Eikonal 12-lead ECG-QRS on Mesh-3 with low (**A**), high (**B**), and hybrid (**C**) resolution of the root node space discretisation. The stars indicate the ground-truth root node locations. The endocardial surface is coloured as a heatmap showing the percentage of experiments that inferred each location as a root node.

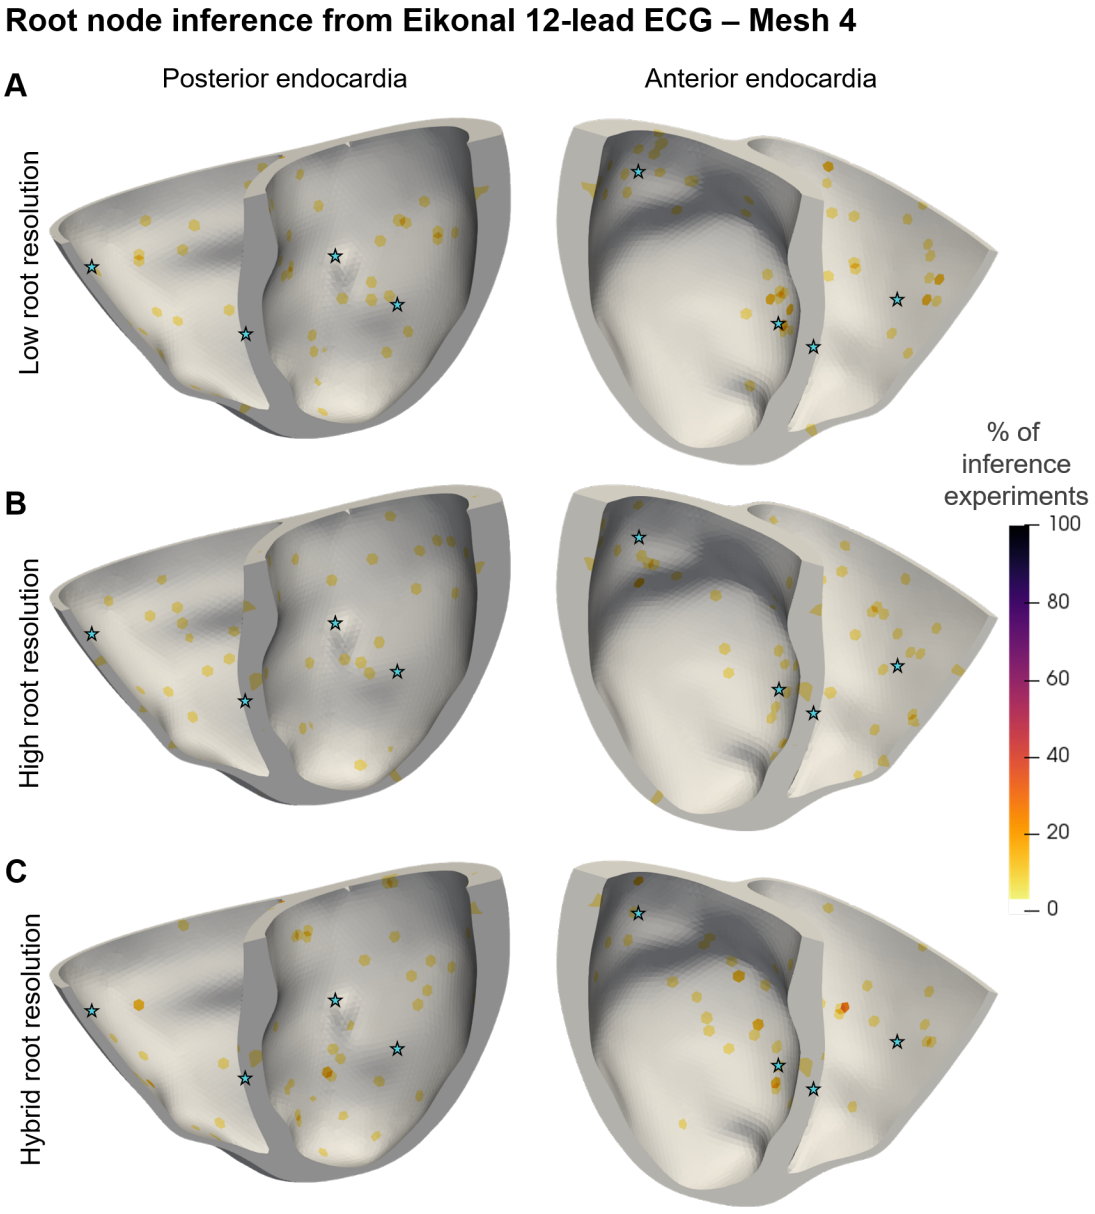

Figure 6.5: Root nodes inferred from noise-contaminated Eikonal 12-lead ECG-QRS on Mesh-4 with low (**A**), high (**B**), and hybrid (**C**) resolution of the root node space discretisation. The stars indicate the ground-truth root node locations. The endocardial surface is coloured as a heatmap showing the percentage of experiments that inferred each location as a root node.

Root node errors	Low	High	Hybrid
LV: Closest root node (cm)	1.9 ± 0.5	2.0 ± 0.6	1.8 ± 0.5
RV: Closest root node (cm)	1.7 ± 0.4	1.7 ± 0.4	1.9 ± 0.4
LV: Number of root nodes	0.6 ± 0.6	0.6 ± 0.5	1.0 ± 0.8
RV: Number of root nodes	1.6 ± 0.8	2.0 ± 1.1	0.8 ± 0.7

Table 6.1: Root nodes errors metrics (Section 4.6.1) from noise-contaminated Eikonal ECGs. Absolute mean ± standard deviation of the distance from the ground-truth and error in the number of root nodes inferred from all meshes. LV – left ventricle; RV – right ventricle; low, high, and hybrid resolutions of the discretisation of the root node locations.

maps (Figure 5.3, Figure 5.4, Figure 5.5 and Figure 5.6). Moreover, all three resolutions reported similar root node recovery accuracies from ECGs (Table 6.1). On average, the inference identified the root node locations 0.7 cm more accurately from epicardial activation maps (Table 5.1) than from ECGs (Table 6.1).

6.3.3 Speed inference from noise-contaminated ECGs

Analogously to Figure 5.7, Figure 6.6 reports the errors for the inference of conduction speeds (Equation 4.6.1) from ECGs. The results are shown for low (A), high (B), and hybrid (C) resolution of root node discretisation and grouped by torso-biventricular mesh in the virtual subject (see colour legend in Figure 6.6). The left side of the figure provides a numerical summary of the right side to ease readability.

Similarly to the inference from noise-contaminated epicardial activation maps (Figure 5.7), the conduction speeds inferred from ECGs (Figure 6.6) also demonstrated that the endocardial and sheet-directed speeds were recovered more accurately than the fibre and sheet-normal directed speeds (p-values < 0.0003, 59 samples). This finding suggests that the endocardial and sheet-directed speeds determined the activation wavefront's propagation speed (Section 5.3.5). The endocardial speed was recovered on average 2.6% more accurately from 12-lead ECG (Figure 6.6) than from epicardial activation maps (Figure 5.7). The hybrid root node resolution improved the recovery of the conduction speeds from 12-lead ECG (Figure 6.6), whereas, as I anticipated (Section 4.5.1), it played the opposite effect for the inference guided by activation maps (Figure 5.7).

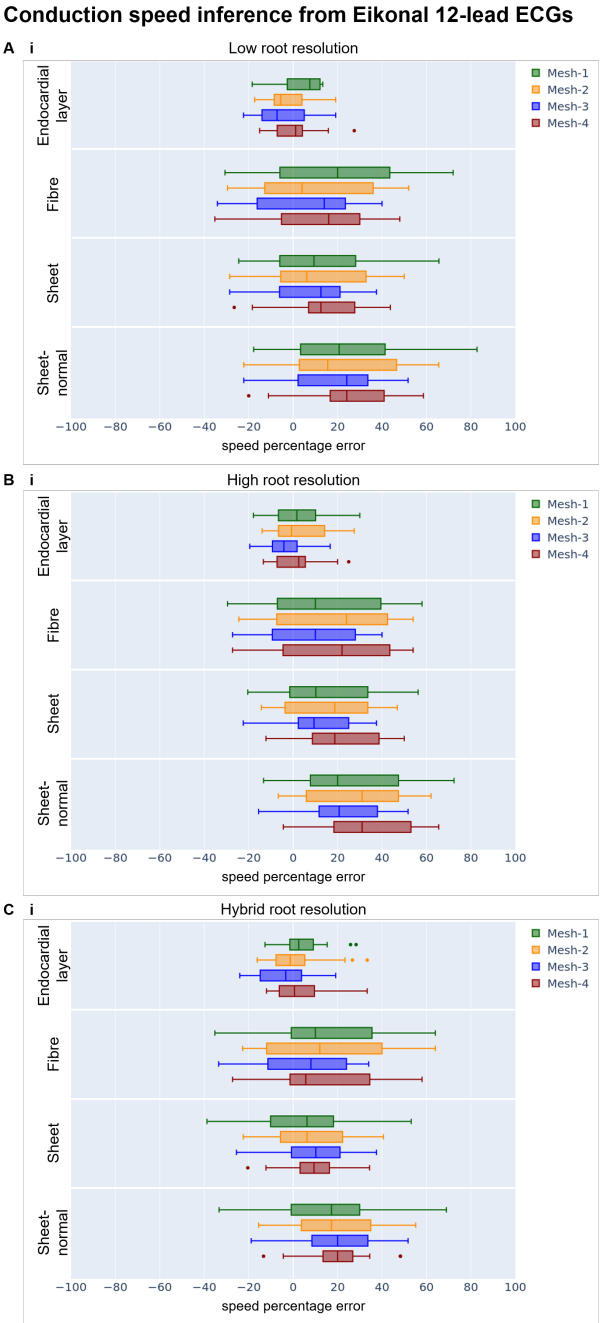

A ii

Conduction speed	Low root resolution
Endocardial	9.6 ± 6.2
Fibre	25.4 ± 16.0
Sheet	20.5 ± 14.5
Sheet-normal	27.7 ± 19.1

B ii

Conduction speed	High root resolution
Endocardial	9.6 ± 7.4
Fibre	25.3 ± 16.1
Sheet	21 ± 15
Sheet-normal	29.5 ± 19.9

C ii

Conduction speed	Hybrid root resolution
Endocardial	9.8 ± 8.6
Fibre	22.9 ± 17.5
Sheet	15.9 ± 12.1
Sheet-normal	23.0 ± 15.4

Figure 6.6: Error in the conduction speeds inference (Equation 4.6.1) from noise-contaminated Eikonal ECGs using low (**A**), high (**B**), and hybrid (**C**) resolution of the root node discretisation. The errors are computed as (Equation 4.6.1) the percentage over the ground truth conduction speeds (x-axis) represented as (**i**) box-plots grouped by anatomy (colour) and conduction speed (y-axis); (**ii**) and represented as absolute mean ± standard deviation.

Torso-biventricular anatomy	**Low**	**High**
Mesh-1	0.87 ± 0.15	0.92 ± 0.09
Mesh-2	0.91 ± 0.10	0.90 ± 0.10
Mesh-3	0.95 ± 0.06	0.96 ± 0.04
Mesh-4	0.93 ± 0.07	0.94 ± 0.07

Table 6.2: Prediction accuracy of the inference from bidomain 12-lead ECGs (mean ± standard deviation of Pearson's correlation coefficients). **Low** and **High** resolution of the discretisation of the root node parameter space.

Root node errors	**Low**	**High**
LV: Closest root node (cm)	1.8 ± 0.5	1.6 ± 0.5
RV: Closest root node (cm)	1.5 ± 0.5	1.6 ± 0.6
LV: Number of root nodes	0.7 ± 0.6	0.6 ± 0.9
RV: Number of root nodes	1.5 ± 1.2	0.9 ± 0.9

Table 6.3: Root node recovery error metrics (Section 4.6.1) (absolute mean ± standard deviation) from bidomain epicardial ECGs. Distance from the ground truth and error in the number of root nodes inferred. **LV** – left ventricle; **RV** – right ventricle; **Low** and **High** resolution of the discretisation of the root node parameter space.

6.3.4 Inference from noise-free bidomain ECGs

Increasing the resolution of the root node's discretisation consistently improved the accuracy for reproducing the 'target data' from noise-free bidomain simulations (Table 5.2 and Table 6.2). Conversely, uniformly increasing this resolution had had the opposite effect on the inference from noise-contaminated Eikonal ECGs (Figure 6.1.B).

The high resolution of the root node location discretisation also helped recover root node locations closer to the ground truth from ECGs (Table 6.3). In agreement with the results from Eikonal 'target data' (Table 5.1 and Table 6.1), the inference of the root node locations from bidomain data was consistently more accurate from activation maps (Table 5.3) than from ECGs (Table 6.3).

In agreement with my previous findings from Eikonal (Figure 5.7 and Figure 6.6) and bidomain 'target data' (Table 5.4), the endocardial and sheet speeds were also consistently better recovered than the fibre and sheet-normal speeds from bidomain ECGs (Table 6.4). Moreover, also in agreement with my previous results from bidomain epicardial activation maps (Table 5.4), increasing the resolution of

Conduction speed	**Low**	**High**
Endocardial	14.2 ± 14.3	13.9 ± 13.2
Fibre	28.4 ± 22.7	30.2 ± 24.2
Sheet	19.5 ± 14.8	14.0 ± 12.6
Sheet-normal	30.7 ± 13.7	34.1 ± 12.3

Table 6.4: Conduction speed recovery error metrics (Equation 4.6.1) (absolute mean ± standard deviation) from bidomain ECGs. **Low** and **High** resolution of the discretisation of the root node parameter space.

the root nodes' discretisation improved the estimation of the conduction speeds (Table 6.4), supporting my finding on the better identification of the root nodes enabling for a better estimation of the conduction speeds (Figure 5.8).

6.4 Discussion

This study extends the inference method I presented in Chapters 4 and 5 to recover the root nodes and conduction speeds from 12-lead ECG recordings instead of epicardial activation time maps. This required integrating the pseudo-ECG algorithm (Gima and Rudy, 2002) and a novel DTW-based ECG discrepancy metric (Section 4.5) into my original inference pipeline. I evaluated the performance of the extended method on the same cohort of twenty virtual subjects (Section 4.3) considered in my previous study (Chapter 5) for comparison.

Through this study, I presented an efficient tool for recovering the human ventricular activation properties directly from standard clinical data modalities (i.e. 12-lead ECG and CMR data). I show consistent results with those from epicardial activation maps (Figure 5.7 and Table 5.4) on the endocardial and sheet speeds having a more significant influence on the ECG than the fibre and sheet-normal directed speeds (Figure 6.6 and Table 6.4). I demonstrate that the identifiability of the root nodes from 12-lead ECG is better on the anterior side of the heart than the posterior side. Finally, I observe that the inference from epicardial activation maps (Section 5.3) was more accurate than the inference from ECGs (Section 6.3).

The inference of root nodes from 12-lead ECGs (Figure 6.2, Figure 6.3, Figure 6.4, and Figure 6.5) demonstrated that the anterior locations were better identified than

the posterior ones. This asymmetric identifiability of the root nodes manifests the asymmetry of the electrode positioning protocol in the standard 12-lead ECG test. More precisely, each root node has a local effect on the region's activation sequence on the heart that the ECG represents. However, the amplitude of the recorded regional-ECG is inversely proportional to the distance between that region and the electrode, while most electrodes are typically positioned on the left-anterior surface of the torso. Consequently, most changes in the activation sequence patterns in the posterior walls get masked by the electrical activity taking place in the anterior half of the organ for the precordial electrograms.

