



Learning cardiac electrophysiology dynamics with PDE-based physiological constraints for data-driven personalised predictions

Victoriya Kashtanova

► To cite this version:

Victoriya Kashtanova. Learning cardiac electrophysiology dynamics with PDE-based physiological constraints for data-driven personalised predictions. Artificial Intelligence [cs.AI]. Université Côte d'Azur, 2023. English. NNT : 2023COAZ4043 . tel-04176798v1

HAL Id: tel-04176798

<https://hal.science/tel-04176798v1>

Submitted on 3 Aug 2023 (v1), last revised 14 Sep 2023 (v2)

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



THÈSE DE DOCTORAT

Apprentissage de la dynamique électrophysiologique
cardiaque avec contraintes physiologiques de modèle
pour une prédiction personnalisée basée sur les données

Victoriya KASHTANOVA

INRIA, Équipe EPIONE

Thèse dirigée par Maxime SERMESANT et co-encadrée par Patrick GALLINARI

Soutenue le 14 Juin 2023

Présentée en vue de l'obtention du grade de DOCTEUR EN AUTOMATIQUE, TRAITEMENT
DU SIGNAL ET DES IMAGES de l'UNIVERSITÉ CÔTE D'AZUR.

Devant le jury composé de :

Stéphane LANTERI	Centre Inria d'Université Côte d'Azur	Président
Ender KONUKOGLU	ETH-Zurich	Rapporteur
Luca DEDE'	Politecnico di Milano	Rapporteur
Patrick GALLINARI	Sorbonne Université	Co-encadrant
Maxime SERMESANT	Centre Inria d'Université Côte d'Azur	Directeur de thèse

PHD THESIS

Learning Cardiac Electrophysiology Dynamics with PDE-based Physiological Constraints for Data-Driven Personalised Predictions

Victoriya KASHTANOVA

CENTRE INRIA D'UNIVERSITÉ CÔTE D'AZUR, EPIONE Team

Supervised by Maxime SERMESANT
and Patrick GALLINARI

Defended on the 14th of June, 2023

Presented to obtain the title of DOCTEUR EN AUTOMATIQUE, TRAITEMENT DU SIGNAL ET
DES IMAGES of the UNIVERSITÉ CÔTE D'AZUR.

Jury:

Stéphane LANTERI	Centre Inria d'Université Côte d'Azur	President
Ender KONUKOGLU	ETH-Zurich	Reviewer
Luca DEDE'	Politecnico di Milano	Reviewer
Patrick GALLINARI	Sorbonne Université	Co-supervisor
Maxime SERMESANT	Centre Inria d'Université Côte d'Azur	Supervisor

Résumé

Un défi scientifique majeur actuel consiste à combiner la polyvalence des approches intensives basées sur les données avec les approches de modélisation à base physique développées dans des domaines scientifiques tels que la biophysique. La modélisation biophysique du cœur humain offre un cadre bien développé pour introduire la physiologie dans l'analyse prédictive des données cliniques. Cependant, les modèles numériques de haute précision développés souffrent souvent de coûts de calcul importants sur le terrain. En conséquence, les chercheurs utilisent souvent des modèles simplifiés basés sur des hypothèses idéalisées qui ne peuvent pas faire face à la diversité et à la complexité des conditions spécifiques aux patients. D'autre part, les méthodes de Machine Learning et de Deep Learning (qui ont déjà obtenu des résultats impressionnants dans plusieurs domaines comme la vision, le traitement du langage naturel, la bio-informatique, etc.) manquent souvent d'interprétabilité et de robustesse, et ne permettent pas une intégration aisée des connaissances préalables disponibles dans de nombreux domaines scientifiques.

L'objectif scientifique de cette thèse est de combiner les avantages de la biophysique et des méthodes de Deep Learning (DL), afin de développer des modèles hybrides exploitant la complémentarité des deux approches. Notre objectif est d'introduire des a priori physiologiques dans les systèmes d'apprentissage par la modélisation biophysique en apprenant des dynamiques spatio-temporelles à partir de simulations et en introduisant des contraintes physiquement motivées relatives à ces dynamiques.

Pour ce faire, nous proposons deux frameworks basés sur DL pour l'apprentissage automatique de la dynamique de l'électrophysiologie cardiaque (EP) à partir de données.

Nous commençons par présenter un framework EP-Net 2.0 entièrement basé sur les données, capable d'apprendre et de prédire la dynamique de l'EP cardiaque en présence de cicatrices non excitables modélisées dans une plaque de tissu cardiaque. À l'aide de données synthétiques, nous démontrons que ce framework peut reproduire une dynamique complexe du potentiel transmembranaire cardiaque, même en dehors de son domaine d'entraînement.

Ensuite, nous améliorons ce framework avec des a priori physiques afin de viser des horizons de prévision plus larges et d'obtenir une certaine interprétabilité. Ce nouveau framework d'apprentissage en profondeur basé sur la physique (APHYN-EP) peut apprendre la dynamique de l'EP cardiaque à partir de données de différentes complexités.

En utilisant des données synthétiques, nous démontrons que le framework APHYN-EP peut reproduire la dynamique complexe du potentiel transmembranaire même en présence de bruit dans les données. De plus, en utilisant des données ex vivo de cartographie optique du potentiel d'action, nous montrons que notre framework : a) peut identifier les paramètres physiques clés pour différentes zones anatomiques ayant une fonction électrique anormale ; b) est capable de reproduire les caractéristiques de propagation du potentiel d'action obtenues à différents emplacements de stimulation.

Dans l'ensemble, nos nouvelles approches combinées basées sur des modèles et axées sur les données ont démontré le potentiel d'amélioration de la modélisation de l'EP cardiaque et d'obtention d'outils de calcul prédictifs robustes et reproductibles. Nous prévoyons que notre framework polyvalent et translationnel sera intégré dans des modèles informatiques cliniques d'EP pour des prédictions plus précises de l'arythmie.

Mots-clés: Apprentissage basé sur la physique, Apprentissage profond, Electrophysiologie cardiaque, Personnalisation de modèle, EDP, Simulations.

Abstract

A current major scientific challenge consists in combining the versatility of intensive data-based approaches with the physically grounded modelling approaches developed in scientific fields such as biophysics. Biophysical modelling of the human heart offers a well-developed framework for introducing physiology into the predictive analysis of clinical data. However, high precision numerical models developed in the field often suffer from large computational costs. As a consequence, researchers often make use of simplified models based on idealized assumptions that cannot cope with the diversity and complexity of patient-specific conditions. On the other hand, Machine Learning and Deep Learning methods (which have already achieved impressive results in several domains like vision, natural language processing, bioinformatics, etc.) often have a lack of interpretability and robustness, and do not allow an easy integration of prior knowledge available in many scientific fields.

The scientific objective of this thesis is to combine the advantages of biophysics and Deep Learning (DL), methods so as to develop hybrid models exploiting the complementarity of the two approaches. Our goal is to introduce physiological priors in learning systems through biophysical modelling by learning spatiotemporal dynamics from simulations and by introducing physically motivated constraints relative to these dynamics.

To achieve this, we propose two DL based frameworks for the automatic learning of cardiac electrophysiology (EP) dynamics from data.

We start by presenting a fully data-driven EP-Net 2.0 framework which is able to learn and forecast the cardiac EP dynamics in the presence of unexcitable scars modelled in a cardiac tissue slab. Using synthetic data, we demonstrate that this framework can reproduce complex cardiac transmembrane potential dynamics, even outside of its training domain.

Next, we extend this framework with physical priors so as to target larger forecasting horizons and to obtain a certain interpretability of corresponding models. This novel physics-based deep learning framework (APHYN-EP) can learn cardiac EP dynamics from data of different complexities.

Using synthetic data, we demonstrate that the APHYN-EP framework can reproduce the complex dynamics of transmembrane potential even in presence of noise in the data. Additionally, using ex vivo optical mapping data of action potential, we show that our framework: a) can identify key physical parameters for different anatomical zones having abnormal electrical function; b) is capable to reproduce the action potential wave characteristics obtained at different pacing locations.

Overall, our novel combined model-based and data-driven approaches have demonstrated the potential to improve cardiac EP modelling and to provide predictive robust and reproducible computational tools. We envision that our versatile and translational framework will be integrated into clinical computational EP models for more accurate predictions of arrhythmia.

Keywords: Physics-based learning, Deep learning, Cardiac electrophysiology, Model personalisation, PDE, Simulations.

Acknowledgements

I would like to thank first my thesis supervisors, Dr. Maxime Sermesant and Dr. Patrick Gallinari. Thank you for giving me this opportunity and for all your help and support during these graduate training years. Thank you very much for your ideas, motivation and, especially, your patience which helped me to grow as a researcher. I would like to thank Dr. Mihaela Pop for your colossal expertise in cardiac electrophysiology that you shared with me to cover my lack of knowledge in biology and medicine. Thanks also to Dr. Nicholas Ayache for allowing me to be part of the Epione team, allowing me to enjoy such an inspiring research environment in the wonderful Côte d’Azur.

I would like to thank the members of my thesis jury, and especially my reviewers Dr. Ender Konukoglu and Dr. Luca Dede’ who were kind enough to take time to read this manuscript. I am grateful to Dr. Stephane Lanteri for accepting to be a president of my jury.