Overall, I attained higher accuracy and consistent predictions from using epicardial activation maps (Section 5.3) for the inference of the root nodes than from ECG recordings (Section 6.3). The activation map's main advantage is that the data are spatially distributed on a surface (epicardium) that strongly relates to the root node parameter space (except for the septal area). More precisely, the inference from epicardial activation maps can be subdivided into regional sub-problems as local root nodes only influence local patterns in the activation sequence. This phenomenon enables the algorithm to quickly identify partial solutions to the inference problem from epicardial activation maps. On the other hand, all root nodes influence each signal's morphology in the ECG simultaneously.

6.4.1 Limitations and future work

This study accounted for the same limitations as my previous version of the inference (Section 5.5) except for the need for an intermediate tool to generate the input data for my extended inference pipeline since 12-lead ECG is the most widely used clinical test for cardiac diagnosis. However, the translation of my extended ECG-guided inference methodology into clinical applications will require an automatic QRS multi-lead delineation method to preprocess real ECG recordings before they can be used to guide the inference. Moreover, the exact location of the electrodes used during 12-lead ECG tests are rarely recorded in clinical practice and clinical QRS mythologies can display beat-to-beat variability. Thus, I suggest the use

of approximate electrode locations using general guidelines, as in Cardone-Noott et al., 2016, and a beat averaging strategy, as in Lyon et al., 2018a to adapt the current method for working with clinical 12-lead ECG recordings. Additionally, the inference method should account for the effect of the uncertainty in the exact positions of these electrodes (Kania et al., 2014).

Overall, this study and the former (Chapter 5) presented the foundations of a novel pipeline capable of non-invasively calibrating cardiac 'digital twins' for healthy subjects from synthetic epicardial activation maps (Chapter 5) and 12-lead ECG recordings (Chapter 6). My approach was designed to easily accommodate disease conditions, such as pathological tissue heterogeneities (scars or fibrosis). Moreover, the results using bidomain-generated 'target data' to guide the inference (Section 5.3.4 and Section 6.3.4) suggest that the presented approach can be directly translated to work on different source models, for example, biophysically detailed simulations or real clinical recordings. Consequently, I anticipate that my methodology can be adapted to work with real ECG recordings from disease patients, as in Giffard-Roisin et al., 2017a, to generate cardiac 'digital twins' in daily clinical practice.

Profound study of nature is the most fertile source of mathematical discoveries.

— Joseph Fourier (Fourier, 1878)

7

Deep learning QRS delineation in multi-lead clinical electrocardiograms

7.1 Introduction

The previous chapter (Chapter 6) presented the inference of the ventricular activation properties from 12-lead ECG-QRS complexes. As outlined in Section 6.4.1, the extraction of QRS complexes requires automatically delineating the ECG. The automatic delineation of ECG recordings is a challenging task (Section 2.4.2) that remains partially unsolved (Section 3.2). The state-of-the-art ECG delineation systems (Martínez et al., 2004; Bote et al., 2018) require experts to manually adapt their rules to the data's characteristics (Section 3.2) and display significant variability in their results. Deep learning algorithms can automatically adapt to the data without the need for manual calibration and can generalise across datasets, producing robust predictions (Section 2.8). This chapter presents and evaluates a novel deep learning-based ECG-QRS multi-lead delineation approach to accurately extract the QRS complexes in clinical ECG recordings. The integration of this deep learning-based delineation approach with the previously presented QRS-guided inference methodology will help their translation to clinical applications.

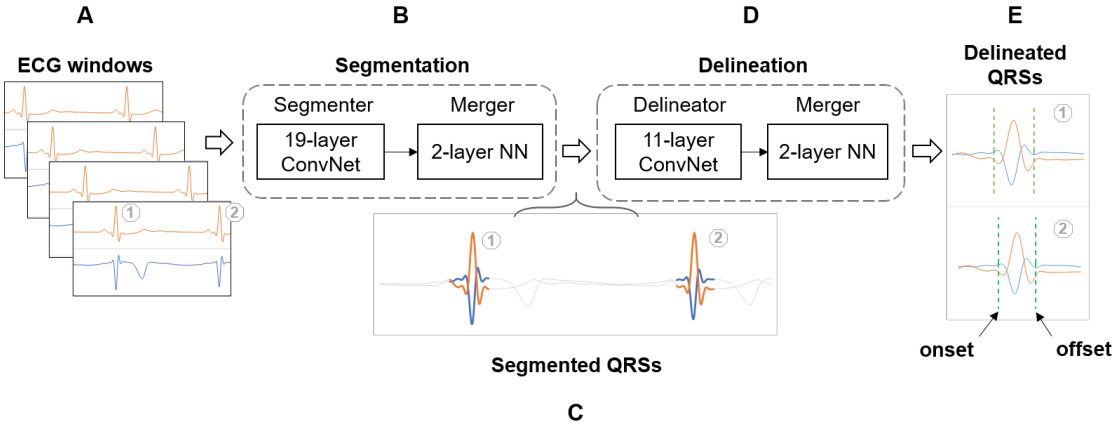

Figure 7.1: Diagram of the deep learning-based ECG multi-lead delineator. From left-to-right: (**A**) Windowed multi-lead ECG signals; (**B**) Segmentation module (i.e. Segmenter + Merger); (**C**) segmented ECGs (intermediate output; bottom - centre); (**D**) Delineation module (i.e. delineator + merger); and (**E**) delineated QRSs. Abbreviations: ConvNet - convolutional neural network; NN - fully-connected neural network; onset - QRS onset; offset - QRS offset.

7.2 Methods

7.2.1 Overview

Figure 7.1 illustrates the proposed delineation system. This deep learning-based multi-lead QRS delineation segments (Figure 7.1.C) and then delineates (Figure 7.1.E) the QRS complexes from an ECG signal. This approach is formed by successive Segmentation (Figure 7.1.B) and Delineation (Figure 7.1.D) modules, each of which is composed of a 1D convolutional neural network (Section 2.8) and a fully connected neural network.

7.2.2 Clinical ECG recordings

I considered the QT database (Section 3.2) (Laguna et al., 1997) to evaluate the delineation method presented in Figure 7.1 as it contains recordings from seven databases recorded in different medical centres from patients with different cardiac conditions. The QT database includes 105 fifteen-minute ECG recordings with two leads at a 250 Hz sampling rate. These recordings come from six pre-existing databases, namely, the MIT-BIH Arrhythmia database (Moody and Mark, 2001), the MIT-BIH ST Change database (Albrecht, 1983), MIT-BIH Supraventricular

Arrhythmia database (Greenwald and Patil, 1990), the MIT-BIH Normal Sinus Rhythm database (Goldberger et al., 2000), the European ST-T database (Taddei et al., 1992), and the MIT-BIH Long-Term ECG database (Goldberger et al., 2000). Additionally, the QT database includes ambulatory ECG recordings from patients who suffered sudden cardiac death (during the recording) and control subjects that matched the former in age and sex.

The QT database includes delineation annotations for 3622 beats, a subset of which was annotated a second time independently to enable intra-annotator variability analysis. Moreover, all annotations were audited to correct potential inconsistencies (Laguna et al., 1997; Goldberger et al., 2000). The QT database remains the standard benchmarking database for validating ECG delineation approaches.

7.2.3 ECG pre-processing and data augmentation

As previously mentioned, the data considered for this study were the QT database ECG recordings. This database contains 3600 delineated beats within its 2-lead ECG signals sampled at 250 Hz. Deep learning strategies can operate on raw data; however, I preprocessed the data according to domain knowledge to aid the machine learning algorithms to distinguish relevant information. I also implemented data augmentation techniques to maximise the information learned from the data available (e.g. using overlapping windowing) and others to teach the algorithm-specific rules (e.g. duplicating the recording but with the order of the leads swapped).

I filtered the ECG signals using a high-pass filter as baseline wander (0.5 Hz) and a low-pass filter at 50 Hz (Section 2.4.1) (Sörnmo and Laguna, 2005) to remove non-QRS related information. I windowed the ECG recordings at two different lengths to train the deep learning algorithms. I used a 2.048 s (i.e. 512 instances or time-steps per lead) long window for the Segmentation module (Figure 7.1.B) and a 0.256 s (i.e. 64 instances) long window for the Delineation module (Figure 7.1.D).

The 2.048 s (i.e. 512 instances) window size for the Segmentation module (Figure 7.1.B) ensured that there would be at least one QRS complex in each segment according to RR-interval physiological ranges for humans. Whereas the

0.256 s (i.e. 64 instances) window accounted only for one QRS complex by itself since the Delineation module (Figure 7.1.D) uses segmented QRSs (Figure 7.1.C) produced by the Segmentation module (Figure 7.1.B) as input.

Once windowed, each ECG segment was standardised (i.e. subtract its mean and divide by its standard deviation) using its mean and (sample) standard deviation. Hence, the presence of at least one QRS complex in each window generated for the Segmentation module (2.048 s) ensures having similar amplitude scales for all signals across different segments and leads after standardising the data.

Then, I augmented the data by using overlapping windows. Namely, the following window started one instance (time-step) to maximise the information learned from each recording. I shifted and scaled the ECG signals by small amounts (random $\pm$ 0.3 and random 1. $\pm$ 0.25, respectively) in each lead independently, augmenting the data about 100 times. This forced the algorithm to learn patterns invariant to scaling or shift differences, which would be in the data due to the standardisation of each segment individually. Lastly, I duplicated each recording with the order in the leads swapped and with the same lead in both positions (in all combinations possible) to teach the approach to delineate the signals regardless of the leads shown or their order in the recordings.

7.2.4 Deep learning architecture

The deep learning architectures composing our delineation pipeline are hereafter outlined following Figure 7.1 from left-to-right. As previously outlined, the pipeline is composed of two modules, a Segmentation module (Figure 7.1.B) and a Delineation module (Figure 7.1.D), each of which includes two deep learning architectures. The Segmentation module (Figure 7.1.B) is formed by a Segmenter architecture (Figure 7.2) and a Merger architecture.

The Segmenter (Figure 7.2) is a 19-layered 1D convolutional neural network composed of sixteen convolutional and three fully connected layers. These 19 layers enabled the algorithm to learn 132 ms (33 instances) long patterns, which is approximately the duration of a healthy QRS complex. Each convolutional

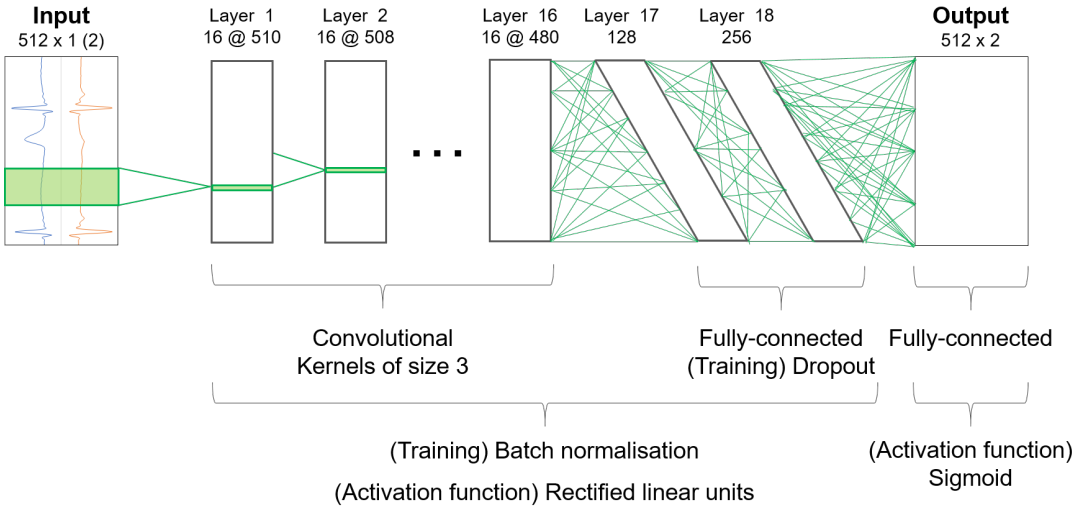

Figure 7.2: Diagram of the Segmenter (in Figure 7.1.B). The Segmenter is a 19-layered 1D convolutional neural network composed of sixteen convolutional and three fully-connected layers, implementing dropout (Section 2.8.2) at 0.5 (during training). All layers implement rectified linear unit (Section 2.8.1) as the activation function - except the last layer that implemented the sigmoid activation function (Section 2.8.1) - and batch normalisation (Section 2.8.3) (during training).

layer is formed by sixteen kernels of length three. These layers implement batch normalisation (Section 2.8.3) and use rectified linear unit activation functions (Section 2.8.1). These are followed by the two 'hidden' (i.e. not output) fully-connected layers, which have 128 and 256 neurons, respectively. These hidden layers also implement batch normalisation, rectified linear units as the activation functions, and dropout (Section 2.8.2) at a rate of 0.5 (i.e. each weight has a 50% chance of being dropped during training). Finally, the output layer was another fully connected layer with as many sigmoid units (Section 2.8.1) as the size of the input window to enable the model to perform classification from one-hot encoded labels.

The second model of the Segmentation module is the segmentation's Merger. This model is a neural network composed of two fully connected layers, with 512 neurons and 1 neuron in the first and second layer, respectively. The first layer implements batch normalisation, dropout at rate 0.5 and rectified linear unit as the activation functions. The second layer is a sigmoid neuron (output layer).

Analogously to the Segmentation module, the Delineation module is also composed of a convolutional (Delineator) and a fully-connected (Merger) architectures

(Figure 7.1). The Delineator, similarly to the Segmenter (Figure 7.2), is also a 1D convolutional neural network composed of the same elements as the Segmenter (Figure 7.2) except for having eight convolutional layers with kernels of size five, instead of the Segmenter's sixteen layers of size three.

Lastly, the delineation's (module) Merger is identical to the Segmentation's one, namely, two fully-connected layers with 512 neurons and 1 neuron, respectively.

The number of layers and kernel sizes in the Segmenter and the Delineator architectures were chosen to enable the algorithms to learn patterns with duration up to 132 ms (i.e. 33 instances), which physiologically accounts for most human QRS complexes. QRS delineation is a challenging task, and the QT database is a relatively small database (within deep learning standards). Thus, I kept the number of layers of all deep learning architectures to a minimum to ease their training process (Goodfellow et al., 2016).

7.2.5 Training of the deep learning algorithms

Once pre-processed, the windowed and augmented ECG recordings are split into training, validation, and test (set aside) datasets. The training data serve to adjust the weights through backpropagation, and the validation data serves to evaluate the generalisation capability of the algorithm during training and prevent overfitting.

I trained the Segmentation module (Figure 7.1.B) and the Delineation module (Figure 7.1.D) separately to avoid fitting the latter to the (during) training error of the former. Hereafter I describe the training process for the Segmentation module (Figure 7.1.B).

I labelled the data using a one-hot-encoding strategy to train the Segmentation module. One-hot-encoding labelling gives binary labels to each instance (i.e. time-step) within the recording segment. I defined the instances within the QRS as the positive class and the negative class elsewhere. This labelling strategy aids the algorithm to learn the morphology of the QRS with respect to other morphologies (e.g. T-wave) in the recording.

Firstly, I trained the Segmenter algorithm to identify the QRS instances within ECG recordings by using the augmented data in a completely randomised order and regenerating the data augmentation continuously. Secondly, I froze the trained weights of the Segmenter and progressed to train the Segmentation module's Merger architecture.

The Merger's purpose was to act as an ensemble learning unit. Namely, it was trained using the (already trained) Segmeter's outputs to correct its prediction errors. When windowing the data, some QRSs got cut at the start or end of the segment, preventing the Segmenter to observe the full complex and, thus, lowering its accuracy. Therefore, the Merger (in the Segmentation module) was trained using two-window long ECG segments (Figure 7.1.C). These segments were internally windowed (regular window-length), sliding one instance at a time (as shown in Figure 7.1.A), and fed to the Segmenter. The Segmenter produced a prediction for each instance in each segment (Figure 7.1.C). These predictions were stacked as a matrix, where the first instance would only have one prediction from the first window, but the second one would have two, then three, and successively until the middle of the matrix where there would be 512 predictions per instance and then monotonically decreasing again back to one. The Merger used this matrix of stacked predictions as input data to determine an ensemble prediction for each instance in the window at the centre of the two-window long ECG segment. This padding ensured that the Merger worked with predictions generated (by the Segmenter) from complete QRS complexes.

On the other hand, as outlined in Section 7.2.5, I trained the Delineation module (Figure 7.1.D) using the 0.256 s (i.e. 64 instances) long windowed ECGs (Figure 7.1.E) that contained only one QRS complex plus some additional signal as padding. These windows were computed using the ground-truth labels in the data (during training). Moreover, besides the data augmentation strategies mentioned for training the Segmentation module, these shorter windows were also augmented by displacing the QRS complex within the window by random amounts so that it

would not always be in the centre of the window to emulate the effect of potential segmentation errors that were likely to be produced by the Segmentation module.

These single-QRS windows were relabelled to represent the delineation task. Only the onset and offset regions of the QRS were set as the positive class, whereas other parts of the signal, including the R-peak, were set as the negative class. However, whereas the labelling strategy for the Segmentation module was binary (one-hot-encoding), I used regression-type labelling for the Delineation module. More precisely, I set the exact location where the onset and offset were as one. Then, the surrounding instances were given labels greater than zero but decreasing as they distanced from the exact location. This should aid the learning process by positively rewarding small misplacement of the delineation marks of the QRS while not rewarding the algorithms identifying the R-peak region for wide QRS complexes. In other words, the algorithms had to learn the patterns defining the onset and offset of these complexes and identify nearby positions to lower their prediction errors.

Analogously to the Segmenter, the Delineator algorithm was trained first on the (just described) augmented data segments. Then, I froze the Delineator's weights and used its predictions to generate input matrices for the Merger (in the Delineation module) to produce ensemble predictions.

The weights for the deep learning algorithms in the Segmentation (Figure 7.1.B) and Delineation (Figure 7.1.D) modules were adjusted using Adam (Kingma and Ba, 2014b) and binary cross-entropy.

As previously mentioned, the data considered in this study (i.e. the QT database) contained ambulatory ECG recordings from seven different databases. Therefore, I trained the deep learning-based delineation method using leave-one-database-out nested cross-validation (Section 2.8.3). Adopting the nested extension combined with grid search enabled automatically adjusting hyperparameters (e.g. learning rate). Nevertheless, most hyperparameters (e.g. architecture, window sizes, and dropout rate) were chosen according to domain knowledge about ECG signals and their waves. Furthermore, we used Adam (Kingma and Ba, 2014b) and binary cross-entropy.

7.2.6 Delineation process and its evaluation

As just mentioned in Section 7.2.5, I evaluated the deep learning-based delineation method using leave-one-database-out nested cross-validation (Section 2.8.3). The adoption of the leave-one-database-out extension enabled assessing the generalisation capability of the delineation method across databases and cardiac conditions. This evaluation represented the performance of the method if translated into a clinical application. Moreover, the training and assessment were repeated ten times to ensure the reproducibility of the results.

As illustrated in Figure 7.1, the process started with the windowing of the ECG recordings with half-window signal of padding in front and behind (Figure 7.1.A) the current window. The QRSs in the window were segmented by the Segmentation module (Figure 7.1.B). I used an inner-segment sliding window (as during training in Section 7.2.5) and the Segmenter to produce a matrix of predictions. Then, the Merger computed the final segmentation for each QRS (Figure 7.1.C) in the original window (without padding) from that matrix as an ensemble unit.

Then, I added some padding to the segmented QRS complexes to allow inner-segment windowing while ensuring that no P-waves or T-waves are included. Next, I inputted the segments one at a time into the Delineation module (Figure 7.1.D). Analogously to the Segmentation module, I also used the Delineator algorithm with an inner-segment sliding window to produce a matrix of predictions, which was then combined into a single onset and offset prediction (Figure 7.1.E) by the Merger.

7.2.7 Error metrics and post-processing

The predictions from the delineation method were regression-like predictions per each instance in the input window. However, calculating the performance of our method required exact QRS onset and offset positions. Thus, I calculated these positions within each window as the maximal density prediction points, using a Gaussian-shaped weighted averaging strategy. Finally, I transformed these locations into time values to compute error metrics.

For this study, I reported the mean ± (sample) standard deviation for the root mean square error (RMSE) (7.1) to measure the performance of the delineation method, and the QRS miss rate using the Segmentation module's (Figure 7.1.B) intermediate predictions (Figure 7.1.C) as QRS detection. Moreover, I also computed the (signed) mean error (ME) (7.2) and the error's sample standard deviation (STE) (7.3) for comparison with the literature (Martínez et al., 2004; Bote et al., 2018).

The equations for these error metrics are

$$\text{RMSE} = \sqrt{\frac{\sum_{i=1}^{n}(\hat{y}_i - y_i)^2}{n}} \tag{7.1}$$

$$\text{ME} = \frac{\sum_{i=1}^{n}(\hat{y}_i - y_i)}{n} \tag{7.2}$$

$$\text{STE} = \sqrt{\frac{\sum_{i=1}^{n}(\hat{y}_i - y_i - \text{ME})^2}{n-1}} \tag{7.3}$$

where n is the number of samples and, $\hat{y}_i$ and y_i are the predicted and the real labels of the ith sample, respectively.

7.3 Results and discussion

This study proposes a deep learning-based multi-lead delineation system for QRS complexes in ECG signals. This section presents and discusses the evaluation of the delineation method on the QT database using a leave-one-database-out strategy. The (mean ± standard deviation) root mean square error of the deep learning-based delineation approach were 12.1 ± 0.5 ms and 18.5±1.1 ms for the QRS onset and offset. Moreover, the root mean square error calculated from the differences between the marks given by the two ECG experts in this database were 14.7 ms for the QRS onset and 17.2 ms for the offset, demonstrating that our errors were within the range of inter-cardiologist variability.

For comparison with the literature, Table 7.1 reports the (signed) mean error ± error standard deviation of the deep learning-based delineation method compared to the Wavelet transform-based delineator by Martínez et al., 2004 and the low complexity thresholding-based delineator by Bote et al., 2018 (both reviewed in Section 3.2) on the QT database.

Delineation method	**QRS onset**	**QRS offset**
Martínez et al., 2004	4.6 ± 7.70	0.8 ± 8.70
Bote et al., 2018	7.0 ± 4.30	-5.0 ± 9.90
Deep learning-based	-2.6 ± 10.8	4.4 ± 15.2

Table 7.1: Comparison of the performance of the delineation approach to Martínez et al., 2004 and Bote et al., 2018 on the QT database. Signed mean ± standard deviation of the error in milliseconds (ms) for the onset and offset of the QRS complex.

Martínez et al., 2004 and Bote et al., 2018 proposed single-lead delineation methods (Section 3.2). Therefore, they produced two predictions for each QRS onset and offset on the QT database as it includes two-lead ECG recordings. For their reported metrics to be directly comparable to mine, Martínez et al., 2004 and Bote et al., 2018 should have combined their two individual predictions before computing the error metrics. However, they first compute the error of each prediction (QRS onset or offset) and then use the most accurate of the two to compute their metrics (Martínez et al., 2004). Despite not invalidating their conclusions, always choosing the best prediction can significantly increase the measured performance if the two predictions are different from one another. Thus, without further information, a direct comparison between our delineation methods is not possible. Nevertheless, our performances (Table 7.1) were similar enough to conclude that my deep learning-based delineation method was comparable to the state-of-the-art for QRS delineation from ECG recordings in new databases.

Overall, the data augmentation strategies combined with the division of the responsibility of the delineation task into a segmentation and a delineation sub-tasks enabled the algorithms to successfully learn how to delineate QRS complexes in new databases from a scarce and highly diverse set of two-lead ECG recordings. The delineation approach demonstrated human-level performance by achieving errors within the range of the inter-cardiologist variability computed from the QT database.

7.3.1 Limitations and future work

I considered the QT database to train and evaluate the multi-lead deep learning-based ECG-QRS delineation method. However, the delineation marks included

in this database were generated using only one of the two leads in the ECG recordings. The chosen lead was different for different recordings and the two leads in each recording were different subsets from the twelve leads in the 12-lead ECG test. These inconsistencies in the data prevented the algorithm from learning optimal patterns. Thus, I expect that retraining the delineation method on a new ECG-QRS delineation database that has the twelve leads in the standard ECG tests independently labelled would increase the method's performance by enabling it to learn the optimal QRS delineation features while effectively exploiting information from multiple leads.

Wide QRS complexes (duration $\geq$120 ms) (Das et al., 2008) can be indicative of Purkinje-related pathological conditions, such as bundle branch block (Sarubbi et al., 2000). The convolutional architectures included in the delineation pipeline can learn 132 ms long patterns. These patterns will be long enough to capture and represent full QRS complexes in most cases, even under disease conditions. However, increasing the number of layers to account for up to 200 ms long patterns would ensure that the method has the capacity to represent almost any individual in the human population.

This is only a foretaste of what is to come and only the shadow of what is going to be. We have to have some experience with the machine before we really know its capabilities. It may take years before we settle down to the new possibilities, but I do not see why it should not enter any one of the fields normally covered by the human intellect and eventually compete on equal terms.

— Alan Turing 1949

8 Conclusions and future directions

8.1 Contributions

This thesis presents and evaluates an inference method (Chapter 4) combined with CMR-based torso-biventricular Eikonal models to estimate the root nodes and conduction speeds from 12-lead ECGs (Chapter 6) or epicardial activation maps (Chapter 5). It also presents a multilead QRS delineation system for clinical ECG recordings (Chapter 7). These methodologies aim to serve as the foundations of an efficient tool for generating cardiac 'digital twins' (Corral-Acero et al., 2020) in the clinic.

I evaluated my inference method on a twenty virtual subject cohort representing the healthy human population's variability (Section 4.3) and the QRS deep learning-based delineation system on 105 clinical ambulatory two-lead ECG recordings from the QT database (Laguna et al., 1997). Combining the inference and delineation approaches will enable the translation of my work to clinical applications using standard clinical data modalities (i.e. 12-lead ECG and CMR).

The other main contributions of this thesis are the following:

- I present an independent sampling strategy for mixed continuous and discrete parameter spaces (Section 4.6).

- I propose a DTW-based distance metric for ECG comparison that highlights differences originated from the root node locations and conduction speeds in the subject's heart (Section 4.5).

- I demonstrate the relevance of the different fibre-directed speeds in the cardiac activation sequence and the relevance of the discretisation of the root node parameter space (Chapter 5 and Chapter 6).

8.2 Future directions

Here I propose three possible continuations to generalise my methodologies to account for healthy and cardiac disease conditions simultaneously. The resulting pipeline will automatically identify and represent complex cardiac conditions, such as combinations of scars and abnormalities in the Purkinje network using only non-invasive standard clinical recordings.

8.2.1 Motivation

Cardiac diseases are the most common non-communicable disease, accounting for 17.8 million deaths in 2017 (Kaptoge et al., 2019). These diseases manifest functional and anatomical alterations that promote lethal arrhythmias. The methodologies developed and evaluated for this thesis can generate cardiac digital twins (Corral-Acero et al., 2020) for healthy subjects by inferring the properties in the subject's heart that cannot be clinically measured from routinely acquired clinical data. Moreover, my methodologies can be extended to account for pathological tissue conduction heterogeneities (Antzelevitch, 2007) and structural abnormalities in the Purkinje network (Ideker et al., 2009; Haissaguerre et al., 2016) to treat patients at risk of arrhythmic death.

Thus, I propose to expand my work to account for high-risk (e.g. electrical conduction heterogeneities) and poorly diagnosed (e.g. Purkinje network abnormalities) arrhythmic conditions. The outcome of this work will be the first efficient end-to-end automatic cardiac digital twin generation pipeline.

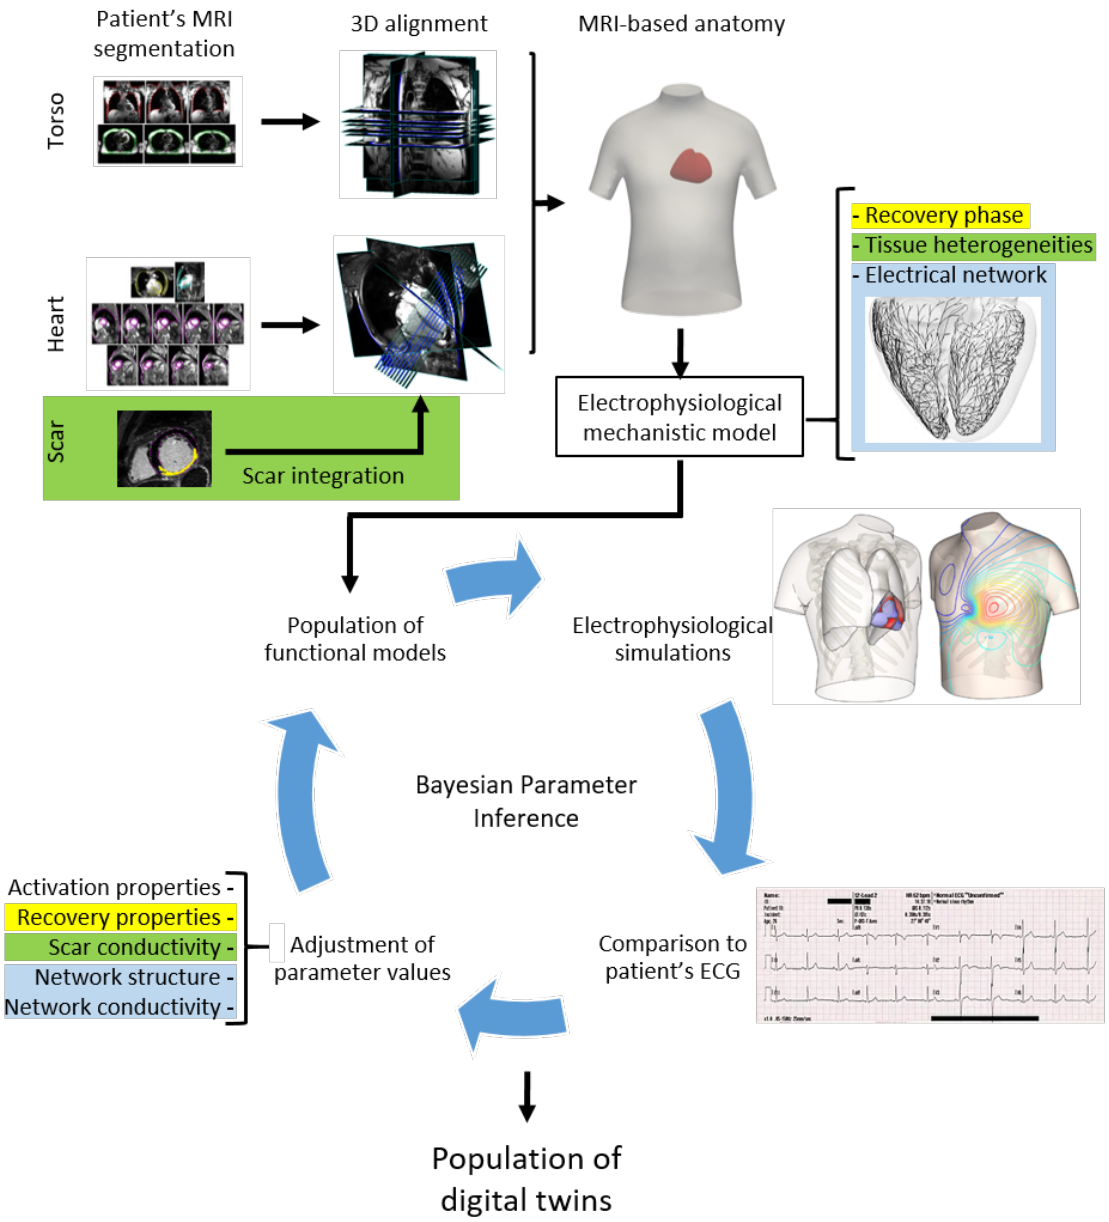

Figure 8.1: Extended automatic cardiac digital twin generation pipeline. The three proposed extensions to the current pipeline are colour-coded as follows: recovery phase – yellow, tissue heterogeneities – green, and realistic Purkinje network – blue.