I would like to thank the members of MLIA team from Sorbonne University for your support, help, fresh ideas and availability for long scientific discussions during my work trips to Paris.

Many thanks to all the Epione team and the people I have met at INRIA, it has been a pleasure to share meals/coffee breaks and to participate in various scientific seminars and conferences together. Thank you for your support and friendly environment, which helped me to overcome all difficulties during my PhD and to forget that I’m very far away from home.

I would like to thank Dr. Mohammed El Rhabi and Dr. Matthieu Vandamme who believed in me and opened up the possibility of studying in France.

Karim, thank you for being by my side every day and making my PhD days easier. Thank you for your support, your patience, your expertise in French and English and your superpower to solve my dilemmas. A big thanks to Leya, our dog, who is always making me smile and comfort me after long tiring bus trips from the office.

And many thanks to my family, my mother, my grandmother and my sister, I hope you are very proud today.

Financial Support

This work has been supported by the French government, through the 3IA Cote d’Azur Investments in the Future project managed by the National Research Agency (ANR) with the reference number ANR-19-P3IA-0002. The authors are grateful to the OPAL infrastructure from Universite Cote d’Azur.

Contents

1	Introduction	1
1.1	Clinical Context	1
1.1.1	Cardiac Physiology	1
1.1.2	Infarction and Arrhythmia	4
1.1.3	Imaging the Heart Electrical Activity	6
1.2	Cardiac Electrophysiology Modelling: State-of-the-art	9
1.2.1	Physical Models	9
1.2.2	Data-driven Approaches	11
1.2.3	Physics-Based Deep Learning Approaches	11
1.3	Motivation and Scientific Objectives	12
1.4	Main Contributions	13
1.5	Manuscript organisation	14
2	EP-Net 2.0: Data-Driven Model for Cardiac Electrophysiology Learning	15
2.1	Introduction	17
2.2	Problem Formulation	19
2.2.1	Learning Method	20
2.3	Experiments	21
2.3.1	Data Collection	21
2.3.2	Training Settings	21
2.4	Results	21
2.4.1	Testing on rectangular scars	22
2.4.2	Generalisation: scars of various shapes and multiple onsets	22
2.4.3	Generalisation: various conduction velocities	23
2.5	Limitations and Discussion	25
2.6	Conclusion and Future Work	26
2.A	Modelling real scars with domain mask containing multiple tissue conductivities	27
3	APHYN-EP: Physics-Based Deep Learning for Model Error Correction	29
3.1	Introduction	30
3.2	Learning Framework	32
3.3	Experimental Settings	35
3.3.1	In silico datasets creation	35

3.3.2	Training settings	37
3.4	Results	38
3.4.1	0D synthetic data	38
3.4.2	2D synthetic data	40
3.5	Conclusion	42
3.A	Learning and Forecasting from Small Time Sequences	43
3.A.1	Results	43
4	Personalisation of Cardiac Electrophysiology Model on Experimental Signals	47
4.1	Introduction	48
4.2	Materials and Methods	49
4.2.1	Data collection	49
4.2.2	Full personalisation of cardiac EP model	51
4.2.3	Fast personalisation of cardiac EP model	53
4.2.4	Training settings	56
4.3	Experiments and Results	56
4.3.1	Heart with an ischaemic region	56
4.3.2	Different pacing frequencies	63
4.4	Discussion and Conclusion	66
5	Deep Personalisation of Cardiac Electrophysiology Model on 2D Optical Mapping Data.	69
5.1	Introduction	70
5.2	Materials and Methods	72
5.2.1	Data collection	72
5.2.2	Diffusion tensor	72
5.2.3	Deep personalisation of cardiac EP model	75
5.2.4	Training settings	79
5.3	Experiments and Results	80
5.3.1	Isotropic diffusion	80
5.3.2	Anisotropic diffusion	80
5.3.3	Generalisation: prediction extension to the whole heart	83
5.4	Discussion and Conclusion	90
6	Conclusion	93
6.1	Contributions Summary	93
6.1.1	Fully data-driven framework	93
6.1.2	Physics-based deep learning framework	94
6.2	Publications	96
6.3	Discussion and Perspectives	96
	Bibliography	101

Abbreviations:

Medical

EP	electrophysiology
AP	action potential
LV	left ventricle
RV	right ventricle
VT	ventricular tachycardia
RFA	radiofrequency ablation
ROI	region of interest

Networks

ML	machine learning
DL	deep learning
NN	neural network
ResNet	residual neural network
CNN	convolutional neural network
MLP	multi-layer perceptron

Metrics

MSE	mean squared error
-----	--------------------

Introduction

Contents

1.1	Clinical Context	1
1.1.1	Cardiac Physiology	1
1.1.2	Infarction and Arrhythmia	4
1.1.3	Imaging the Heart Electrical Activity	6
1.2	Cardiac Electrophysiology Modelling: State-of-the-art	9
1.2.1	Physical Models	9
1.2.2	Data-driven Approaches	11
1.2.3	Physics-Based Deep Learning Approaches	11
1.3	Motivation and Scientific Objectives	12
1.4	Main Contributions	13
1.5	Manuscript organisation	14

This thesis explores novel methods for computational frameworks to leverage Deep Learning of to improve electrophysiology modeling approaches as well as the personalisation of cardiac action potential (i.e., the transmembrane voltage) features and associated propagating waves.

1.1 Clinical Context

1.1.1 Cardiac Physiology

1.1.1.1 Heart Anatomy

The main function of the heart is to pump the blood into the circulatory system, thus providing oxygenated blood to the organs and in parallel sending de-oxygenated blood to the lungs. It is composed of four chambers; two atria and two ventricles, see Figure 1.1. The right atrium and ventricle provide blood to the lungs while the left ventricle, being a stronger muscle, provides blood to all organs. The left ventricle (LV) and the right ventricle (RV) are separated by an intraventricular muscular wall, called septum. The inferior part of the heart is called the apex. The ventricular cavities are closed by valves, which separate the ventricles from the atria. Finally, the pulmonary and aortic valves

allow the blood to leave the ventricles via aorta and large arteries. The muscular wall of the heart is called myocardium, and has an inner surface (endocardium) and an outer surface (epicardium), as illustrated in Figure 1.2. The surrounding membrane is called pericardium. The myocardium is mainly composed of myocyte cells (or muscle fibers). The heart contraction (based on synchronous RV and LV contraction) ensures an efficient expulsion of the blood, shortening along axes as determined by the fiber orientation. Figure 1.3 shows at the macroscopic level this supra-cellular structure which forms the cardiac fibres.




Fig. 1.1: Anatomy and conduction system of the heart. Illustration from [Open Stax Anatomy & Physiology](#), CC-BY license.

A full cardiac cycle lasts approximately one second and is divided into two phases: the systole during which the ventricles contract, and the diastole during which the ventricles relax. There are also two iso-volumetric phases.

1.1.1.2 Cardiac Electrical Conduction System

The cardiac contraction is driven by an electrical impulse (wave) propagating through the myocardium and depolarising each cell in a sequential order: first the atria and then the ventricles. Due to the particular myocardial fibers' organisation and the density of the gap junctions at the cell end, the electrical propagation is anisotropic and thus the action potential propagates 2-3 times faster along the fibers.

The cardiac rhythm is driven by the Sinoatrial Node (SN), which is capable of spontaneously generating electrical signals due to a chemical imbalance in SN cells. In a normal sinus rhythm, the stimulus reaches the ventricles from the Atrioventricular Node




Fig. 1.2: Anatomy and conduction system of the heart. Illustration from [Wikimedia commons](#), CC-BY license.

through the His bundle where the network splits into two branches: the Left bundle and the Right bundle, respectively (see Figure 1.1). These bundle branches are connected to a fast electrical conduction system called Purkinje fibers, composed of specialised conduction cells. The Purkinje fibers are located just under the endocardial layer and are able to conduct the electrical impulse quickly and efficiently: typical conduction speed ranges from 2m/s to 3m/s, while for myocardial cells it ranges from 0.3 to 0.4m/s [Durrer, 1970].

The resting (polarised) state of a myocyte consists in a voltage difference of approximately 90mV between the intracellular and extracellular medium, mainly defined by the Na^+ , Ca^{++} and K^+ ions. When a stimulus is applied to a myocyte, the transmembrane potential (and polarity) is modified with respect to the gradient of ionic concentrations. The concentrations of different ions (sodium, calcium and potassium) and their flow through the membrane are responsible of particular cell state changes. The evolution of the transmembrane potential is called action potential (AP) and is composed of 3 main phases: the depolarisation phase (phase 0, see Figure 1.4), the plateau phase (phase 2) and the repolarisation phase (phase 3); and is followed by a resting phase (phase 4). It is important to mention here that the sodium channels involved in the depolarisation phase are voltage-dependent: their gates only open if the membrane potential is raised above a certain threshold value. This initial rise in voltage is mediated by nearby ionic movements in the same cell or between cells through specific structures named gap junctions located in the intercalated discs. Thus, the influence of a cell's neighbourhood is what causes the propagation of AP wave through the myocardial tissue.