8.2.2 Proposed investigations

The proposed cardiac digital twin generation pipeline produces a population of digital twins to account for how more than one condition can reproduce the same behaviour in the ECG and CMR recordings.

To extend the pipeline to consider healthy and disease patients by including the recovery phase of the cardiac cycle (Figure 8.1 – yellow).

The ability to reproduce arrhythmias is of paramount importance for the clinical relevance of digital twins. Thus, the pipeline can be extended to include the recovery phase of the cardiac cycle (Neic et al., 2017), and validate it on the EDGAR database (Aras et al., 2015), which provides simultaneously acquired epicardial surface and body surface potential measurements during abnormal cardiac activity.

To extend the pipeline to treat patients at high arrhythmic risk due to abnormal tissue heterogeneities (Figure 8.1 – green).

The most prevalent pro-arrhythmic characteristic is having pathological conduction heterogeneities in regions of the cardiac tissue. Therefore, a deep learning-based segmentation algorithm could automatically extract scar and fibrosis regions from late gadolinium enhancement CMR data, available in the UK Biobank (Sudlow et al., 2015). Then, these segmented regions could be integrated into the 3D reconstruction of the heart. Next, the inference pipeline can be extended to estimate different conduction speeds in these pathological tissue regions. The cardiac digital twins generated from these extended methodologies would treat high-risk patients and aid in their preventative therapy selection.

To extend the pipeline to stratify patients at high arrhythmic risk due to abnormalities in the heart's electrical network (Figure 8.1 blue).

A key contribution of this proposal is the estimation of the patient's realistic Purkinje network. Neither tissue conduction heterogeneities nor Purkinje network abnormalities can be non-invasively measured in the clinic. However, while scars, visible from medical imaging data, are associated with tissue conduction heterogeneities (Dastidar et al., 2019). The structure of the Purkinje network cannot be observed in living people; thus, the inclusion of a realistic Purkinje network will enable the identification of arrhythmic risk in otherwise clinically misclassified patients. I propose that the properties defining a realistic Purkinje network are

inferred as in Barber et al., 2018 but from ECG recordings. This extension should be validated in virtual and real (Aras et al., 2015) patient cohorts.

8.3 Concluding remarks

This thesis presents the foundations of an efficient methodology for routinely generating cardiac digital twins from healthy people using standard non-invasive clinical recordings (e.g. CMR and ECG). Further development of the presented methods will enable their generalisation to patients with the most common life-threatening cardiac pathologies and patients suffering conditions that would otherwise be misclassified in the clinic. Consequently, continuing this work will cost-effectively improve diagnostic and therapeutical decision-making by generating digital twins to identify complex risk factors and simulate treatment outcomes.

References

Acharya, U. Rajendra, Hamido Fujita, Oh Shu Lih, Yuki Hagiwara, Jen Hong Tan, and Muhammad Adam (2017a). "Automated detection of arrhythmias using different intervals of tachycardia ECG segments with convolutional neural network". In: *Information Sciences* 405, pp. 81–90. URL: https://www.sciencedirect.com/science/article/pii/S0020025517306539.

Acharya, U. Rajendra, Hamido Fujita, Shu Lih Oh, U. Raghavendra, Jen Hong Tan, Muhammad Adam, Arkadiusz Gertych, and Yuki Hagiwara (2018). "Automated identification of shockable and non-shockable life-threatening ventricular arrhythmias using convolutional neural network". In: *Future Generation Computer Systems* 79, pp. 952 –959. URL: http://www.sciencedirect.com/science/article/pii/S0167739X17315248.

Acharya, U. Rajendra, Shu Lih Oh, Yuki Hagiwara, Jen Hong Tan, Muhammad Adam, Arkadiusz Gertych, and Ru San Tan (2017b). "A deep convolutional neural network model to classify heartbeats". In: *Computers in Biology and Medicine* 89, pp. 389–396. URL: https://www.sciencedirect.com/science/article/pii/S0010482517302810.

Al-Zaiti, Salah, Lucas Besomi, Zeineb Bouzid, Ziad Faramand, Stephanie Frisch, Christian Martin-Gill, Richard Gregg, Samir Saba, Clifton Callaway, and Ervin SejdiÄ‡ (Aug. 2020). "Machine learning-based prediction of acute coronary syndrome using only the pre-hospital 12-lead electrocardiogram". en. In: *Nature Communications* 11.1. Bandiera_abtest: a Cc_license_type: cc_by Cg_type: Nature Research Journals Number: 1 Primary_atype: Research Publisher: Nature Publishing Group Subject_term: Acute coronary syndromes;Diagnostic markers Subject_term_id: acute-coronary-syndromes;diagnostic-markers, p. 3966. URL: https://www.nature.com/articles/s41467-020-17804-2 (visited on 09/02/2021).

Albrecht, P (1983). "ST segment characterization for long term automated ECG analysis". In: *M.S. thesis, MIT Dept. of Electrical Engineering and Computer Science*. URL: https://scholar.google.es/scholar?hl=es{\&}as{_}sdt=0{\%}2C5{\&}q=S-T+segment+characterization+for+long-term+automated+ECG+analysis{\&}btnG=.

Aliev, Rubin R. and Alexander V. Panfilov (1996). "A Simple Two-variable Model of Cardiac Excitation". In: *Chaos, Solitons and Fractals* 7.3, pp. 293–301. URL: https://www.sciencedirect.com/science/article/pii/0960077995000895.

Almeida, R, JP Martínez, and AP Rocha (2006). "Automatic multilead VCG based approach for QT interval measurement". In: *2006 Computers in Cardiology*. URL: https://ieeexplore.ieee.org/abstract/document/4511865/.

Almeida, R., J.P. Martinez, A.P. Rocha, and P. Laguna (2009). "Multilead ECG Delineation Using Spatially Projected Leads From Wavelet Transform Loops". In:

IEEE Transactions on Biomedical Engineering 56.8, pp. 1996–2005. URL: http://ieeexplore.ieee.org/document/4915796/.

Almeida, R, J P Martinez, a P Rocha, and P Laguna (2010). "P wave delineation using spatially projected leads from wavelet transform loops". In: *Computers in Cardiology 2010* XX, pp. 1003–1006. URL: http://www.cinc.org/archives/2010/pdf/1003.pdf.

Almeida, R., E. Pueyo, J. P. Martinez, A. P. Rocha, and P. Laguna (Sept. 2004). "Quantification of the QT variability related to HRV: robustness study facing automatic delineation and noise on the ECG". In: *Computers in Cardiology, 2004*, pp. 769–772.

Andrews Christopher M., Srinivasan Neil T., Rosmini Stefania, Bulluck Heerajnarain, Orini Michele, Jenkins Sharon, Pantazis Antonis, McKenna William J., Moon James C., Lambiase Pier D., and Rudy Yoram (July 2017). "Electrical and Structural Substrate of Arrhythmogenic Right Ventricular Cardiomyopathy Determined Using Noninvasive Electrocardiographic Imaging and Late Gadolinium Magnetic Resonance Imaging". In: *Circulation: Arrhythmia and Electrophysiology* 10.7. Publisher: American Heart Association, e005105. URL: https://www.ahajournals.org/doi/10.1161/CIRCEP.116.005105 (visited on 10/24/2020).

Antzelevitch, Charles (Oct. 2007). "Role of spatial dispersion of repolarization in inherited and acquired sudden cardiac death syndromes". In: *American Journal of Physiology-Heart and Circulatory Physiology* 293.4. Publisher: American Physiological Society, H2024–H2038. URL: https://journals.physiology.org/doi/full/10.1152/ajpheart.00355.2007 (visited on 03/30/2021).

Aras, Kedar, Wilson Good, Jess Tate, Brett Burton, Dana Brooks, Jaume Coll-Font, Olaf Doessel, Walther Schulze, Danila Potyagaylo, Linwei Wang, Peter van Dam, and Rob MacLeod (2015). "Experimental Data and Geometric Analysis Repository - Edgar". In: *Journal of electrocardiology* 48.6, pp. 975–981. URL: https://www.ncbi.nlm.nih.gov/pmc/articles/PMC4624576/ (visited on 11/16/2020).

Arevalo, Hermenegild J., Fijoy Vadakkumpadan, Eliseo Guallar, Alexander Jebb, Peter Malamas, Katherine C. Wu, and Natalia A. Trayanova (May 2016). "Arrhythmia risk stratification of patients after myocardial infarction using personalized heart models". en. In: *Nature Communications* 7.1. Number: 1 Publisher: Nature Publishing Group, p. 11437. URL: https://www.nature.com/articles/ncomms11437 (visited on 07/06/2020).

Barber, Fernando, Miguel Lozano, Ignacio García-Fernández, and Rafael Sebastián (2018). "Inverse estimation of terminal connections in the cardiac conduction system". en. In: *Mathematical Methods in the Applied Sciences* 41.6. _eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1002/mma.4280, pp. 2340–2349. URL: http://undefined/doi/abs/10.1002/mma.4280 (visited on 11/16/2020).

Bayer, J. D., R. C. Blake, G. Plank, and N. A. Trayanova (Oct. 2012). "A Novel Rule-Based Algorithm for Assigning Myocardial Fiber Orientation to Computational Heart Models". en. In: *Annals of Biomedical Engineering* 40.10, pp. 2243–2254. URL: https://doi.org/10.1007/s10439-012-0593-5 (visited on 01/05/2021).

Bear, Laura R., Oumayma Bouhamama, Matthijs Cluitmans, Josselin Duchateau, Richard D. Walton, Emma Abell, et al. (2019). "Advantages and pitfalls of

noninvasive electrocardiographic imaging". In: *Journal of Electrocardiology* 57, S15–S20.

Bear, Laura R, Y Serinagaoglu Dogrusoz, J Svehlikova, J Coll-Font, W Good, E van Dam, R Macleod, E Abell, R Walton, R Coronel, Michel Haissaguerre, and R Dubois (Sept. 2018). "Effects of ECG Signal Processing on the Inverse Problem of Electrocardiography". In: *Computing in cardiology* 45. URL: https://www.ncbi.nlm.nih.gov/pmc/articles/PMC6424339/ (visited on 01/20/2021).

Bers, Donald M. (Jan. 2002). "Cardiac excitation–contraction coupling". en. In: *Nature* 415.6868. Number: 6868 Publisher: Nature Publishing Group, pp. 198–205. URL: https://www.nature.com/articles/415198a (visited on 04/07/2021).

Berte, Benjamin, Mattias Duytschaever, Juliana Elices, Vikas Kataria, Liesbeth Timmers, Frédéric Van Heuverswyn, Roland Stroobandt, Jan De Neve, Karel Watteyne, Elke Vandensteen, Yves Vandekerckhove, and Rene Tavernier (Sept. 2015). "Variability in interpretation of the electrocardiogram in young athletes: an unrecognized obstacle for electrocardiogram-based screening protocols". en. In: *Europace* 17.9, pp. 1435–1440. URL: https://academic.oup.com/europace/article-lookup/doi/10.1093/europace/euu385 (visited on 06/09/2020).

Biesbroek, P Stefan, Alexander Hirsch, Alwin Zweerink, Peter M van de Ven, Aernout M Beek, Maarten Groenink, Fons Windhausen, R Nils Planken, Albert C van Rossum, and Robin Nijveldt (Dec. 2018). "Additional diagnostic value of CMR to the European Society of Cardiology (ESC) position statement criteria in a large clinical population of patients with suspected myocarditis". In: *European Heart Journal - Cardiovascular Imaging* 19.12, pp. 1397–1407. URL: https://doi.org/10.1093/ehjci/jex308 (visited on 02/01/2021).

Bleijendaal, H, RR Van Der Leur, K Taha, T Mast, JMIH Gho, MM Winter, AH Zwinderman, PA Doevendans, YM Pinto, FW Asselbergs, R Van Es, and FVY Tjong (May 2021a). "Electrocardiogram-based mortality prediction in patients with COVID-19 using machine learning". In: *EP Europace* 23.Supplement_3. URL: https://doi.org/10.1093/europace/euab116.512 (visited on 09/02/2021).

Bleijendaal, Hidde, Lucas A. Ramos, Ricardo R. Lopes, Tom E. Verstraelen, Sarah W. E. Baalman, Marinka D. Oudkerk Pool, Fleur V. Y. Tjong, Francisco M. Melgarejo-Meseguer, F. Javier Gimeno-Blanes, Juan R. Gimeno-Blanes, Ahmad S. Amin, Michiel M. Winter, Henk A. Marquering, Wouter E. M. Kok, Aeilko H. Zwinderman, Arthur A. M. Wilde, and Yigal M. Pinto (Jan. 2021b). "Computer versus cardiologist: Is a machine learning algorithm able to outperform an expert in diagnosing a phospholamban p.Arg14del mutation on the electrocardiogram?" en. In: *Heart Rhythm* 18.1, pp. 79–87. URL: https://www.sciencedirect.com/science/article/pii/S1547527120308614 (visited on 09/02/2021).

Bliek, Erwin Christian de (2018). "ST elevation: Differential diagnosis and caveats. A comprehensive review to help distinguish ST elevation myocardial infarction from nonischemic etiologies of ST elevation". In: *Turkish Journal of Emergency Medicine* 18.1, pp. 1–10. URL: https://www.sciencedirect.com/science/article/pii/S2452247317302571.

Bote, Jose Manuel, Joaquin Recas, Francisco Rincon, David Atienza, and Roman Hermida (2018). "A Modular Low-Complexity ECG Delineation Algorithm

for Real-Time Embedded Systems". In: *IEEE Journal of Biomedical and Health Informatics* 22.2, pp. 429–441.

Bottou, Léon (1998). "Online learning and stochastic approximations". In: *On-line learning in neural networks* 17.9, p. 142.

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. URL: https://degruyter.com/view/j/bmte.1995.40.issue-s1/bmte.1995.40.s1.317/bmte.1995.40.s1.317.xml.

BOXYM (2018). *What is Blood Circulation? Understanding Blood Circulation System - BOXYM.COM*. en-US. Section: oximeter. URL: https://boxym.com/what-is-blood-circulation/ (visited on 12/08/2020).

Boyle, Patrick M., Tarek Zghaib, Sohail Zahid, Rheeda L. Ali, Dongdong Deng, William H. Franceschi, Joe B. Hakim, Michael J. Murphy, Adityo Prakosa, Stefan L. Zimmerman, Hiroshi Ashikaga, Joseph E. Marine, Aravindan Kolandaivelu, Saman Nazarian, David D. Spragg, Hugh Calkins, and Natalia A. Trayanova (Nov. 2019). "Computationally guided personalized targeted ablation of persistent atrial fibrillation". en. In: *Nature Biomedical Engineering* 3.11, pp. 870–879. URL: http://www.nature.com/articles/s41551-019-0437-9 (visited on 08/30/2020).