Fig. 1.3: Cardiac Muscle organisation: a) Schematic structure of cardiac muscle cell connections; b) Photomicrograph of cardiac muscle cells showing the intercalated discs; c) An intercalated disc connects cardiac muscle cells and consists of desmosomes and gap junctions (which allow the ionic flow and propagation of action potential from one myocyte to another). Illustration taken from [Open Stax Anatomy & Physiology](#), CC-BY license.




Fig. 1.4: Ventricular action potential evolution in mV. The main ions involved in each phase are represented next to the curve. Arrows indicate whether the ions enter (up) or leave (down) the cytoplasm through the specific ionic channels located in the cellular membrane. From [\[Cedilnik, 2020\]](#).

1.1.2 Infarction and Arrhythmia

The pumping function of the heart is also responsible for transporting oxygen and nutrients to the heart itself via the arterial coronary system. When this process is

perturbed by occlusions of one or more vessels (Figure 1.5), the ischemic cardiomyocytes' activity becomes inefficient: this process represents the acute phase of infarction. If the occlusion lasts more than several tens of minutes or hours, the myocardial tissue suffers irreversible changes (chronic infarction): the cells become necrotic and are gradually replaced within several weeks by fibrotic tissue (i.e., unexcitable collagen fibrils), whose contractility and electrochemical properties significantly differ from those of healthy myocytes.

The past decades have seen dramatic improvements in the survival rates of acute infarcts, due to quick interventional response such as the immediate re-vascularisation. However, the permanent damage of the cardiac tissue is responsible for a large part of infarct-related mortality and morbidity. These chronic affections include heart failure when they are related to the pumping function of the heart, as well as arrhythmia events which are related to the pathological remodelling of electrical properties (e.g. modification of AP characteristics, and the speed of AP wave in ischemic tissue).

One particular infarct-related ventricular arrhythmia case is the re-entrant ventricular tachycardia (VT). Especially in the case of chronic infarct-related VT, the depolarisation wave re-excites (or re-enters) segments of the myocardial tissue that are out of their refractory period. The wave propagates through channels of slow conduction (i.e., reduced speed values due to the closure of some gap junctions) and revolves around unexcitable fibrotic scars. As a result, the ventricle enters a self-sustained, rapid but inefficient contraction of the heart, short-circuiting the physiological role of the Sinoatrial Node. In the worst cases, this cascade of events lead to a totally disorganised and potentially lethal arrhythmia: ventricular fibrillation, which is characterised by a chaotic and disorganised electrical activity during which the heart 'shivers' and does not contract. Untreated VF leads to sudden cardiac death, a major cause of mortality in the world.

1.1.2.1 Ventricular Tachycardia Therapy

Different therapeutic approaches are possible for ischaemia-related VT.

On one hand, a possible treatment of symptomatic VT relies on the implantation of an implantable cardioverter defibrillator (ICD), a small device placed under the chest skin that detects arrhythmia and delivers electric shocks (cardioversion) to the patient when it senses an abnormally high heart rate.

On the other hand, the etiologic treatment for VT is the catheter-based ablation of foci through the destruction of zones responsible for the re-entry waves. Usually, the VT ablation is performed by employing radio frequency energy. While radiofrequency thermal ablation (RFA) has already demonstrated great benefits [Ghanbari, 2014], the




Fig. 1.5: Infarct : a) Overview of the coronary system; b) Cross-section of the coronary artery with details on the occlusion process. In this thesis, we use the term "infarct scar" instead of "dead heart muscle" to refer to the necrotic tissue deprived from oxygen. Illustration from [https: National Heart, Lung and Blood Institute](https://www.nationalheart.org/), public domain.

technique still lacks clinical consensus regarding the optimal treatment strategy [Aliot, 2009] mainly due to the inadequate detection of the VT substrate. Moreover, this invasive procedure is extremely time-consuming (6-10 hours), containing two long phases: i) the electrical mapping of the endocardial or epicardial potentials; and, ii) the ablation therapy delivery followed by scar re-mapping and VT re-inducibility. Therefore, this procedure is often incomplete or not applicable especially to patients presenting with poor hemodynamic conditions.

1.1.3 Imaging the Heart Electrical Activity

1.1.3.1 Electro-anatomical in vivo mapping

Electro-anatomical mapping (EAM) is a catheter-based invasive technique used to record in-vivo the cardiac electrical activity at specific locations inside the heart. The method of catheter insertion is chosen according to the cardiac chamber to be studied: through the femoral vein for the right side of the heart, or, through the femoral artery for the left side. Figure 1.6 illustrates a general procedural setup for this invasive procedure. This

procedure is initially performed in the first phase of RFA therapy, but it can be repeated after the ablation procedure in order to validate the therapy success.




Fig. 1.6: Illustration of the clinical set-up for an x ray image-guided procedure in the EP lab, using a 3-dimensional electro-anatomical mapping system. Illustration taken from [Rhythmia HDx, Boston scientific, USA](#).

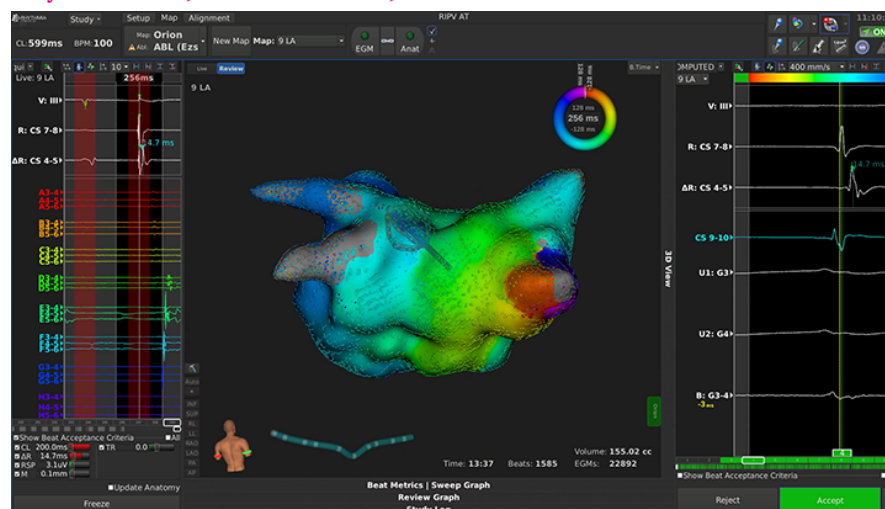


Fig. 1.7: Translation of captured electrical signals into a 3D visualisation system, with the position of catheter being clearly visible. Illustration taken from [Rhythmia HDx, Boston scientific, USA](#).

During the procedure, a catheter is navigated under image guidance along the endocardial and/or epicardial surfaces, capturing local electrical signals named intracardiac electrograms (iEGMs). When the catheter either comes in contact with the tissue of interest or is located within its vicinity, the iEGMs and the local coordinates of catheter's position in space are both recorded and projected onto a heart 3D anatomical shell model, as is illustrated in Fig. 1.7. Using this 3D model and the recorded iGMs, a clinical

electrophysiologist will try to diagnose and locate the arrhythmia foci. However, the data points are rather sparse, often leading to an inadequate identification of these foci.

1.1.3.2 Ex vivo optical mapping data

Another type of data used to study the cardiac electrical activity and to model the heart pathologies is optical mapping data. Originally developed to study brain electrophysiology, optical fluorescence imaging provides an excellent platform for fully recording the depolarization and repolarization phases of action potential with high temporal ($<1\text{ms}$) and spatial resolution ($<1\text{mm}$). This technique is an ideal basic science tool for exploring in detail the normal propagation of cardiac impulse and underlying pathological changes in electrical properties, as well as for calibrating mathematical models. However, due to its toxicity, the voltage-sensitive optical dye can only be used in ex vivo experiments (e.g., in cell culture, tissue patches, or freshly explanted hearts). The optical data give an excellent visual surfacic representation of the electrical wave propagation through the entire heart, being useful for scientific pre-clinical studies and model validation under precisely controlled conditions (e.g. heart pacing at different frequencies or locations).

In this thesis, to validate our developed models in real settings, we use an experimental database that contains optical fluorescence signals recorded ex vivo on several hearts freshly explanted from healthy swine [Pop, 2009]. For the optical recordings, the experimental protocol is described below. All hearts were excised from juvenile swine (weighing approximately 25kg), using an open-chest procedure approved by the animal research protocol guidelines at Sunnybrook Research Institute, Toronto (Canada).

After the explantation, the aorta of each heart was rapidly cannulated and attached to a Langendorff perfusion system. The Langendorff system allowed the heart to be constantly perfused with a mixture of oxygenated blood and Tyrodes' solution maintained at normal physiological conditions (i.e., $\text{pH}=7.3\pm0.3$ and 37°C temperature). A 20-ml bolus of a voltage-sensitive fluorescence dye solution (0.2ml di4-ANEPPS Biotium Inc, dissolved in saline solution) was injected into the heart via the perfusion line. Next, a bolus of an electro-mechanical uncoupler (2,3-butanediol monoxime, Sigma Aldrich) was injected to stop the heart contraction. This was specifically done in order to avoid heart motion artifacts, which can otherwise severely disrupt the pixel-wise recordings of each action potential (AP) wave. For the optical imaging of action potential, the fluorescence dye was excited with green light ($530\pm20\text{nm}$) using a 150-W halogen light source (Moritex Corp, Japan). The photons (i.e., optical signals) emitted from the epicardium were filtered through a high-pass red filter ($> 610\text{ nm}$) and then collected by a high-speed dual-CCD camera system (MICAM02, BrainVision Inc, Scimedica Japan). For all recordings a high-temporal resolution (3.7ms) as well as a high-spatial resolution (i.e., pixel size 0.7mm) were used. For the electrophysiological stimulation protocol, the

hearts were paced using a bipolar stimulating Ag/AgCl₂ electrode, with the duration of the square-wave stimulus being set to 5-10ms.

For each pixel, the AP wave was then derived from the relative change in the intensity of optical fluorescence signal. All recorded signals were exported from the BV acquisition software (<https://www.scimedia.com>). Figure 1.8 shows the experimental set-up, exemplary recordings of action potential waves from one pixel, and the associated 2D depolarization maps obtained following stimulation.




Fig. 1.8: Optical fluorescence imaging of action potential in an explanted porcine heart perfused via a Langendorff system: (a) snapshot of the optical set-up using two CCD cameras; (b) 2D optical image showing the location of the pacing electrode; (c) example of depolarization map recorded (red represents the depolarization phase, while blue represents the repolarization phase); and (d) example of normalized action potential waves recorded from one pixel, at a pacing frequency of 1.1 Hz (shown over 3.7 s). From [Camara, 2011].

1.2 Cardiac Electrophysiology Modelling: State-of-the-art

Cardiac electrophysiology (EP) (see section 1.1.1.2) can be mathematically modelled in order to reproduce the behaviour of AP wave propagation and to give a mechanistic understanding of arrhythmia genesis per individual heart/patient. Such computational modelling could also help us predict the response to a spectrum of treatments. This *in silico* approach makes the personalisation of a certain cardiac EP model, along with the estimation of the associated patient-specific parameters, a crucial point for predicting the response to therapies on an individual heart basis.

1.2.1 Physical Models

Physiological and multi-physics phenomena characterising the function of the heart in normal and pathological conditions, can be mechanistically described by mathematical models of different complexity. Notably, several EP models are able to accurately reproduce the electrical behaviour of the heart at different scales (i.e., cellular, tissue or organ).

There are three principal types of physical EP models describing the action potential (AP): biophysical models, phenomenological models and Eikonal models, depending on the degree of complexity and the scale of phenomena modelled. Cardiac EP models can be also classified as bidomain or monodomain, depending on the considered intra/extracellular spaces and the associated electrical potentials of the cell.

1.2.1.1 Biophysical Models

Complex detailed models [Ten Tusscher, 2004; Ten Tusscher, 2006] are used to describe the dynamics of transmembrane voltage, the flowing current, as well as the different ionic concentrations inside and outside the cardiac cell. However, the ionic models are not only intricate but also computationally expensive, requiring costly resources such as super-computers and clusters. In addition, these sophisticated models also employ numerous hidden variables that are impossible to be all measured, making difficult to accurately identify and calibrate all model parameters.

1.2.1.2 Phenomenological Models

An alternative to using complex computational models, is to employ phenomenological models involving descriptions derived from simple biophysical models adapted to cardiac EP. Examples of such models are the FitzHugh-Nagumo [FitzHugh, 1961; Nagumo, 1962], Aliev-Panfilov [Aliev, 1996; Nash, 2004], and Mitchell-Schaeffer [Mitchell, 2003] models, which only make use of a very few variables (two or three) and a reduced number of parameters in their mathematical equations. These models are used for rapid simulations of the AP wave propagation at tissue level (e.g. on a 2D tissue layer or a 3D tissue slab), and organ level (i.e., through the entire heart). However, being less realistic, these simplified models may be also less accurate.

1.2.1.3 Eikonal Models

Eikonal models developed by Keener [Keener, 1991] are simplistic models corresponding to non-linear partial differential equations of the activation time, typically modelling only the propagation of the depolarisation wave of the action potential. These categories of models ignore the complexity of the dynamics of voltage changes in the cardiac membrane. They make the assumption that the propagation of the APs in the myocardium can be described as a wave. This makes them a lot easier to parameterise, and extremely fast to compute, at the cost of losing the possibility to simulate some complex phenomena, although several variants have been proposed to overcome this limitation [Pernod, 2011].

1.2.2 Data-driven Approaches

With the recent advances in artificial intelligence, pure Machine Learning (ML) and Deep Learning (DL) methods have been increasingly used in order to learn dynamical models from data and make intelligent decision without human intervention, motivating several works relevant to the simplification of EP models and their representation through equations. Notably, it has been shown that neural networks are capable to learn and to model various complex (e.g. nonlinear) relationships between input and output data [Willard, 2020].

The idea of leveraging ML methods in order to learn data-driven models of dynamical systems is not quite new, for instance [Nelles, 2001] gives a thorough introduction to the closely related field of Nonlinear System Identification, while [Crutchfield, 1987] gives an early example of such endeavours. More recently, these topics have seen a renewed interest with works such as [Raissi, 2017; Alvarez, 2013; Rudy, 2017; Zhang, 2018] proposing to use alternative statistical learning tools such as Gaussian Processes and sparse linear regressions to learn the explicit form of differential equations.

Neural Networks have been also used for similar problems. For example, a framework useful for automatic learning PDEs from data [Long, 2018; Long, 2019] has been proposed. One group used Deep Neural Network models for solving differential equations [Raissi, 2018a; Sirignano, 2018]. Another group used an adjoint method to learn differential equations parameterised with neural networks [Chen, 2018], while [Ayed, 2019a] proposed a framework for learning models using a purely data-driven approach in partially observable settings.

However, despite achieving good progress and producing promising results in cardiac EP simulations [Ayed, 2019b; Kashtanova, 2021], data-driven models alone were not able to reproduce complex unseen dynamics (such as the repolarisation phase of action potential); thus, the maximum forecasting horizon remains limited. These limitations are discussed in detail in this manuscript (section 2.5).

1.2.3 Physics-Based Deep Learning Approaches

Physics-aware deep learning is a recent field of research aiming at promoting the use of data-intensive methods for the modeling of complex physical phenomena [Karniadakis, 2021; Willard, 2022; Wang, 2022b]. This research topic motivates works from different disciplines, ranging from climate to aeronautics and biology, by encompassing diverse objectives including accelerating numerical simulation, improvement over physical models, building of emulators, solving differential equations in large variable spaces, discovering physical laws from data, etc. The methods developed for reaching

these objectives are also multiple-folded: incorporating prior physical background in the loss functions [Raissi, 2019; Raissi, 2018b; Yang, 2021] or as strict constraints in the NN architecture [Jiang, 2019], enforcing the invariance and equivariance properties characteristic of the physical laws in DL models [Wang, 2021; Wang, 2022a], combining DL and physical components into hybrid systems [Yin, 2021].

Subsequently, several research groups have started to use coupled physico-statistical approaches for cardiac EP simulations, in order to achieve a high precision at low computational cost. For example, one group designed a neural network that approximates the FitzHugh-Nagumo model [Court, 2021], while others used a physics-informed neural networks to construct cardiac activation maps by accounting for the underlying wave propagation dynamics [Sahli Costabal, 2020], and to estimate the cardiac fiber architecture of the human atria from catheter recordings of the electrical activation times [Ruiz Herrera, 2022]. Furthermore, another group proposed an approach to create a nonlinear reduced order model by employing deep learning algorithms (DL-ROM) designed for cardiac EP simulations [Fresca, 2021]. Finally, other researchers presented a physics-informed neural network for an accurate simulation of action potential wave and a correct estimation of the model parameters [Herrero Martin, 2022].

However, the majority of these coupled approaches use a high-fidelity physical model as a core component of its structure. As a result, fitting those models to the real data may not only be computationally expensive, but also difficult especially in order to properly deal with the frequently observed large discrepancies between simulated and real data.

1.3 Motivation and Scientific Objectives

Computational cardiology is a multi-disciplinary field that has seen extensive progress in the past decade. In particular, recent advances in numerical analysis and the development of virtual patient-specific models (known as ‘digital twin’) have allowed researchers to address critical challenges related to limitations of clinical methods that are routinely employed to diagnose arrhythmia, as well as to help planning the best therapy on an individual heart basis. However, in order to build such accurate predictive heart models, one needs to select the most suitable theoretical framework, balancing the following aspects: the degree of mathematical complexity needed for the specific problem studied; the correct parameterisation of model from measurements; and, the validation of predictions.

As explained in detail in section 1.2.1, physical models alone can not be used as sufficient framework for patient-based modeling because accurate but complex models are difficult to personalise, while simpler and faster models are less realistic and, consequently, need

a complementary mechanism to enable their fitting to measured data. Furthermore, it is still difficult to reduce the discrepancy between the output of idealised mathematical models and clinical measurements, which are usually noisy.

Given the increasing availability of simulated and observation data, an important research issue is to consider how ML and, more specifically, DL methods could help complementing simple physical models in order to improve their accuracy.