Britton, Oliver J., Alfonso Bueno-Orovio, Karel Van Ammel, Hua Rong Lu, Rob Towart, David J. Gallacher, and Blanca Rodriguez (June 2013). "Experimentally calibrated population of models predicts and explains intersubject variability in cardiac cellular electrophysiology". en. In: *Proceedings of the National Academy of Sciences* 110.23. Publisher: National Academy of Sciences Section: PNAS Plus, E2098–E2105. URL: https://www.pnas.org/content/110/23/E2098 (visited on 09/09/2020).

Bukhari, H. A., F. Palmieri, D. Ferreira, M. Potse, J. Ramírez, P. Laguna, C. Sánchez, and E. Pueyo (Sept. 2019). "Transmural Ventricular Heterogeneities Play a Major Role in Determining T-Wave Morphology at Different Extracellular Potassium Levels". In: *2019 Computing in Cardiology (CinC)*. ISSN: 2325-887X, pp. 1–4.

Bycroft, Clare, Colin Freeman, Desislava Petkova, Gavin Band, Lloyd T Elliott, Kevin Sharp, Allan Motyer, Damjan Vukcevic, Olivier Delaneau, Jared O'Connell, et al. (2018). "The UK Biobank resource with deep phenotyping and genomic data". In: *Nature* 562.7726, pp. 203–209.

Calaprice, Alice (Oct. 2010). *The Ultimate Quotable Einstein*. en. Princeton University Press. (Visited on 11/29/2020).

Caldwell, Bryan, Mark Trew, Gregory Sands, Darren Hooks, Ian LeGrice, and Bruce Smaill (Aug. 2009). "Three Distinct Directions of Intramural Activation Reveal Nonuniform Side-to-Side Electrical Coupling of Ventricular Myocytes". In: *Circulation: Arrhythmia and Electrophysiology* 2.4. Publisher: American Heart Association, pp. 433–440. URL: https://www.ahajournals.org/doi/10.1161/CIRCEP.108.830133 (visited on 01/04/2021).

Camps, Julià, Blanca Rodríguez, and Ana Mincholé (Sept. 2018). "Deep Learning Based QRS Multilead Delineator in Electrocardiogram Signals". In: *2018 Computing in Cardiology Conference (CinC)*. Vol. 45. ISSN: 2325-887X, pp. 1–4.

Cardone-Noott, Louie, Alfonso Bueno-Orovio, Ana Mincholé, Nejib Zemzemi, and Blanca Rodriguez (Dec. 2016). "Human ventricular activation sequence and the simulation of the electrocardiographic QRS complex and its variability in healthy and

intraventricular block conditions". en. In: *EP Europace* 18.suppl_4, pp. iv4–iv15. URL: https://academic.oup.com/europace/article/18/suppl_4/iv4/2737164 (visited on 06/09/2020).

Cerqueira Manuel D., Weissman Neil J., Dilsizian Vasken, Jacobs Alice K., Kaul Sanjiv, Laskey Warren K., Pennell Dudley J., Rumberger John A., Ryan Thomas, and Verani Mario S. (Jan. 2002). "Standardized Myocardial Segmentation and Nomenclature for Tomographic Imaging of the Heart". In: *Circulation* 105.4. Publisher: American Heart Association, pp. 539–542. URL: https://www.ahajournals.org/doi/10.1161/hc0402.102975 (visited on 10/28/2020).

Cheng, Maowei, Worku J. Sori, Feng Jiang, Adil Khan, and Shaohui Liu (2017). "Recurrent Neural Network Based Classification of ECG Signal Features for Obstruction of Sleep Apnea Detection". In: *Proceedings - 2017 IEEE International Conference on Computational Science and Engineering and IEEE/IFIP International Conference on Embedded and Ubiquitous Computing, CSE and EUC 2017*. Vol. 2. IEEE, pp. 199–202. URL: http://ieeexplore.ieee.org/document/8005997/.

Chinchapatnam, Phani, Kawal S. Rhode, Matthew Ginks, C. Aldo Rinaldi, Pier Lambiase, Reza Razavi, Simon Arridge, and Maxime Sermesant (2008). "Model-based imaging of cardiac apparent conductivity and local conduction velocity for diagnosis and planning of therapy". In: *IEEE Transactions on Medical Imaging* 27.11, pp. 1631–1642. URL: http://www.ncbi.nlm.nih.gov/pubmed/18955178http://ieeexplore.ieee.org/document/4608725/.

Cluitmans, Matthijs, Dana H. Brooks, Rob MacLeod, Olaf Dössel, María S. Guillem, Peter M. van Dam, Jana Svehlikova, Bin He, John Sapp, Linwei Wang, and Laura Bear (Sept. 2018). "Validation and Opportunities of Electrocardiographic Imaging: From Technical Achievements to Clinical Applications". In: *Frontiers in Physiology* 9. URL: https://www.ncbi.nlm.nih.gov/pmc/articles/PMC6158556/ (visited on 10/24/2020).

Cluitmans, M.J.M., R.L.M. Peeters, R.L. Westra, and P.G.A. Volders (June 2015). "Noninvasive reconstruction of cardiac electrical activity: update on current methods, applications and challenges". en. In: *Netherlands Heart Journal* 23.6, pp. 301–311. URL: https://doi.org/10.1007/s12471-015-0690-9 (visited on 08/30/2020).

Coll-Font, Jaume, Matthijs Cluitmans, Paul Volders, Rob S. MacLeod, and Dana H. Brooks (Nov. 2019). "Spatiotemporal Estimation of Activation Times in Cardiac Electrograms". en. In: *Journal of Electrocardiology* 57, S111. URL: http://www.sciencedirect.com/science/article/pii/S0022073619308362 (visited on 07/14/2020).

Colli Franzone, P., L. Guerri, and S. Rovida (Feb. 1990). "Wavefront propagation in an activation model of the anisotropic cardiac tissue: asymptotic analysis and numerical simulations". en. In: *Journal of Mathematical Biology* 28.2, pp. 121–176. URL: http://link.springer.com/10.1007/BF00163143 (visited on 06/09/2020).

Colli-Franzone, P., L. Guerri, S. Tentoni, C. Viganotti, S. Baruffi, S. Spaggiari, and B. Taccardi (Dec. 1985). "A mathematical procedure for solving the inverse potential problem of electrocardiography. analysis of the time-space accuracy from in vitro experimental data". en. In: *Mathematical Biosciences* 77.1, pp. 353–396. URL: http://www.sciencedirect.com/science/article/pii/0025556485901063 (visited on 01/11/2021).

Commons, Wikimedia (2020). *Wikimedia Commons, the free media repository*. [Online; accessed 7-December-2020]. URL: https://commons.wikimedia.org.

Computing in Cardiology. URL: http://www.cinc.org/ (visited on 09/21/2018).

Corral-Acero, Jorge et al. (Mar. 2020). "The 'Digital Twin' to enable the vision of precision cardiology". en. In: *European Heart Journal*, ehaa159. URL: https://academic.oup.com/eurheartj/advance-article/doi/10.1093/eurheartj/ehaa159/5775673 (visited on 11/17/2020).

Creative Commons Attribution 3.0 Unported CC BY 3.0. URL: https://creativecommons.org/licenses/by/3.0/deed.en (visited on 07/08/2021).

Creative Commons Attribution 4.0 International CC BY 4.0. URL: https://creativecommons.org/licenses/by/4.0/ (visited on 07/08/2021).

Creative Commons Attribution-ShareAlike 4.0 International CC BY-SA 4.0. URL: https://creativecommons.org/licenses/by-sa/4.0/ (visited on 07/08/2021).

Creative Commons CC0 1.0 Universal. URL: https://creativecommons.org/publicdomain/zero/1.0/deed.en (visited on 07/08/2021).

Cuculich, Phillip S., Matthew R. Schill, Rojano Kashani, Sasa Mutic, Adam Lang, Daniel Cooper, Mitchell Faddis, Marye Gleva, Amit Noheria, Timothy W. Smith, Dennis Hallahan, Yoram Rudy, and Clifford G. Robinson (Dec. 2017). "Noninvasive Cardiac Radiation for Ablation of Ventricular Tachycardia". In: *New England Journal of Medicine* 377.24. Publisher: Massachusetts Medical Society _eprint: https://doi.org/10.1056/NEJMoa1613773, pp. 2325–2336. URL: https://doi.org/10.1056/NEJMoa1613773 (visited on 12/08/2020).

Cuiwei Li, Chongxun Zheng, and Changfeng Tai (1995). "Detection of ECG characteristic points using wavelet transforms". In: *IEEE Transactions on Biomedical Engineering* 42.1, pp. 21–28. URL: http://ieeexplore.ieee.org/document/362922/.

Curie, Eve (1937). *Madame Curie. A biography*. Doubleday, Doran and Company.

Dall'Armellina, Erica, Theodoros D. Karamitsos, Stefan Neubauer, and Robin P. Choudhury (Nov. 2010). "CMR for characterization of the myocardium in acute coronary syndromes". en. In: *Nature Reviews Cardiology* 7.11. Number: 11 Publisher: Nature Publishing Group, pp. 624–636. URL: https://www.nature.com/articles/nrcardio.2010.140/email/correspondent/c1/ (visited on 01/04/2021).

Das, Mithilesh K., Hussam Suradi, Waddah Maskoun, Mark A. Michael, Changyu Shen, Jonathan Peng, Gopi Dandamudi, and Jo Mahenthiran (Oct. 2008). "Fragmented Wide QRS on a 12-Lead ECG". In: *Circulation: Arrhythmia and Electrophysiology* 1.4. Publisher: American Heart Association, pp. 258–268. URL: https://www.ahajournals.org/doi/full/10.1161/CIRCEP.107.763284 (visited on 07/05/2021).

Dastidar, Amardeep Ghosh, Anna Baritussio, Estefania De Garate, Zsofia Drobni, Giovanni Biglino, Priyanka Singhal, Elena G. Milano, Gianni D. Angelini, Stephen Dorman, Julian Strange, Thomas Johnson, and Chiara Bucciarelli-Ducci (Oct. 2019). "Prognostic Role of CMR and Conventional Risk Factors in Myocardial Infarction With Nonobstructed Coronary Arteries". en. In: *JACC: Cardiovascular Imaging* 12.10, pp. 1973–1982. URL: https://www.sciencedirect.com/science/article/pii/S1936878X19300658 (visited on 03/19/2021).

Dastidar, Amardeep Ghosh, Jonathan C. L. Rodrigues, Nauman Ahmed, Anna Baritussio, and Chiara Bucciarelli-Ducci (2015). "The Role of Cardiac MRI in Patients with Troponin-Positive Chest Pain and Unobstructed Coronary Arteries". In: *Current Cardiovascular Imaging Reports* 8.8. URL: https://www.ncbi.nlm.nih.gov/pmc/articles/PMC4483181/ (visited on 03/19/2021).

Dhamala, Jwala, Hermenegild J. Arevalo, John Sapp, Milan Horacek, Katherine C. Wu, Natalia A. Trayanova, and Linwei Wang (2017). "Spatially adaptive multi-scale optimization for local parameter estimation in cardiac electrophysiology". In: *IEEE Transactions on Medical Imaging* 36.9, pp. 1966–1978. URL: http://ieeexplore.ieee.org/document/7911242/.

Di Marco, Luigi Y and Lorenzo Chiari (2011). "A wavelet-based ECG delineation algorithm for 32-bit integer online processing". In: *BioMedical Engineering OnLine* 10.1, p. 23. URL: http://biomedical-engineering-online.biomedcentral.com/articles/10.1186/1475-925X-10-23.

Dijkstra, E. W. (Dec. 1959). "A note on two problems in connexion with graphs". en. In: *Numerische Mathematik* 1.1, pp. 269–271. URL: https://doi.org/10.1007/BF01386390 (visited on 09/08/2020).

Doste, Ruben, David Soto-Iglesias, Gabriel Bernardino, Alejandro Alcaine, Rafael Sebastian, Sophie Giffard-Roisin, Maxime Sermesant, Antonio Berruezo, Damian Sanchez-Quintana, and Oscar Camara (2019). "A rule-based method to model myocardial fiber orientation in cardiac biventricular geometries with outflow tracts". en. In: *International Journal for Numerical Methods in Biomedical Engineering* 35.4. _eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1002/cnm.3185, e3185. URL: https://onlinelibrary.wiley.com/doi/abs/10.1002/cnm.3185 (visited on 01/05/2021).

Drovandi, C. C. and A. N. Pettitt (2011). "Estimation of Parameters for Macroparasite Population Evolution Using Approximate Bayesian Computation". en. In: *Biometrics* 67.1. _eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1111/j.1541-0420.2010.01410.x, pp. 225–233. URL: https://onlinelibrary.wiley.com/doi/abs/10.1111/j.1541-0420.2010.01410.x (visited on 09/14/2020).

Duchi, John, Elad Hazan, and Yoram Singer (2011). "Adaptive subgradient methods for online learning and stochastic optimization". In: *Journal of Machine Learning Research* 12.Jul, pp. 2121–2159.

Durrer, Dirk, Van Dam R. Th., Freud G. E., Janse M. J., Meijler F. L., and Arzbaecher R. C. (June 1970). "Total Excitation of the Isolated Human Heart". In: *Circulation* 41.6. Publisher: American Heart Association, pp. 899–912. URL: https://www.ahajournals.org/doi/10.1161/01.CIR.41.6.899 (visited on 07/13/2020).

Dutta, Sara, Ana Mincholé, T. Alexander Quinn, and Blanca Rodriguez (Oct. 2017). "Electrophysiological properties of computational human ventricular cell action potential models under acute ischemic conditions". en. In: *Progress in Biophysics and Molecular Biology*. Validation of Computer Modelling 129, pp. 40–52. URL: http://www.sciencedirect.com/science/article/pii/S0079610716300918 (visited on 01/05/2021).

Ebrahimi, Zahra, Mohammad Loni, Masoud Daneshtalab, and Arash Gharehbaghi (Sept. 2020). "A review on deep learning methods for ECG arrhythmia classification". en.

In: *Expert Systems with Applications: X* 7, p. 100033. URL: https://www.sciencedirect.com/science/article/pii/S2590188520300123 (visited on 09/02/2021).

Englewood Hospital (2020). *Arrhythmia Treatment |Englewood Health Physician Network |NJ & NY*. URL: https://www.englewoodhealthphysicians.org/medical-services/heart-care/cardiovascular-conditions-we-treat/arrhythmias/ (visited on 12/08/2020).

Faust, Oliver, Yuki Hagiwara, Tan Jen Hong, Oh Shu Lih, and U Rajendra Acharya (2018). *Deep learning for healthcare applications based on physiological signals: A review*. URL: https://linkinghub.elsevier.com/retrieve/pii/S0169260718301226.

Ferreira, Vanessa M. (Oct. 2019). "CMR Should Be a Mandatory Test in the Contemporary Evaluation of "MINOCA"". en. In: *JACC: Cardiovascular Imaging* 12.10, pp. 1983–1986. URL: https://www.sciencedirect.com/science/article/pii/S1936878X19304644 (visited on 03/19/2021).

Finn, Chelsea, Ian Goodfellow, and Sergey Levine (2016). "Unsupervised Learning for Physical Interaction through Video Prediction". In: *Advances in Neural Information Processing Systems 29*. Ed. by D. D. Lee, M. Sugiyama, U. V. Luxburg, I. Guyon, and R. Garnett. Curran Associates, Inc., pp. 64–72. URL: http://papers.nips.cc/paper/6161-unsupervised-learning-for-physical-interaction-through-video-prediction.pdf.

Fourier, Joseph (June 1878). "The Analytical Theory of Heat". en. In: *Nature* 18.451. Number: 451 Publisher: Nature Publishing Group, pp. 192–192. (Visited on 11/30/2020).

FRANK, E. (1956). "An accurate, clinically practical system for spatial vectorcardiography." In: *Circulation* 13.5, pp. 737–749. URL: https://www.ahajournals.org/doi/10.1161/01.CIR.13.5.737.

Frazier, D W, W Krassowska, P S Chen, P D Wolf, N D Danieley, W M Smith, and R E Ideker (July 1988). "Transmural activations and stimulus potentials in three-dimensional anisotropic canine myocardium." en. In: *Circulation Research* 63.1, pp. 135–146. URL: https://www.ahajournals.org/doi/10.1161/01.RES.63.1.135 (visited on 07/05/2021).

Ghaffari, A., M. R. Homaeinezhad, M. Akraminia, M. Atarod, and M. Daevaeiha (2009). "A robust wavelet-based multi-lead electrocardiogram delineation algorithm". In: *Medical Engineering and Physics* 31.10, pp. 1219–1227. URL: https://www.sciencedirect.com/science/article/pii/S1350453309001647?via{\%}3Dihub.

Ghanem, Raja N., Ping Jia, Charulatha Ramanathan, Kyungmoo Ryu, Alan Markowitz, and Yoram Rudy (Apr. 2005). "Noninvasive Electrocardiographic Imaging (ECGI): Comparison to intraoperative mapping in patients". In: *Heart rhythm : the official journal of the Heart Rhythm Society* 2.4, pp. 339–354. URL: https://www.ncbi.nlm.nih.gov/pmc/articles/PMC1949041/ (visited on 08/21/2020).

Giffard-Roisin, Sophie, Hervé Delingette, Thomas Jackson, Lauren Fovargue, Jack Lee, Aldo Rinaldi, Nicholas Ayache, Reza Razavi, and Maxime Sermesant (2017a). "Sparse Bayesian Non-linear Regression for Multiple Onsets Estimation in Non-invasive Cardiac Electrophysiology". en. In: *Functional Imaging and Modelling of the Heart*.

Ed. by Mihaela Pop and Graham A Wright. Lecture Notes in Computer Science. Cham: Springer International Publishing, pp. 230–238.

Giffard-Roisin, Sophie, Thomas Jackson, Lauren Fovargue, Jack Lee, Hervé Delingette, Reza Razavi, Nicholas Ayache, and Maxime Sermesant (Sept. 2017b). "Noninvasive Personalization of a Cardiac Electrophysiology Model From Body Surface Potential Mapping". In: *IEEE Transactions on Biomedical Engineering* 64.9. Conference Name: IEEE Transactions on Biomedical Engineering, pp. 2206–2218.

Gilks, W. R. (2005). "Markov Chain Monte Carlo". en. In: *Encyclopedia of Biostatistics*. _eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1002/0470011815.b2a14021. American Cancer Society. URL: https://onlinelibrary.wiley.com/doi/abs/10.1002/0470011815.b2a14021 (visited on 09/10/2020).

Gima, Kazutaka and Yoram Rudy (May 2002). "Ionic Current Basis of Electrocardiographic Waveforms: A Model Study". en. In: *Circulation Research* 90.8, pp. 889–896. URL: https://www.ahajournals.org/doi/10.1161/01.RES.0000016960.61087.86 (visited on 07/06/2020).

Goethe, Johann Wolfgang von (1829). *Wilhelm Meisters Wanderjahre oder die Entsagenden*. de. Cotta.

Goldberger, Ary L, Luis A N Amaral, ; Leon Glass, Jeffrey M Hausdorff, ; Plamen, Ch Ivanov, Roger G Mark, Joseph E Mietus, George B Moody, Chung-Kang Peng, and ; H Eugene Stanley (2000). *PhysioBank, PhysioToolkit, and PhysioNet Components of a New Research Resource for Complex Physiologic Signals*. Tech. rep. URL: http://www.physionet.org.

Goodfellow, Ian, Yoshua Bengio, and Aaron Courville (2016). *Deep Learning*. http://www.deeplearningbook.org. MIT Press.

Grandits, Thomas, Karli Gillette, Aurel Neic, Jason Bayer, Edward Vigmond, Thomas Pock, and Gernot Plank (July 2020). "An Inverse Eikonal Method for Identifying Ventricular Activation Sequences from Epicardial Activation Maps". en. In: *Journal of Computational Physics*, p. 109700. URL: http://www.sciencedirect.com/science/article/pii/S0021999120304745 (visited on 07/07/2020).

Grant Augustus O. (Apr. 2009). "Cardiac Ion Channels". In: *Circulation: Arrhythmia and Electrophysiology* 2.2. Publisher: American Heart Association, pp. 185–194. URL: https://www.ahajournals.org/doi/full/10.1161/circep.108.789081 (visited on 04/07/2021).

Graves, Alex, Greg Wayne, Malcolm Reynolds, Tim Harley, Ivo Danihelka, Agnieszka Grabska-Barwińska, Sergio Gómez Colmenarejo, Edward Grefenstette, Tiago Ramalho, John Agapiou, et al. (2016). "Hybrid computing using a neural network with dynamic external memory". In: *Nature* 538.7626, pp. 471–476.

Gray, Henry and Warren Harmon Lewis (1918). *Anatomy of the Human Body*. en. Philadelphia: Lea & Febiger. URL: https://www.bartleby.com/107/ (visited on 12/07/2020).

Greenwald, SD and RS Patil (1990). "Improved detection and classification of arrhythmias in noise-corrupted electrocardiograms using contextual information". In: *1990 Computers in Cardiology*. URL: https://ieeexplore.ieee.org/abstract/document/144257/.

Haissaguerre, Michel, Edward Vigmond, Bruno Stuyvers, Meleze Hocini, and Olivier Bernus (Mar. 2016). "Ventricular arrhythmias and the His–Purkinje system". en. In: *Nature Reviews Cardiology* 13.3. Number: 3 Publisher: Nature Publishing Group, pp. 155–166. URL: https://www.nature.com/articles/nrcardio.2015.193 (visited on 06/16/2020).

Harvey, William (1628). *Exercitatio Anatomica de Motu Cordis et Sanguinis in Animalibus*. Latin. 1st. Frankfort: Sumptibus Guilielmi Fitzeri.

Hausvater, Anais, Sivabaskari Pasupathy, Per Tornvall, Himali Gandhi, Rosanna Tavella, John Beltrame, Stefan Agewall, Christina Ekenbäck, Elin Bacsovics Brolin, Judith S. Hochman, Olov Collste, and Harmony R. Reynolds (July 2019). "ST-segment elevation and cardiac magnetic resonance imaging findings in myocardial infarction with non-obstructive coronary arteries". en. In: *International Journal of Cardiology* 287, pp. 128–131. URL: https://www.sciencedirect.com/science/article/pii/S0167527318374266 (visited on 03/19/2021).

Hayn, D, A Kollmann, and G Schreier (2006). "Automated QT interval measurement from multilead ECG signals". In: *2006 Computers in Cardiology*, pp. 381–384. URL: https://ieeexplore.ieee.org/abstract/document/4511868/.

He, Kaiming, Xiangyu Zhang, Shaoqing Ren, and Jian Sun (2015). "Deep residual learning for image recognition". In: *arXiv preprint arXiv:1512.03385*.

Homaeinezhad, M.R., M. ErfanianMoshiri-Nejad, and H. Naseri (2014). "A correlation analysis-based detection and delineation of ECG characteristic events using template waveforms extracted by ensemble averaging of clustered heart cycles". In: *Computers in Biology and Medicine* 44, pp. 66–75. URL: https://www.sciencedirect.com/science/article/pii/S0010482513003132.

Hughes, Nicholas P., Lionel Tarassenko, and Stephen J. Roberts (2004). "Markov models for automated ECG interval analysis". In: *Advances in Neural Information Processing Systems* 16, pp. 611–618. URL: http://machinelearning.wustl.edu/mlpapers/paper{_}files/NIPS2003{_}AP06.pdf.

Ideker, Raymond E., Wei Kong, and Steven Pogwizd (Mar. 2009). "Purkinje Fibers and Arrhythmias". In: *Pacing and clinical electrophysiology : PACE* 32.3, pp. 283–285. URL: https://www.ncbi.nlm.nih.gov/pmc/articles/PMC2743071/ (visited on 07/07/2020).

Iman, Ronald L. and W. J. Conover (Jan. 1980). "Small sample sensitivity analysis techniques for computer models.with an application to risk assessment". In: *Communications in Statistics - Theory and Methods* 9.17. Publisher: Taylor & Francis _eprint: https://doi.org/10.1080/03610928008827996, pp. 1749–1842. URL: https://doi.org/10.1080/03610928008827996 (visited on 09/08/2020).

Itakura, F. (Feb. 1975). "Minimum prediction residual principle applied to speech recognition". In: *IEEE Transactions on Acoustics, Speech, and Signal Processing* 23.1. Conference Name: IEEE Transactions on Acoustics, Speech, and Signal Processing, pp. 67–72.

Jimenez-Perez, Guillermo, Alejandro Alcaine, and Oscar Camara (Sept. 2019). "U-Net Architecture for the Automatic Detection and Delineation of the Electrocardiogram". In: *2019 Computing in Cardiology (CinC)*. ISSN: 2325-887X, Page 1–Page 4.

— (Jan. 2021). "Delineation of the electrocardiogram with a mixed-quality-annotations dataset using convolutional neural networks". en. In: *Scientific Reports* 11.1. Bandiera_abtest: a Cc_license_type: cc_by Cg_type: Nature Research Journals

Number: 1 Primary_atype: Research Publisher: Nature Publishing Group Subject_term: Biomedical engineering;Computer science Subject_term_id: biomedical-engineering;computer-science, p. 863. URL: https://www.nature.com/articles/s41598-020-79512-7 (visited on 07/06/2021).

Kania, MichaÅ,, HervÃ© Rix, MaÅ‚gorzata Fereniec, Heriberto Zavala-Fernandez, Dariusz Janusek, Tomasz Mroczka, GÃ¼nter Stix, and Roman Maniewski (Feb. 2014). "The effect of precordial lead displacement on ECG morphology". en. In: *Medical & Biological Engineering & Computing* 52.2, pp. 109–119. URL: https://doi.org/10.1007/s11517-013-1115-9 (visited on 07/05/2021).

Kaptoge, Stephen et al. (Oct. 2019). "World Health Organization cardiovascular disease risk charts: revised models to estimate risk in 21 global regions". English. In: *The Lancet Global Health* 7.10. Publisher: Elsevier, e1332–e1345. URL: https://www.thelancet.com/journals/langlo/article/PIIS2214-109X(19)30318-3/abstract (visited on 01/05/2021).

Kara, Vinay, Haibo Ni, Erick Andres Perez Alday, and Henggui Zhang (2019). "ECG Imaging to Detect the Site of Ventricular Ischemia Using Torso Electrodes: A Computational Study". English. In: *Frontiers in Physiology* 10. Publisher: Frontiers. URL: https://www.frontiersin.org/articles/10.3389/fphys.2019.00050/full (visited on 06/16/2020).

Kingma, Diederik P. and Jimmy Ba (2014b). "Adam: A Method for Stochastic Optimization". In: *CoRR* 1412.6980.

— (2014a). "Adam: A Method for Stochastic Optimization". In: *CoRR* abs/1412.6980. URL: http://arxiv.org/abs/1412.6980.

Kiranyaz, Serkan, Turker Ince, and Moncef Gabbouj (2016). "Real-Time Patient-Specific ECG Classification by 1-D Convolutional Neural Networks". In: *IEEE Transactions on Biomedical Engineering* 63.3, pp. 664–675. URL: http://ieeexplore.ieee.org/document/7202837/.

Konukoglu, Ender, Jatin Relan, Ulas Cilingir, Bjoern H. Menze, Phani Chinchapatnam, Amir Jadidi, Hubert Cochet, Mélèze Hocini, Hervé Delingette, Pierre Jaïs, Michel Haïssaguerre, Nicholas Ayache, and Maxime Sermesant (2011). "Efficient probabilistic model personalization integrating uncertainty on data and parameters: Application to Eikonal-Diffusion models in cardiac electrophysiology". In: *Progress in Biophysics and Molecular Biology* 107.1, pp. 134–146. URL: https://www.sciencedirect.com/science/article/pii/S0079610711000654.

Konukoglu, Ender, Maxime Sermesant, Olivier Clatz, Jean Marc Peyrat, Hervé Delingette, and Nicholas Ayache (2007). "A recursive anisotropic fast marching approach to reaction diffusion equation: Application to tumor growth modeling". In: *Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)*. Vol. 4584 LNCS. Berlin, Heidelberg: Springer Berlin Heidelberg, pp. 687–699. URL: http://link.springer.com/10.1007/978-3-540-73273-0{_}57.

Krizhevsky, Alex, Ilya Sutskever, and Geoffrey E Hinton (2012). "Imagenet classification with deep convolutional neural networks". In: *Advances in neural information processing systems*, pp. 1097–1105.

Kunisch, Karl, Aurel Neic, Gernot Plank, and Philip Trautmann (Dec. 2019). "Inverse localization of earliest cardiac activation sites from activation maps based on the

viscous Eikonal equation". en. In: *Journal of Mathematical Biology* 79.6, pp. 2033–2068. URL: https://doi.org/10.1007/s00285-019-01419-3 (visited on 07/07/2020).

Laguna, P., R.G. Mark, A. Goldberg, and G.B. Moody (1997). "A database for evaluation of algorithms for measurement of QT and other waveform intervals in the ECG". In: *Computers in Cardiology 1997* 24, pp. 673–676. URL: http://ieeexplore.ieee.org/document/648140/.

Laguna, P, N V Thakor, P Caminal, R Jané, H R Yoon, A Bayés de Luna, V Marti, and J Guindo (1990). "New algorithm for QT interval analysis in 24-hour Holter ECG: performance and applications". In: *Medical & biological engineering & computing* 28.1, pp. 67–73. URL: http://www.ncbi.nlm.nih.gov/pubmed/2325452.

Laguna, Pablo, Raimon Jané, and Pere Caminal (1994). "Automatic Detection of Wave Boundaries in Multilead ECG Signals: Validation with the CSE Database". In: *Computers and Biomedical Research* 27.1, pp. 45–60. URL: http://www.sciencedirect.com/science/article/pii/S0010480984710068http://linkinghub.elsevier.com/retrieve/pii/S0010480984710068.

LeCun, Yann, Yoshua Bengio, and Geoffrey Hinton (2015). "Deep learning". In: *Nature* 521.7553, pp. 436–444.

LeCun, Yann, Léon Bottou, Yoshua Bengio, and Patrick Haffner (1998). "Gradient-based learning applied to document recognition". In: *Proceedings of the IEEE* 86.11, pp. 2278–2324.

Leur, Rutger R. van de, Karim Taha, Max N. Bos, Jeroen F. van der Heijden, Deepak Gupta, Maarten J. Cramer, Rutger J. Hassink, Pim van der Harst, Pieter A. Doevendans, Folkert W. Asselbergs, and RenĂ© van Es (Feb. 2021). "Discovering and Visualizing Disease-Specific Electrocardiogram Features Using Deep Learning". In: *Circulation: Arrhythmia and Electrophysiology* 14.2. Publisher: American Heart Association, e009056. URL: https://www.ahajournals.org/doi/full/10.1161/CIRCEP.120.009056 (visited on 09/02/2021).

Lew, Seok, Carsten H. Wolters, Alfred Anwander, Scott Makeig, and Rob S. MacLeod (2009). "Improved EEG source analysis using low-resolution conductivity estimation in a four-compartment finite element head model". en. In: *Human Brain Mapping* 30.9. _eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1002/hbm.20714, pp. 2862–2878. URL: https://onlinelibrary.wiley.com/doi/abs/10.1002/hbm.20714 (visited on 09/10/2020).

Lines, G.T., M.L. Buist, P. Grottum, A.J. Pullan, J. Sundnes, and A. Tveito (Apr. 2003). "Mathematical models and numerical methods for the forward problem in cardiac electrophysiology". en. In: *Computing and Visualization in Science* 5.4, pp. 215–239. URL: https://doi.org/10.1007/s00791-003-0101-4 (visited on 04/09/2021).