In sections 1.2.2 and 1.2.3, we presented existing data-driven and physics-based DL approaches applied in cardiac EP modelling. Unfortunately, the pure data-driven methods encounter difficulties with respect to generalisation on new conditions. Moreover, the majority of the coupled approaches are based on high-fidelity physical models. As a result, fitting those models to real data may not only be computationally expensive, but also difficult especially in order to properly deal with the frequently observed large discrepancies between simulated and real data.

Therefore, the scientific objectives of this thesis are:

- to develop an automatic framework to learn the cardiac EP dynamics from data (personalisation);
- to validate this framework on EP data of different complexity (flexibility);
- to prove that this framework can generalise to new conditions (generalisation).

1.4 Main Contributions

Our main contributions, presented throughout this manuscript are the following:

- a fully data-driven DL framework (section 2.2) that is able to learn and forecast simple cardiac EP dynamics in the presence of unexcitable scars (modelled in a virtual cardiac tissue slab), and also to generalise to new complex and unseen conditions (section 2.4);
- a fast physics-based DL framework for automatic learning cardiac EP dynamics from data of different complexity (section 3.2);
- a method for full and fast personalisation of cardiac EP model (section 4.2.3), evaluated on real OD ex vivo data (i.e., OD optical AP waves);

- a method for deep personalisation of cardiac EP model (section 5.2.3), evaluated on real 2D ex vivo data (i.e., 2D optical AP waves including depolarisation wave propagation).

1.5 Manuscript organisation

The manuscript is organised as follows, in accordance with the above-mentioned scientific objectives:

In Chapter 2, we introduced a principled data-driven DL approach to learn the cardiac EP dynamics in the presence of unexcitable scars in the cardiac tissue slab. Using synthetic data, we explore the ability of this framework to learn and forecast assimilated dynamics, and to generalise to new conditions including more complex scar geometries, multiple signal onsets and various conduction velocities.

In Chapter 3, we presented a refinement of the previous framework with physical priors. In this work, we proposed a fast physics-based DL framework, which is able to learn cardiac EP dynamics from data of different complexity. Using synthetic data, we demonstrated that this framework allows us to reproduce the complex dynamics of the transmembrane potential, even when only incomplete (partial) measurements are available.

In Chapter 4, the physics-based DL framework of Chapter 3 is used to identify the key physical parameters for different anatomical zones of the porcine heart (using 0D ex vivo data). Alongside, we demonstrated examples of faster personalisation of cardiac EP model on 0D dynamics data using our novel framework.

Chapter 5 is an extension of Chapter 4 to 2D ex vivo data, by including a depolarisation wave propagation through the cardiac tissue in the overall learned dynamics. In particular, we demonstrated that the learning framework of Chapter 3 is able to learn and forecast the real ex vivo recorded 2D dynamic. Additionally, we show the framework's potential to automatic learning of local conduction velocity of cardiac fibers and its extrapolation to the entire organ.

In Chapter 6, the main contributions of this thesis are summarised. Finally, potential future works and perspectives are discussed.

EP-Net 2.0: Data-Driven Model for Cardiac Electrophysiology Learning

Contents

2.1	Introduction	17
2.2	Problem Formulation	19
2.2.1	Learning Method	20
2.3	Experiments	21
2.3.1	Data Collection	21
2.3.2	Training Settings	21
2.4	Results	21
2.4.1	Testing on rectangular scars	22
2.4.2	Generalisation: scars of various shapes and multiple onsets	22
2.4.3	Generalisation: various conduction velocities	23
2.5	Limitations and Discussion	25
2.6	Conclusion and Future Work	26
2.A	Modelling real scars with domain mask containing multiple tissue conductivities	27

Abstract Cardiac electrophysiology (EP) models have achieved good progress in simulating cardiac electrical activity for normal cases and pathologic scenarios. However, it still remains challenging to leverage clinical measurements due to the discrepancy between idealised models and realistic patient-specific conditions. Machine learning (ML), or more particularly Deep learning (DL) approaches could help alleviate these difficulties, due its data-orientated structure.

In this chapter, we included the full version of our publication, which was published in Lecture Notes in Computer Science Proceedings for 11th International Conference on Functional Imaging and Modeling of the Heart (FIMH 2021) [[Kashtanova, 2021](#)].

Specifically, in this publication, we proposed a principled data-driven DL approach to learn the cardiac EP dynamics in the presence of scars in a slab of cardiac tissue (section 2.2). Using simulated data, we demonstrated that this technique is indeed able to reproduce the transmembrane potential dynamics in situations close to the

training context (section 2.4.1). We then focused on evaluating the ability of the trained networks to generalise outside their training domain. In the sections 2.4.2 and 2.4.3, we showed through in silico experiments that our model is able to generalise to new conditions such as: more complex scar geometries; multiple signal onsets; and, various conduction velocities.

Our original contributions in this article are:

- the development and implementation of an improved framework (EP-Net 2.0) for automatic learning cardiac EP dynamics from data (section 2.2);
- an evaluation of the framework ability to learn and forecast simple cardiac EP dynamics (section 2.4.1);
- an evaluation of the framework ability to generalise to complex and unseen conditions (sections 2.4.2 and 2.4.3).

In the Appendix 2.A we explored further the abilities of framework to learn and generalise the transmembrane potential dynamics in presence of real-like scar tissues modelled via scar masks with heterogeneous distribution of tissue electrical conductivities.

EP-Net 2.0: Data-Driven Model for Automatic Learning of Cardiac Electrophysiology Dynamics

Kashtanova, V., Ayed, I., Cedilnik, N., Gallinari, P., and Sermesant, M.

Published in *Functional Imaging and Modelling of the Heart* [Kashtanova, 2021]

2.1 Introduction

Mathematical modelling of the heart has been an active research area for the last decades, and it is nowadays more and more coupled with artificial intelligence approaches, see for instance [Mansi, 2020]. Among the multi-physics phenomena involved in the cardiac function, cardiac EP models can accurately reproduce electrical behaviour of cardiac cells. However, it is still challenging to leverage clinical measurements due to the discrepancy between idealised models and realistic patient-specific conditions. Machine learning (ML) approaches could help alleviate these difficulties.

The idea of leveraging ML methods in order to learn data-driven models of dynamical systems is not new: [Nelles, 2001] gave a thorough introduction to the closely related field of Nonlinear System Identification while [Crutchfield, 1987] gave an early example of such endeavours. More recently, those questions have seen a renewed interest, with works such as [Raissi, 2017; Raissi, 2018b; Sirignano, 2018] proposing to use Deep Neural Network models for solving differential equations, while [Alvarez, 2013; Rudy, 2017; Zhang, 2018] used alternative statistical learning tools such as Gaussian Processes and sparse linear regressions to learn the explicit form of differential equations. In the last few years, Neural Networks have been increasingly used in order to learn dynamical models from data: [Long, 2018; Long, 2019] endow neural layers with additional structure, beneficial for learning PDEs while [Chen, 2018; Ayed, 2019a] used the adjoint method to learn differential equations parametrised with neural models and learn them in fully and partially observable settings. More generally, [Willard, 2020] propose a broader survey of ML in the field of physics-based modeling.

We propose here a framework for learning cardiac electrophysiology (EP) dynamics from data and we experimentally evaluate (via *in silico* studies) its ability to forecast cardiac dynamics on new conditions, unseen during training. Our models were trained and evaluated using data simulated from an electrophysiology model [Mitchell, 2003]. This is a classical experimental setting in the domain [Fresca, 2020] which, although offering a simplification over real cardiac data, it allows us to assess our framework using controlled conditions. This work builds on initial results [Ayed, 2019a; Ayed, 2019b] that were

evaluated in an idealised setting with simple boundary conditions corresponding to a healthy slab of cardiac tissue having a uniform conductivity and a single onset of transmembrane potential.




Fig. 2.1: Example of transmembrane potential (yellow) propagation in the cardiac tissue slab in absence (a) and the presence (b) of scar tissue, through successive time steps.

Our extension here introduced more complex conditions. First, we considered diffusion in tissues with ischaemic (non conductive) regions denoted scars in the following. In clinical practice it is essential to be able to recognise and to estimate the impact of scars, because these are the main cause of cardiac arrhythmias. For example, in Fig. 2.1 we can clearly observe the changes in the dynamics of the depolarisation wave in the presence of scar tissue (black area). Second, we introduced in our simulations multiple onsets and colliding wave fronts, as this is a classical situation in cardiac electrophysiology. The particular focus of our work is on the evaluation of the ability of our model to generalise to unseen conditions. The model was then trained on simulated data corresponding to relatively simple context (i.e., one type of scar; one wave front; and, a set of several discrete conduction velocities) and its generalisation ability was further evaluated on more challenging contexts such as: more complex scars; multiple fronts; and, any real conduction velocity sampled from a given interval. Fig. 2.2 presents the general setting of the in silico experimental used in the manuscript.

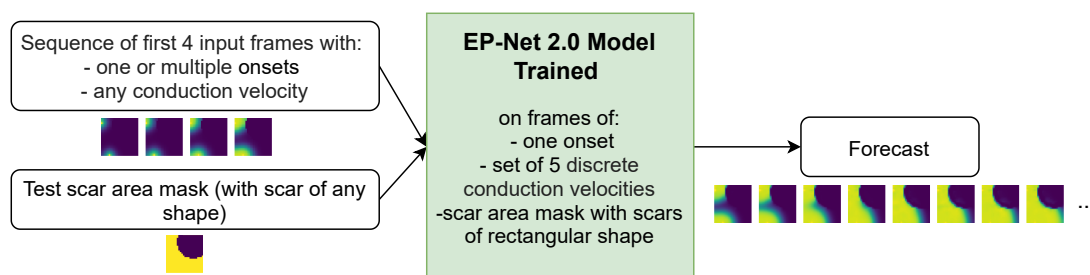


Fig. 2.2: General setting used throughout the manuscript. Once trained, the EP-Net 2.0 model takes as input a context consisting of a few (i.e., $n=4$ here) observations of transmembrane potential dynamics, plus an indication of the scar area (left), and forecasts transmembrane potential depolarisation wave propagation (right).

2.2 Problem Formulation

We used the Mitchell–Schaeffer biophysical model [Mitchell, 2003] for cardiac EP simulations. This two-variable model has been successfully used in patient-specific modelling of scar-related ventricular arrhythmia [Relan, 2011a]. The variable v in equation 2.1 represents a normalised ($v \in [0, 1]$) dimensionless transmembrane potential, while the “gating” variable h controls the repolarisation phase (i.e., the gradual return to the initial resting state):

$$\begin{aligned} \partial_t v &= \text{div}(\sigma \mathbf{I} \nabla v) + \frac{h v^2 (1 - v)}{\tau_{in}} - \frac{v}{\tau_{out}} + J_{stim} \\ \partial_t h &= \begin{cases} \frac{1-h}{\tau_{open}} & \text{if } v < v_{gate} \\ \frac{-h}{\tau_{close}} & \text{if } v > v_{gate} \end{cases} \end{aligned} \quad (2.1)$$

where J_{stim} is a transmembrane potential activation function, which is equal to 1 during the time the stimulus is applied (t_{stim}) in a certain stimulated area.

In practice, since h is a hidden variable, it is difficult to measure, and only the measurement of the potential v could be made available. Therefore, as in [Ayed, 2019b], we modify the system (2.1) by replacing variable h with an observation operator \mathcal{H} which extracts the corresponding information from the current state X_t . This clever approach allows us to rewrite this model in a vector form:

$$\begin{cases} X_0 = g_\theta(V_{-k}) \\ \frac{dX_t}{dt} = F_\theta(X_t) \\ V_t = \mathcal{H}(X_t) \end{cases} \quad (2.2)$$

where $X = (V, H)^T$ is a spatio-temporal two-dimensional vector field over the domain $\Omega \subset \mathbb{R}^2$, g_θ and F_θ are parameterised functions which allow to model the ODE governing the dynamics of X and $V_{-k} = (V_{-k+1}, \dots, V_0)$ is a sequence of the past observations of transmembrane potential.

We then introduce a constraint corresponding to the presence of the scar:

$$\Omega_{scar} \subset \Omega \subset \mathbb{R}^2 : (X_t)_{\Omega_{scar}} \equiv 0 \quad (2.3)$$

Note that in our setting, for simplification, scars are considered to be binary masks.

In order to enforce the constraints (2.2) and (2.3), we compared the sequence of observations V_t generated by the parameterised model to the data simulated from the actual equations, and minimised the following loss:

$$\mathcal{L}(V, \tilde{V}) = \mathcal{L}_{obs}(V, \tilde{V}) + \lambda_{scar} \mathcal{L}_{scar}(\tilde{V}), \quad (2.4)$$

$$\text{where: } \mathcal{L}_{obs}(V, \tilde{V}) = \int_0^T \|V_t - \tilde{V}_t\|^2 dt, \mathcal{L}_{scar}(\tilde{V}) = \|\Omega_{scar} \odot \tilde{V}_t\|^2$$

with \odot the element-wise product and λ_{scar} a hyper-parameter used to balance the losses. We can then frame the statistical learning problem in the following form:

$$\begin{aligned} & \underset{\theta}{\text{minimize}} && \mathbb{E}_{V \in \text{Dataset}} \mathcal{L}(V, \mathcal{H}(X^\theta)) \\ & \text{subject to} && \frac{dX_t}{dt} = F_\theta(X_t), \\ & && X_0 = g_\theta(V_{-k}) \end{aligned} \quad (2.5)$$

2.2.1 Learning Method

The operators F, g in problem (2.5) are implemented via Deep Neural Networks. We chose to use ResNets [He, 2016] (illustrated in Fig. 2.3) to parameterise both F and g . Optimisation is performed via a stochastic gradient descent, precisely ADAM [Kingma, 2014], according to the following algorithm:

0. Randomly initialise θ (denoting the parameters of F and g);
1. Solve the forward state equation 2.2 to find X^θ with an explicit differentiable solver (Euler scheme);
2. Get the gradient of $\theta \rightarrow \mathbb{E}_{V \in \text{Dataset}} [\mathcal{J}(V, \mathcal{X}^\theta)]$ with automatic differentiation tools and update θ ;
3. Repeat from step 1 until convergence.




Fig. 2.3: The ResNet architecture used in EP-Net 2.0. It has 5 input frames (mask frame plus 4 frames of the transmembrane potential) and 1 output frame of forecast.

2.3 Experiments

2.3.1 Data Collection

We generated 2D data frames using the Lattice Boltzmann method to solve the EP model [Rapaka, 2012b] on a Cartesian grid. The Mitchell-Schaeffer model parameters were taken as in the original paper [Mitchell, 2003]: $\tau_{in} = 0.3$, $\tau_{out} = 6$, $\tau_{open} = 120$, $\tau_{close} = 150$, $v_{gate} = 0.13$. The computational domain was represented by a slab of cardiac tissue, having a selected size of $24 \times 24 \text{ mm}^2$ discretised in small 1 mm^2 pixels. A stimulation current was applied for 10 ms to initiate the propagation of electrical impulse in selected pixels (J_{stim}). In order to simulate scars in the cardiac tissue slab, we superposed a mask with a randomly generated rectangular area (having a random size and position), excluding it from the domain. Training was performed with five different discrete conduction velocities, each corresponding to a given parameter for electrical conductivity ($\sigma = 1, 2, 3, 4, 5$) in equation 2.1. The conductivity was considered uniform on the whole cardiac slab, except for the non-conductive scar zone. One value of σ was used per each simulation. The simulations were conducted for 30 ms, with a discrete time step of 0.1 ms, and stored every ms. Then, random sequences of 10 data frames (one data sample) were extracted at different time points for training / validation data. Overall, we obtained a database of 30000 training samples and 12000 validation samples, respectively.

2.3.2 Training Settings

Parameter λ_{scar} in loss (2.4) was set to 0.1 and a learning rate for ADAM optimiser was set to 10^{-3} . We use dResNet with 64 filters at the initial stage, three downsampling initial layers and three intermediary blocks and start with a re-weighted orthogonal initialisation for its parameters. We also used exponential scheduled sampling [Bengio, 2015] with parameter 0.9999 during training. We trained our EP-Net 2.0 model until full model convergence (about 5000 epochs). In each training (and validation) sequences of data frames, we used the first 4 frames for initialisation and the rest (i.e. 6 frames) to compute the losses.

2.4 Results

Tests were performed for two selected situations: first with scar and current distributions similar to the training set, and, second with different scars and initial current onsets distributions in order to test the model ability to generalise to new situations.

2.4.1 Testing on rectangular scars

Figure 2.4 illustrates the behaviour of our trained EP-Net 2.0 model in test conditions similar to the training ones: rectangular scars with random size and position plus one onset only. Figure 2.4a shows the forecast over nine time frames (9 ms) after assimilating the first four frames (not presented in Fig. 2.4a). Notably, we observed a very good agreement with the ground truth on this forecast, which represents an important part of cardiac dynamics within this virtual slab of tissue, from the early depolarisation phase to the full depolarisation phase. Figure 2.4b shows that EP-Net 2.0 model has a very good precision on depolarisation times during more than 50 ms, an equilibrium state for the model, but cannot predict the repolarisation (Fig. 2.4c). Our quantitative results provided in the table 2.1 for different forecasting horizons T (6, 12 and 24 ms), showed excellent performance, while the training time horizon was only 6 ms.