Luo, Kan, Jianqing Li, Zhigang Wang, and Alfred Cuschieri (2017). "Patient-Specific Deep Architectural Model for ECG Classification". In: *Journal of Healthcare Engineering* 2017. URL: https://www.hindawi.com/journals/jhe/2017/4108720/abs/.

Luzianin, Ivan and Bernd Krause (2016). "Similarity Measurement of Biological Signals Using Dynamic Time Warping Algorithm". en. In: p. 7.

Lyon, Aurore, Rina Ariga, Ana Mincholé, Masliza Mahmod, Elizabeth Ormondroyd, Pablo Laguna, Nando de Freitas, Stefan Neubauer, Hugh Watkins, and

Blanca Rodriguez (2018a). "Distinct ECG Phenotypes Identified in Hypertrophic Cardiomyopathy Using Machine Learning Associate With Arrhythmic Risk Markers". English. In: *Frontiers in Physiology* 9. Publisher: Frontiers. URL: https://www.frontiersin.org/articles/10.3389/fphys.2018.00213/full?report=reader (visited on 06/09/2020).

Lyon, Aurore, Alfonso Bueno-Orovio, Ernesto Zacur, Rina Ariga, Vicente Grau, Stefan Neubauer, Hugh Watkins, Blanca Rodriguez, and Ana Mincholé (Nov. 2018b). "Electrocardiogram phenotypes in hypertrophic cardiomyopathy caused by distinct mechanisms: apico-basal repolarization gradients vs. Purkinje-myocardial coupling abnormalities". en. In: *EP Europace* 20.suppl_3. Publisher: Oxford Academic, pp. iii102–iii112. URL: https://academic.oup.com/europace/article/20/suppl_3/iii102/5202143 (visited on 06/09/2020).

Lyon, Aurore, Ana Mincholé, Juan Pablo Martínez, Pablo Laguna, and Blanca Rodriguez (Jan. 2018c). "Computational techniques for ECG analysis and interpretation in light of their contribution to medical advances". In: *Journal of the Royal Society Interface* 15.138. URL: https://www.ncbi.nlm.nih.gov/pmc/articles/PMC5805987/ (visited on 08/05/2020).

Majumdar, Angshul and Rabab Ward (2017). "Robust greedy deep dictionary learning for ECG arrhythmia classification". In: *Proceedings of the International Joint Conference on Neural Networks*. Vol. 2017-May. IEEE, pp. 4400–4407. URL: http://ieeexplore.ieee.org/document/7966413/.

Mamoshina, Polina, Armando Vieira, Evgeny Putin, and Alex Zhavoronkov (2016). "Applications of Deep Learning in Biomedicine". In: *Molecular Pharmaceutics* 13.5, pp. 1445–1454. URL: http://pubs.acs.org/doi/10.1021/acs.molpharmaceut.5b00982.

Maršánová, Lucie, Marina Ronzhina, Radovan Smíšek, Martin Vítek, Andrea Němcová, Lukas Smital, and Marie Nováková (2017). "ECG features and methods for automatic classification of ventricular premature and ischemic heartbeats: A comprehensive experimental study". In: *Scientific Reports* 7.1, p. 11239. URL: http://www.nature.com/articles/s41598-017-10942-6.

Martínez, Arturo, Raúl Alcaraz, and José Joaquín Rieta (2010). "Application of the phasor transform for automatic delineation of single-lead ECG fiducial points". In: *Physiological Measurement* 31.11, pp. 1467–1485. URL: http://stacks.iop.org/0967-3334/31/i=11/a=005?key=crossref.d09c3cb8e45c1c9ffeff9ca7f03e3d41.

Martínez, Juan Pablo, Rute Almeida, Salvador Olmos, Ana Paula Rocha, and Pablo Laguna (2004). "A Wavelet-Based ECG Delineator Evaluation on Standard Databases". In: *IEEE Transactions on Biomedical Engineering* 51.4, pp. 570–581.

Martinez-Navarro, Hector, Xin Zhou, Alfonso Bueno-Orovio, and Blanca Rodriguez (Feb. 2021). "Electrophysiological and anatomical factors determine arrhythmic risk in acute myocardial ischaemia and its modulation by sodium current availability". In: *Interface Focus* 11.1. Publisher: Royal Society, p. 20190124. URL: https://royalsocietypublishing.org/doi/full/10.1098/rsfs.2019.0124 (visited on 01/11/2021).

Medtronic. *Cardioinsight Mapping Vest*. en. URL: https://europe.medtronic.com/xd-en/healthcare-professionals/products/cardiac-rhythm/cardiac-mapping/cardioinsight-mapping-vest.html (visited on 08/21/2020).

Mincholé, Ana, Julià Camps, Aurore Lyon, and Blanca Rodríguez (Nov. 2019a). "Machine learning in the electrocardiogram". en. In: *Journal of Electrocardiology* 57, S61–S64. URL: http://www.sciencedirect.com/science/article/pii/S0022073619304571 (visited on 12/08/2020).

Mincholé, Ana, Ernesto Zacur, Rina Ariga, Vicente Grau, and Blanca Rodriguez (2019b). "MRI-Based Computational Torso/Biventricular Multiscale Models to Investigate the Impact of Anatomical Variability on the ECG QRS Complex". English. In: *Frontiers in Physiology* 10. Publisher: Frontiers. URL: https://www.frontiersin.org/articles/10.3389/fphys.2019.01103/full#supplementary-material (visited on 06/09/2020).

Mitchell, Colleen C. and David G. Schaeffer (2003). "A two-current model for the dynamics of cardiac membrane". In: *Bulletin of Mathematical Biology* 65.5, pp. 767–793. URL: http://link.springer.com/10.1016/S0092-8240(03)00041-7.

Molléro, Roch, Xavier Pennec, Hervé Delingette, Alan Garny, Nicholas Ayache, and Maxime Sermesant (2018). "Multifidelity-CMA: A multifidelity approach for efficient personalisation of 3D cardiac electromechanical models". In: *Biomechanics and Modeling in Mechanobiology* 17.1, pp. 285–300. URL: http://link.springer.com/10.1007/s10237-017-0960-0.

Moody, G B and R G Mark (2001). "The impact of the MIT-BIH arrhythmia database". In: *IEEE engineering in medicine and biology magazine : the quarterly magazine of the Engineering in Medicine & Biology Society* 20.3, pp. 45–50. URL: http://www.ncbi.nlm.nih.gov/pubmed/11446209.

Moody, GB, H Koch, and U Steinhoff (2006). "The PhysioNet/Computers in Cardiology Challenge 2006: QT Interval Measurement". In: *Computers in Cardiology 2006*.

Neic, Aurel, Fernando O. Campos, Anton J. Prassl, Steven A. Niederer, Martin J. Bishop, Edward J. Vigmond, and Gernot Plank (Oct. 2017). "Efficient computation of electrograms and ECGs in human whole heart simulations using a reaction-eikonal model". In: *Journal of computational physics* 346, pp. 191–211. URL: https://www.ncbi.nlm.nih.gov/pmc/articles/PMC5555399/ (visited on 11/16/2020).

Nerbonne, Jeanne M. and Robert S. Kass (Oct. 2005). "Molecular Physiology of Cardiac Repolarization". In: *Physiological Reviews* 85.4. Publisher: American Physiological Society, pp. 1205–1253. URL: https://journals.physiology.org/doi/full/10.1152/physrev.00002.2005 (visited on 03/23/2021).

Nguyên, Uyên Châu, Mark Potse, François Regoli, Maria Luce Caputo, Giulio Conte, Romina Murzilli, Stefano Muzzarelli, Tiziano Moccetti, Enrico G. Caiani, Frits W. Prinzen, Rolf Krause, and Angelo Auricchio (July 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". en. In: *Journal of Electrocardiology* 48.4, pp. 617–625. URL: http://www.sciencedirect.com/science/article/pii/S0022073615001302 (visited on 01/03/2021).

Niederer, Steven A., Joost Lumens, and Natalia A. Trayanova (Feb. 2019). "Computational models in cardiology". In: *Nature reviews. Cardiology* 16.2, pp. 100–111. URL: https://www.ncbi.nlm.nih.gov/pmc/articles/PMC6556062/ (visited on 03/16/2021).

Nygårds, M. E. and L. Sörnmo (1983). "Delineation of the QRS complex using the envelope of the e.c.g." In: *Medical & Biological Engineering & Computing* 21.5, pp. 538–547. URL: http://link.springer.com/10.1007/BF02442378.

O'Hara, Thomas, László Virág, András Varró, and Yoram Rudy (May 2011). "Simulation of the Undiseased Human Cardiac Ventricular Action Potential: Model Formulation and Experimental Validation". en. In: *PLOS Computational Biology* 7.5. Publisher: Public Library of Science, e1002061. URL: https://journals.plos.org/ploscompbiol/article?id=10.1371/journal.pcbi.1002061 (visited on 12/09/2020).

Oord, Aaron van den, Nal Kalchbrenner, Lasse Espeholt, koray kavukcuoglu, Oriol Vinyals, and Alex Graves (2016a). "Conditional Image Generation with PixelCNN Decoders". In: *Advances in Neural Information Processing Systems 29*. Ed. by D. D. Lee, M. Sugiyama, U. V. Luxburg, I. Guyon, and R. Garnett. Curran Associates, Inc., pp. 4790–4798. URL: http://papers.nips.cc/paper/6527-conditional-image-generation-with-pixelcnn-decoders.pdf.

Oord, Aaron van den, Sander Dieleman, Heiga Zen, Karen Simonyan, Oriol Vinyals, Alex Graves, Nal Kalchbrenner, Andrew Senior, and Koray Kavukcuoglu (2016b). "WaveNet: A Generative Model for Raw Audio". In: *arXiv preprint arXiv:1609.03499*.

Peimankar, Abdolrahman and Sadasivan Puthusserypady (Mar. 2021). "DENS-ECG: A deep learning approach for ECG signal delineation". en. In: *Expert Systems with Applications* 165, p. 113911. URL: https://www.sciencedirect.com/science/article/pii/S0957417420307065 (visited on 07/06/2021).

Pereira, Helder, Steven Niederer, and Christopher A Rinaldi (Aug. 2020). "Electrocardiographic imaging for cardiac arrhythmias and resynchronization therapy". In: *Europace* 22.10, pp. 1447–1462. URL: https://www.ncbi.nlm.nih.gov/pmc/articles/PMC7544539/ (visited on 01/03/2021).

Pernod, E., M. Sermesant, E. Konukoglu, J. Relan, H. Delingette, and N. Ayache (2011). "A multi-front eikonal model of cardiac electrophysiology for interactive simulation of radio-frequency ablation". In: *Computers and Graphics (Pergamon)* 35.2, pp. 431–440. URL: https://www.sciencedirect.com/science/article/pii/S0097849311000173.

Pezzuto, Simone, Peter Kal'avský, Mark Potse, Frits W. Prinzen, Angelo Auricchio, and Rolf Krause (2017). "Evaluation of a Rapid Anisotropic Model for ECG Simulation". English. In: *Frontiers in Physiology* 8. Publisher: Frontiers. URL: https://www.frontiersin.org/articles/10.3389/fphys.2017.00265/full (visited on 01/03/2021).

Piersanti, Roberto, Pasquale C. Africa, Marco Fedele, Christian Vergara, Luca Dedè, Antonio F. Corno, and Alfio Quarteroni (Jan. 2021). "Modeling cardiac muscle fibers in ventricular and atrial electrophysiology simulations". en. In: *Computer Methods in Applied Mechanics and Engineering* 373, p. 113468. URL: http://www.sciencedirect.com/science/article/pii/S0045782520306538 (visited on 10/24/2020).

Pitt-Francis, Joe, Pras Pathmanathan, Miguel O. Bernabeu, Rafel Bordas, Jonathan Cooper, Alexander G. Fletcher, Gary R. Mirams, Philip Murray, James M. Osborne, Alex Walter, S. Jon Chapman, Alan Garny, Ingeborg M. M. van Leeuwen, Philip K. Maini, Blanca Rodríguez, Sarah L. Waters,

Jonathan P. Whiteley, Helen M. Byrne, and David J. Gavaghan (Dec. 2009). "Chaste: A test-driven approach to software development for biological modelling". en. In: *Computer Physics Communications*. 40 YEARS OF CPC: A celebratory issue focused on quality software for high performance, grid and novel computing architectures 180.12, pp. 2452–2471. URL: http://www.sciencedirect.com/science/article/pii/S0010465509002604 (visited on 01/04/2021).

Plonsey, Robert and Roger C. Barr (2000). "Bioelectricity: A Quantitative Approach". en. In: *Bioelectricity: A Quantitative Approach*. 2nd ed. Springer US, pp. 218–222. URL: https://www.springer.com/gp/book/9781475731521 (visited on 10/29/2020).

Potse, M., B. Dube, J. Richer, A. Vinet, and R. M. Gulrajani (Dec. 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. Conference Name: IEEE Transactions on Biomedical Engineering, pp. 2425–2435.

Potse, Mark (2018). "Scalable and Accurate ECG Simulation for Reaction-Diffusion Models of the Human Heart". English. In: *Frontiers in Physiology* 9. Publisher: Frontiers. URL: https://www.frontiersin.org/articles/10.3389/fphys.2018.00370/full (visited on 01/03/2021).

Potse, Mark and Nico H L Kuijpers (Sept. 2010). "Simulation of fractionated electrograms at low spatial resolution in large-scale heart models". In: *2010 Computing in Cardiology*. ISSN: 2325-8853, pp. 849–852.

Potyagaylo, Danila, Elisenda Gil Cortés, Walther H. W. Schulze, and Olaf Dössel (Sept. 2014). "Binary optimization for source localization in the inverse problem of ECG". en. In: *Medical & Biological Engineering & Computing* 52.9, pp. 717–728. URL: https://doi.org/10.1007/s11517-014-1176-4 (visited on 01/11/2021).

Proakis, John G. and Dimitris G. Manolakis (1992). *Digital Signal Processing (2nd Ed.): Principles, Algorithms, and Applications*. USA: Macmillan Publishing Co., Inc.

Rajpurkar, Pranav, Awni Y. Hannun, Masoumeh Haghpanahi, Codie Bourn, and Andrew Y. Ng (2017). "Cardiologist-Level Arrhythmia Detection with Convolutional Neural Networks". In: arXiv: 1707.01836. URL: http://arxiv.org/abs/1707.01836.

Ronneberger, Olaf, Philipp Fischer, and Thomas Brox (May 2015). "U-Net: Convolutional Networks for Biomedical Image Segmentation". en. In: *arXiv:1505.04597 [cs]*. arXiv: 1505.04597. URL: http://arxiv.org/abs/1505.04597 (visited on 07/06/2021).

Rossi, Simone, Toni Lassila, Ricardo Ruiz-Baier, Adélia Sequeira, and Alfio Quarteroni (Nov. 2014). "Thermodynamically consistent orthotropic activation model capturing ventricular systolic wall thickening in cardiac electromechanics". en. In: *European Journal of Mechanics - A/Solids*. Frontiers in Finite-Deformation Electromechanics 48, pp. 129–142. URL: http://www.sciencedirect.com/science/article/pii/S0997753813001228 (visited on 01/05/2021).

Saffitz, J E, H L Kanter, K G Green, T K Tolley, and E C Beyer (June 1994). "Tissue-specific determinants of anisotropic conduction velocity in canine atrial and ventricular myocardium." en. In: *Circulation Research* 74.6, pp. 1065–1070. URL:

https://www.ahajournals.org/doi/10.1161/01.RES.74.6.1065 (visited on 07/05/2021).