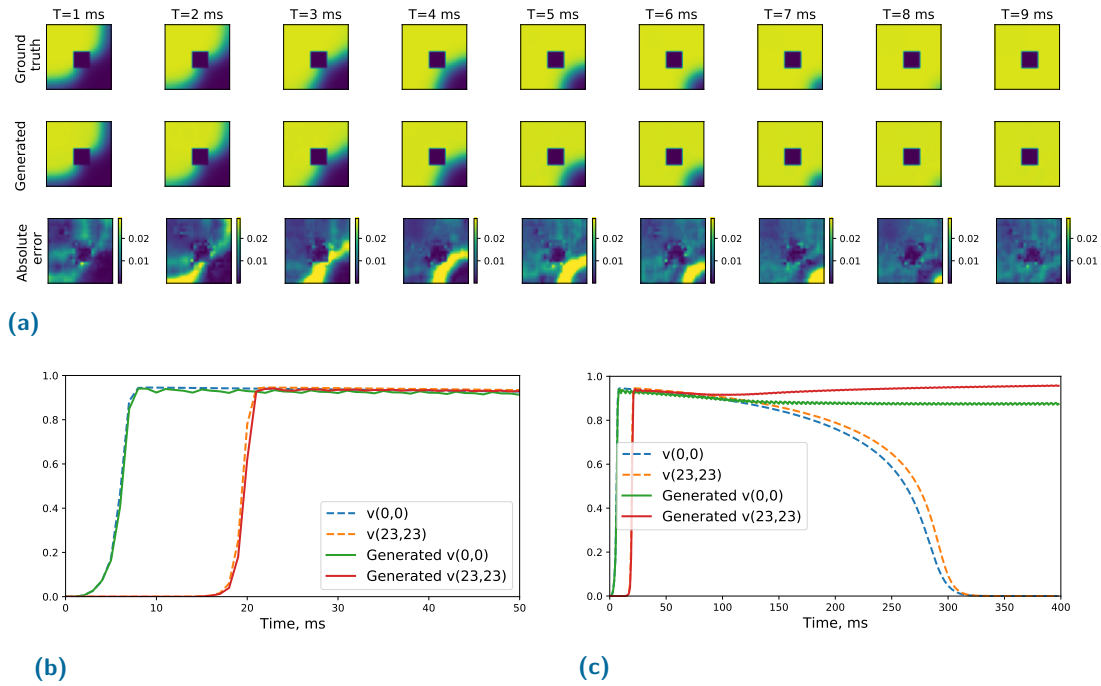


Fig. 2.4: (a) Representative results for the trained EP-Net 2.0 model (9 ms of forecast, conductivity of the cardiac slab tissue $\sigma = 2$). (b,c) Plots of the transmembrane potential at the leftmost upper point (0,0) and the rightmost bottom point (23,23) in the slab with different forecasting horizons.

2.4.2 Generalisation: scars of various shapes and multiple onsets

Our objective was to train models that are able to generalise to conditions outside the training environment. This is an important aspect since, for example, different patients will have different physiological and pathological characteristics. In order to evaluate the capability of our model to generalise we performed two types of tests, one with scars

	MSE (6 ms)	MSE (12 ms)	MSE (24 ms)
Rectangular shape	$1.8 * 10^{-4}$	$4.45 * 10^{-4}$	$6,8 * 10^{-4}$
Triangular shape	$3.1 * 10^{-4}$	$8 * 10^{-4}$	$1.36 * 10^{-3}$
Circular shape	$2.7 * 10^{-4}$	$8.2 * 10^{-4}$	$3.4 * 10^{-3}$
Complex shape	$4.6 * 10^{-4}$	$1.9 * 10^{-3}$	$6.36 * 10^{-3}$

Tab. 2.1: Relative mean-squared error (MSE) of transmembrane potential forecasting in the presence of scars of various forms, computed for different forecasting horizons (cardiac slab conductivity $\sigma = 2$).

	MSE (6 ms)	MSE (12 ms)	MSE (24 ms)
One Onset	$1.8 * 10^{-4}$	$4.45 * 10^{-4}$	$6,8 * 10^{-4}$
Multiple Onsets	$4.7 * 10^{-4}$	$5.8 * 10^{-4}$	$6.9 * 10^{-4}$

Tab. 2.2: Relative mean-squared error (MSE) of potential forecasting in presence of multiple onsets and the scar of rectangular shape for different forecasting horizons (cardiac slab conductivity $\sigma = 2$).

with different shapes when training was performed only with rectangular shapes, and one with multiple onsets when training considered only one onset.

Moreover, for the generalisation to different scar shapes, we evaluated our model with the following scar shapes: triangular, circular and more complex scars (see Fig. 2.5). Table 2.1 shows that the model performs well on these different shapes. The errors are slightly larger than for the rectangular scars used for training, but still remain small overall. However, the errors increased for long term predictions (i.e., 24 ms in this example). Figure 2.5 illustrates the behaviour of our model for typical test sequences.

We next performed tests with multiple onsets. The model showed good results for forecasting of multiple depolarisation waves on one cardiac tissue slab (Fig. 2.6), which is essential to capture correctly for ventricular tachycardia simulations. As per the results presented in the table 2.2, the relative mean-squared error (MSE) is larger for multiple onsets than for one onset (as it was used for training), but still within an acceptable range. However, this error did not increase proportionately with time of forecast (as shown in the table 2.1), because the virtual slab reached faster the full depolarisation state with multiple onsets (an equilibrium state for EP-Net 2.0 model).

2.4.3 Generalisation: various conduction velocities

To estimate the ability of EP-Net 2.0 model to learn the conduction velocity of cardiac tissue we performed tests with various cardiac slab conductivity values (σ). The tests




Fig. 2.5: Results of trained EP-Net 2.0 model on scar with circular (top three rows) and complex (bottom three rows) shape (9 ms of forecast, cardiac slab conductivity $\sigma = 2$).




Fig. 2.6: Results of trained EP-Net 2.0 model with two stimulation currents applied on different pixels and at different times (9 ms of forecast, cardiac slab conductivity $\sigma = 2$).




Fig. 2.7: Results of EP-Net 2.0 model using a scar of circular shape and cardiac slab conductivity $\sigma = 3.8$ (top three rows), and a scar of triangular shape, two onsets and cardiac slab conductivity $\sigma = 1.5$ (bottom three rows).

	MSE (6 ms)	MSE (12 ms)	MSE (24 ms)
$\sigma = 0.7$	$4.65 * 10^{-4}$	$3.95 * 10^{-3}$	$1,9 * 10^{-2}$
$\sigma = 2$	$1.8 * 10^{-4}$	$4.45 * 10^{-4}$	$6,8 * 10^{-4}$
$\sigma = 2.5$	$3.5 * 10^{-4}$	$1.4 * 10^{-3}$	$1,6 * 10^{-4}$
$\sigma = 6$	$2 * 10^{-3}$	$4.7 * 10^{-3}$	$3 * 10^{-3}$

Tab. 2.3: Relative mean-squared error (MSE) of potential forecasting in the presence of various conduction velocities of cardiac slab and scar of rectangular form for different forecasting horizons.

have been performed with sigma values used for training ($\sigma \in 1, 2, 3, 4, 5$), and sigma values uniformly sampled from the interval $[0.7, 6]$. (i.e., outside the training set).

As shown in figure 2.7, the EP-Net 2.0 model preserves its capability to generalise to unseen conditions, such as scars of various shapes and multiple onsets. Our quantitative results provided in table 2.3 show that the model achieved good precision in forecasting depolarisation waves in cardiac tissue slabs for any conductivity (real number) chosen within the interval $[0.7, 6]$.

2.5 Limitations and Discussion

Section 2.4 shows the ability of model to learn the local dynamics and to generalise to unseen conditions.

Although our approach can achieve compelling results in many cases, there are still a few limitations. For example, as shown in the figure 2.8, our EP-Net 2.0 model does not perform well on thin scars (i.e., wall thickness less than 2 pixels) and produces an additional transmembrane potential diffusion through the scar from the generated noise. In addition, we acknowledge that the current model has been trained only to model the depolarisation phase of the cardiac slab tissue and cannot predict tissue repolarisation (see Fig. 2.4c), in this state. This task will be addressed in future work.

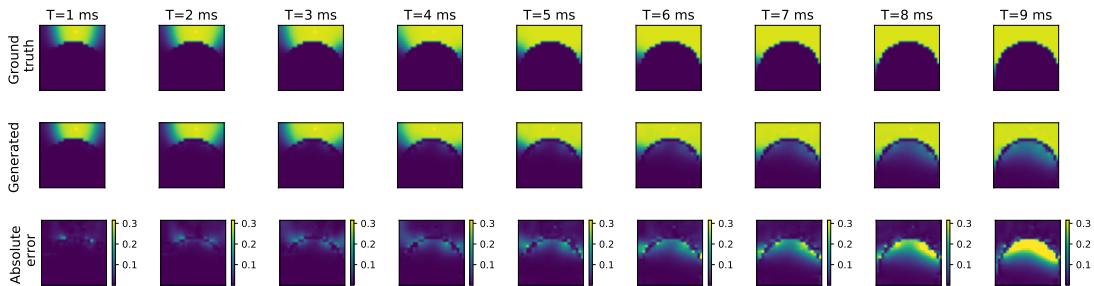


Fig. 2.8: Results of trained EP-Net 2.0 model on thin scar with circular shape.

2.6 Conclusion and Future Work

In this chapter, we proposed an EP-Net 2.0 Deep Learning model to learn the cardiac EP dynamics in the presence of complex initial boundary conditions (e.g. scar zone of a select shape, multiple onsets and various conduction velocities). Our novel results demonstrated a great generalisation ability of this model to unseen conditions. Despite training the model on data with scars with only rectangular shape and one onset activation, EP-Net 2.0 model performed well on data with scars of any chosen geometric shape and multiple onsets, even when the normalised transmembrane potential stimulations were applied at different time points¹.