Sakoe, H. and S. Chiba (Feb. 1978). "Dynamic programming algorithm optimization for spoken word recognition". In: *IEEE Transactions on Acoustics, Speech, and Signal Processing* 26.1. Conference Name: IEEE Transactions on Acoustics, Speech, and Signal Processing, pp. 43–49.

Sarubbi, Berardo, Wei Li, and Jane Somerville (Aug. 2000). "QRS width in right bundle branch block. Accuracy and reproducibility of manual measurement". en. In: *International Journal of Cardiology* 75.1, pp. 71–74. URL: https://www.sciencedirect.com/science/article/pii/S0167527300002990 (visited on 07/05/2021).

Schaufelberger, Matthias, Steffen Schuler, Laura Bear, Matthijs Cluitmans, Jaume Coll-Font, Önder Nazim Onak, Olaf Dössel, and Dana Brooks (Sept. 2019). "Comparison of Activation Times Estimation for Potential-Based ECG Imaging". In: *Computing in cardiology* 46. URL: https://www.ncbi.nlm.nih.gov/pmc/articles/PMC7079739/ (visited on 01/20/2021).

Schreier, Günter, Dieter Hayn, and Suave Lobodzinski (2003). "Development of a New QT Algorithm with Heterogenous ECG Databases". In: *Journal of Electrocardiology*. Vol. 36. SUPPL. Churchill Livingstone, pp. 145–150. URL: https://www.sciencedirect.com/science/article/pii/S0022073603001274.

Schuler, Steffen, Jess D. Tate, Thom F. Oostendorp, Robert S. MacLeod, and Olaf Dössel (2019). "Spatial Downsampling of Surface Sources in the Forward Problem of Electrocardiography". en. In: *Functional Imaging and Modeling of the Heart*. Ed. by Yves Coudière, Valéry Ozenne, Edward Vigmond, and Nejib Zemzemi. Lecture Notes in Computer Science. Cham: Springer International Publishing, pp. 29–36.

Sebastian, R., V. Zimmerman, D. Romero, D. Sanchez-Quintana, and A. F. Frangi (Jan. 2013). "Characterization and Modeling of the Peripheral Cardiac Conduction System". In: *IEEE Transactions on Medical Imaging* 32.1. Conference Name: IEEE Transactions on Medical Imaging, pp. 45–55.

Serinagaoglu Dogrusoz, Yesim, Laura Bear, Jake Bergquist, Remi Dubois, Wilson Good, Rob MacLeod, Ali Rababah, and Job Stoks (Dec. 2019). "Effects of Interpolation on the Inverse Problem of Electrocardiography". en. In: URL: http://www.cinc.org/archives/2019/pdf/CinC2019-100.pdf (visited on 08/04/2020).

Silver, David, Aja Huang, Chris J Maddison, Arthur Guez, Laurent Sifre, George Van Den Driessche, Julian Schrittwieser, Ioannis Antonoglou, Veda Panneershelvam, Marc Lanctot, et al. (2016). "Mastering the game of Go with deep neural networks and tree search". In: *Nature* 529.7587, pp. 484–489.

Singh, Yogendra Narain and Phalguni Gupta (2008). "ECG to individual identification". In: *BTAS 2008 - IEEE 2nd International Conference on Biometrics: Theory, Applications and Systems* 00.

Sisson, S. A., Y. Fan, and Mark M. Tanaka (Feb. 2007). "Sequential Monte Carlo without likelihoods". en. In: *Proceedings of the National Academy of Sciences* 104.6. ISBN: 9780607208108 Publisher: National Academy of Sciences Section: Physical Sciences, pp. 1760–1765. URL: https://www.pnas.org/content/104/6/1760 (visited on 09/09/2020).

Somani, Sulaiman, Adam J Russak, Felix Richter, Shan Zhao, Akhil Vaid, Fayzan Chaudhry, Jessica K De Freitas, Nidhi Naik, Riccardo Miotto, Girish N Nadkarni, Jagat Narula, Edgar Argulian, and Benjamin S Glicksberg (Aug. 2021). "Deep learning and the electrocardiogram: review of the current state-of-the-art". In: *EP Europace* 23.8, pp. 1179–1191. URL: https://doi.org/10.1093/europace/euaa377 (visited on 09/02/2021).

South, L. F., A. N. Pettitt, and C. C. Drovandi (Sept. 2019). "Sequential Monte Carlo Samplers with Independent Markov Chain Monte Carlo Proposals". EN. In: *Bayesian Analysis* 14.3. Publisher: International Society for Bayesian Analysis, pp. 753–776. URL: https://projecteuclid.org/euclid.ba/1560240027 (visited on 09/09/2020).

Srivastava, Nitish, Geoffrey E Hinton, Alex Krizhevsky, Ilya Sutskever, and Ruslan Salakhutdinov (2014). "Dropout: a simple way to prevent neural networks from overfitting." In: *Journal of Machine Learning Research* 15.1, pp. 1929–1958.

Streeter, Daniel D., Henry M. Spotnitz, Dali P. Patel, John Ross, and Edmund H. Sonnenblick (Mar. 1969). "Fiber Orientation in the Canine Left Ventricle during Diastole and Systole". en. In: *Circulation Research* 24.3, pp. 339–347. URL: https://www.ahajournals.org/doi/10.1161/01.RES.24.3.339 (visited on 07/07/2020).

Sudlow, Cathie, John Gallacher, Naomi Allen, Valerie Beral, Paul Burton, John Danesh, Paul Downey, Paul Elliott, Jane Green, Martin Landray, Bette Liu, Paul Matthews, Giok Ong, Jill Pell, Alan Silman, Alan Young, Tim Sprosen, Tim Peakman, and Rory Collins (Mar. 2015). "UK Biobank: An Open Access Resource for Identifying the Causes of a Wide Range of Complex Diseases of Middle and Old Age". In: *PLoS Medicine* 12.3. URL: https://www.ncbi.nlm.nih.gov/pmc/articles/PMC4380465/ (visited on 11/17/2020).

Sun, Yan, Kap Luk Chan, and Shankar Muthu Krishnan (2005). "Characteristic wave detection in ECG signal using morphological transform". In: *BMC Cardiovascular Disorders* 5.1, p. 28. URL: http://bmccardiovascdisord.biomedcentral.com/articles/10.1186/1471-2261-5-28.

Szegedy, Christian, Wei Liu, Yangqing Jia, Pierre Sermanet, Scott Reed, Dragomir Anguelov, Dumitru Erhan, Vincent Vanhoucke, and Andrew Rabinovich (2015). "Going Deeper With Convolutions". In: *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*.

Sörnmo, Leif and Pablo Laguna (2005). *Bioelectrical Signal Processing in Cardiac and Neurological Applications*. en. Elsevier. URL: https://linkinghub.elsevier.com/retrieve/pii/B9780124375529X50004 (visited on 07/13/2020).

Taddei, A, G Distante, M Emdin, P Pisani, GB Moody, C Zeelenberg, and C Marchesi (1992). "The European ST-T database: standard for evaluating systems for the analysis of ST-T changes in ambulatory electrocardiography". In: *European Heart Journal*, pp. 1164–1172. URL: https://academic.oup.com/eurheartj/article-abstract/13/9/1164/451781.

Taigman, Yaniv, Ming Yang, Marc'Aurelio Ranzato, and Lior Wolf (2014). "Deepface: Closing the gap to human-level performance in face verification". In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 1701–1708.

Tan, Jen Hong, Yuki Hagiwara, Winnie Pang, Ivy Lim, Shu Lih Oh, Muhammad Adam, Ru San Tan, Ming Chen, and U. Rajendra Acharya (2018). "Application of stacked convolutional and long short-term memory network for accurate identification of CAD ECG signals". In: *Computers in Biology and Medicine* 94, pp. 19–26. URL: https://www.sciencedirect.com/science/article/pii/S0010482517304201.

Tate, Jess D., Steffen Schuler, Olaf Dössel, Rob S. MacLeod, and Thom F. Oostendorp (June 2019). "Correcting Undersampled Cardiac Sources in Equivalent Double Layer Forward Simulations". In: *Functional imaging and modeling of the heart : ... International Workshop, FIMH ..., proceedings. FIMH* 11504, pp. 147–155. URL: https://www.ncbi.nlm.nih.gov/pmc/articles/PMC6889815/ (visited on 07/13/2020).

The CSE Working Party (1985). "Recommendations for measurement standards in quantitative electrocardiography". In: *European Heart Journal* 6.10, pp. 815–825. URL: https://academic.oup.com/eurheartj/article/6/10/815/460562/Recommendations-for-measurement-standards-in.

Thomas, LJ, KW Clark, CN Mead, and Undefined Ripley, KL (1979). "Automated cardiac dysrhythmia analysis". In: *Proceedings of the IEEE* 67.9, pp. 1322 –1337. URL: https://ieeexplore.ieee.org/abstract/document/1455719/.

Tieleman, Tijmen and Geoffrey Hinton (2012). "Lecture 6.5-rmsprop: Divide the gradient by a running average of its recent magnitude". In: *COURSERA: Neural networks for machine learning* 4.2.

Tikhonov, Andrey N and Vasiliy Y Arsenin (1977). "Solutions of ill-posed problems". In: *New York*, pp. 1–30.

Tomek, Jakub, Alfonso Bueno-Orovio, Elisa Passini, Xin Zhou, Ana Minchole, Oliver Britton, Chiara Bartolucci, Stefano Severi, Alvin Shrier, Laszlo Virag, Andras Varro, and Blanca Rodriguez (Dec. 2019). "Development, calibration, and validation of a novel human ventricular myocyte model in health, disease, and drug block". In: *eLife* 8. Ed. by José D Faraldo-Gómez, Naama Barkai, Thomas Hund, and Molly Maleckar. Publisher: eLife Sciences Publications, Ltd, e48890. URL: https://doi.org/10.7554/eLife.48890 (visited on 12/09/2020).

Tornvall, P., E. Gerbaud, A. Behaghel, R. Chopard, O. Collste, E. Laraudogoitia, G. Leurent, N. Meneveau, M. Montaudon, E. Perez-David, P. Sörensson, and S. Agewall (July 2015). "Myocarditis or "true" infarction by cardiac magnetic resonance in patients with a clinical diagnosis of myocardial infarction without obstructive coronary disease: A meta-analysis of individual patient data". en. In: *Atherosclerosis* 241.1, pp. 87–91. URL: https://www.sciencedirect.com/science/article/pii/S0021915015010345 (visited on 03/19/2021).

Turing, Alan (1949). "The Mechanical Brain. Answer Found to 300 Year Old Problem". In: *The Times newspaper*, p. 4.

Vančura, Vlastimil, Dan Wichterle, Ivan Ulč, Jiří Šmíd, Marek Brabec, Marta Zárybnická, and Richard Rokyta (2016). "The variability of automated QRS duration measurement". In: *Europace* 19.4, euw015. URL: https://academic.oup.com/europace/article-lookup/doi/10.1093/europace/euw015.

Vázquez-Seisdedos, Carlos R, João Neto, Enrique J Marañón Reyes, Aldebaro Klautau, and Roberto C Limão de Oliveira (2011). "New approach for T-wave end detection

on electrocardiogram: Performance in noisy conditions". In: *BioMedical Engineering OnLine* 10.1, p. 77. URL: http://biomedical-engineering-online.biomedcentral.com/articles/10.1186/1475-925X-10-77.

Vigmond, Edward J. and Bruno D. Stuyvers (Jan. 2016). "Modeling our understanding of the His-Purkinje system". en. In: *Progress in Biophysics and Molecular Biology.* Recent Developments in Biophysics & Molecular Biology of Heart Rhythm 120.1, pp. 179–188. URL: http://www.sciencedirect.com/science/article/pii/S007961071500262X (visited on 06/16/2020).

Vítek, M., J. Hrubeš, and J. Kozumplík (2009). "A Wavelet-Based ECG Delineation in Multilead ECG Signals: Evaluation on the CSE Database". In: Springer, Berlin, Heidelberg, pp. 177–180. URL: http://link.springer.com/10.1007/978-3-642-03882-2{_}46.

Vitek, Martin, Jan Hrubes, and Jiri Kozumplik (2010). "A Wavelet-Based ECG Delineation with Improved P Wave Offset Detection Accuracy". In: *Analysis of Biomedical Signals and Images* 20.1, pp. 1–6. URL: http://www.biosignal.cz/bs2010/papers/1042.pdf.

Wallman, Mikael, Nicolas P. Smith, and Blanca Rodriguez (June 2012). "A Comparative Study of Graph-Based, Eikonal, and Monodomain Simulations for the Estimation of Cardiac Activation Times". In: *IEEE Transactions on Biomedical Engineering* 59.6. Conference Name: IEEE Transactions on Biomedical Engineering, pp. 1739–1748.

— (Jan. 2014). "Computational methods to reduce uncertainty in the estimation of cardiac conduction properties from electroanatomical recordings". en. In: *Medical Image Analysis* 18.1, pp. 228–240. URL: http://www.sciencedirect.com/science/article/pii/S1361841513001503 (visited on 07/10/2020).

Wang, Yong and Yoram Rudy (Aug. 2006). "Application of the Method of Fundamental Solutions to Potential-based Inverse Electrocardiography". en. In: *Annals of Biomedical Engineering* 34.8. Company: Springer Distributor: Springer Institution: Springer Label: Springer Number: 8 Publisher: Kluwer Academic Publishers-Plenum Publishers, pp. 1272–1288. URL: https://link.springer.com/article/10.1007/s10439-006-9131-7 (visited on 01/11/2021).

Weston, Jason E (2016). "Dialog-based Language Learning". In: *Advances in Neural Information Processing Systems 29.* Ed. by D. D. Lee, M. Sugiyama, U. V. Luxburg, I. Guyon, and R. Garnett. Curran Associates, Inc., pp. 829–837. URL: http://papers.nips.cc/paper/6264-dialog-based-language-learning.pdf.

Willems, Jos L., Pierre Arnaud, Jan H. van Bemmel, Peter J. Bourdillon, Rosanna Degani, Bernard Denis, Frits M A Harms, Peter W. Macfarlane, Gianfranco Mazzocca, Jürgen Meyer, Henk J. Ritsema van Eck, Etienne O. Robles de Medina, and Christoph Zywietz (1985). "Establishment of a reference library for evaluating computer ECG measurement programs". In: *Computers and Biomedical Research* 18.5, pp. 439–457. URL: https://www.sciencedirect.com/science/article/pii/0010480985900217.

Zeiler, Matthew D. (2012). "ADADELTA: An Adaptive Learning Rate Method". In: *CoRR* abs/1212.5701. URL: http://arxiv.org/abs/1212.5701.

Zettinig, Oliver, Tommaso Mansi, Bogdan Georgescu, Elham Kayvanpour, Farbod Sedaghat-Hamedani, Ali Amr, Jan Haas, Henning Steen, Benjamin Meder,

Hugo Katus, Nassir Navab, Ali Kamen, and Dorin Comaniciu (2013). "Fast data-driven calibration of a cardiac electrophysiology model from images and ECG". In: *Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)*. Vol. 8149 LNCS. PART 1. Springer, Berlin, Heidelberg, pp. 1–8. URL: http://link.springer.com/10.1007/978-3-642-40811-3{_}1.

Zettinig, Oliver, Tommaso Mansi, Dominik Neumann, Bogdan Georgescu, Saikiran Rapaka, Philipp Seegerer, Elham Kayvanpour, Farbod Sedaghat-Hamedani, Ali Amr, Jan Haas, Henning Steen, Hugo Katus, Benjamin Meder, Nassir Navab, Ali Kamen, and Dorin Comaniciu (2014). "Data-driven estimation of cardiac electrical diffusivity from 12-lead ECG signals". In: *Medical Image Analysis* 18.8, pp. 1361–1376. URL: https://www.sciencedirect.com/science/article/pii/S1361841514000656.