We envision that our novel approach could help upgrade and personalise such mathematical models via additional data. However, in clinical practice tissue properties are not binary, and real data is always noisy and sparse. Thus, our future work will aim to include more complex formulations with respect to model parameter identification for continuously varying cardiac tissue properties, and to study possible strategies to complete the real data via in silico simulation data.

¹Visit our [github page with trained EP-Net-2.0](#), for more detailed examples.

Appendix

2.A Modelling real scars with domain mask containing multiple tissue conductivities

Ep-Net 2.0 framework has shown a great potential for modeling real scars, which typically have a heterogeneous distribution of tissue electrical conductivities (within and around the scars), which depend on myocardial tissue thickness [Cedilnik, 2018].

The training for the experiments was conducted following the same training protocol, as described in section 2.3.2, on data generated as described in section 2.3.1, except for conductivity parameter distribution. In order to generate the train masks, we randomly chose two discrete values of σ (from 1 to 5) and a position of its vertical separation plane. Figure 2.A.1 demonstrates examples of such masks used for training data generation, where the randomly generated scar area is also presented.




Fig. 2.A.1: Examples of masks containing 3 values of tissue conductivity parameter (0-conductivity, dark blue area, to model a scar tissue) used for training.

One can observe in Figure 2.A.2 that the EP-Net 2.0 framework has a very good precision for depolarisation times, even when it had been trained on data with simpler distribution. Moreover, the error is mostly produced by a continued depolarisation wave propagation within the scar area in the test data.

This demonstrates that the Ep-Net 2.0 framework can be a very efficient tool for the automatic modelling of depolarization wave propagation, even in the presence of unexcitable scars in cardiac tissue.




Fig. 2.A.2: Test results of EP-Net 2.0 trained on masks with multiple tissue conductivities. Forecast of 21 ms, first frame is an input mask.

APHYN-EP: Physics-Based Deep Learning for Model Error Correction

Contents

3.1	Introduction	30
3.2	Learning Framework	32
3.3	Experimental Settings	35
3.3.1	In silico datasets creation	35
3.3.2	Training settings	37
3.4	Results	38
3.4.1	0D synthetic data	38
3.4.2	2D synthetic data	40
3.5	Conclusion	42
3.A	Learning and Forecasting from Small Time Sequences	43
3.A.1	Results	43

Abstract Biophysically detailed mathematical modeling of cardiac electrophysiology is often computationally demanding, for example, when solving problems associated with complex patient pathological conditions. Furthermore, it is still difficult to reduce the discrepancy between the output of idealised mathematical models and clinical measurements, as the latter are inherently noisy.

In chapter 2, we presented a DL framework that is able to learn and simulate the cardiac EP dynamics, even in the presence of complex initial boundary conditions. However, as we demonstrated in section 2.5, a fully data-driven model alone was not able to reproduce the complex unseen dynamics (such as the re-polarisation phase), and therefore the maximum forecasting horizon is still limited.

In this chapter, we proposed a fast physics-based deep learning framework to learn cardiac EP dynamics from data. This novel framework has two components, decomposing the dynamics into: a physical term and a data-driven term, respectively. This construction allows our framework to learn from data of different complexity. Using the synthetic data, we demonstrated that this framework can learn the full cardiac action potential cycle (i.e., depolarisation and repolarisation phases) and reproduce

the complex dynamics of the transmembrane potential, even in the case when only partial measurements are available.

The main advantage of our proposed framework is its coupled architecture, which allows us to use a simplified low-fidelity EP model as a physical component of the framework. Such advanced framework opens up exciting possibilities in order to introduce prior knowledge in DL approaches through explicit equations, as well as to correct physical model errors from data.

The results of this chapter were partially published in Proceedings of Machine Learning Research for the 5th International Conference on Medical Imaging with Deep Learning (MIDL) [Kashtanova, 2022a] (section 3.A) and in Lecture Notes in Computer Science Proceedings for 13th International Conference on Statistical Atlases and Computational Models of the Heart (STACOM) [Kashtanova, 2022b] (section 3.4.1).

Our main contributions in this chapter are the following:

- the development and implementation of a new physics-based framework (APHYN-EP) for automatic learning of cardiac EP dynamics from data (section 3.2);
- the evaluation of the framework ability to learn and forecast the cardiac EP dynamics from synthetic data of different complexities (section 3.4);
- the evaluation of the framework capacity to generalise to new conditions and to forecast cardiac EP dynamics outside of its training domain (section 3.4.2.1).

3.1 Introduction

Mathematical modelling of cardiac physiology, mechanics or haemodynamics, has been an active research area for decades, and nowadays it is increasingly coupled with artificial intelligence approaches [Mansi, 2020; Giffard-Roisin, 2017a; Karoui, 2021; Trayanova, 2021]. Among the multi-physics phenomena involved in cardiac function, several cardiac electrophysiology (EP) models can accurately reproduce electrical behaviour of cardiac cells; nevertheless, the accuracy often depends on the model complexity and a correct personalisation of parameters.

In order to describe the dynamics of transmembrane voltage, current, and different ionic concentrations in the cardiac cell, biophysically detailed models such as the Ten Tusscher-Panfilov model [Ten Tusscher, 2004; Ten Tusscher, 2006] have been proposed. However, these models are complex and computationally expensive, and have many hidden variables which are impossible to measure, making model parameters difficult to personalise. Another type of models are the phenomenological models, which are simplified models derived from biophysical models. Examples include the FitzHugh-Nagumo, Aliev-Panfilov, and Mitchell-Schaeffer models [FitzHugh, 1961; Nagumo, 1962; Aliev, 1996; Nash, 2004; Mitchell, 2003]. These models employ fewer variables and parameters, and are therefore especially useful for rapid computational modelling of wave propagation at the organ level. However, they are also less realistic and therefore need a complementary mechanism to fit their parameters to real/measured data. Machine learning (ML) and in particular deep learning (DL) approaches could help us implement such correction mechanisms. The combination of rapid phenomenological models and ML components could facilitate the development of rapid and accurate models of transmembrane dynamics.

For this reason, researchers have begun to use coupled physico-statistical approaches for cardiac electrophysiology simulations, with a high precision and at low cost (see section 1.2.3). However, the majority of these coupled approaches is based on high-fidelity physical models and fitting them to real data. This methodology could be computationally expensive and might not be able to manage large discrepancies between simulated and real data.

To mitigate this limitation, here we propose a framework to Augment incomplete PHYsical models with a deep learning component for ideNtifying cardiac ElectroPhysiology dynamics (APHYN-EP) from data, based on a fast low-fidelity (or incomplete) physical model. This framework has two components, decomposing the dynamics into: a physical term and a data-driven term, respectively. The data-driven deep learning component is specifically designed to capture only the information that cannot be modeled by the incomplete physical model. The proposed model closely follows the approach of Yin et al. [Yin, 2021]. However, in contrast to this work which considers fully-observable dynamics and simple tests on use cases, cardiac EP dynamics have a high complexity and represent simultaneously multiple underlying processes. Furthermore, most cardiac EP models lack measurements for several variables, making them partially-observable and requiring the inferring of dynamics from incomplete observations only. Fig. 3.1.1 presents the general framework of our approach. The training amounts to identifying the physical model parameter (inverse problem) and learning the neural network parameters (direct problem) together. After training, the model can be used for forecasting at multiple horizons.




Fig. 3.1.1: General APHYN-EP framework scheme. During the training phase two-component framework learn the parameters for the physical (F_p) and the data-driven (F_d) components from data. Then, via an ODE solver, the framework can further forecast the learned dynamics.

3.2 Learning Framework

In order to learn the cardiac electrophysiology dynamics (X_t), in this work we solved an optimisation problem via our physics-based data-driven APHYN-EP framework. This particular framework combines a physical model (F_p) representing an incomplete description of the underlying phenomenon with a neural network (F_d), where the latter complements the physical model by capturing the information that cannot be modeled by the physics-described component:

$$\min_{F_p \in \mathcal{F}_p, F_d \in \mathcal{F}_d} \|F_d\| \quad \text{subject to } \forall X \in \mathcal{D}, \forall t, \frac{dX_t}{dt} = F(X_t) = (F_p + F_d)(X_t). \quad (3.1)$$

Assuming that \mathcal{F}_p is a Chebyshev set, Propositions 1 and 2 from Yin et al. [Yin, 2021] guarantee the existence and uniqueness of minimising pair for (3.1).

Specifically, our incomplete physical model is the two-variable (v, h) model [Mitchell, 2003] for cardiac EP simulations, as described by equations (2.1) (presented in the section 2.2).

This physical model has been successfully used in patient-specific modelling [Relan, 2011a], covering general EP dynamics. Furthermore, in contrast to the very detailed ionic/cellular models, this model is flexible in terms of spatial and temporal steps that are set for the numerical analysis. Thus, assuming the initial conditions for this system (2.1) $v(t=0) = 0$ and $h(t=0) = 1$ we can compute an approximation of h for any time point t by employing a simple integration scheme.

In the experiments presented later in this chapter (see Section 3.3), \mathcal{F}_p is the set of models spanned by the R.H.S. of the equations 3.1 for varying variables $\sigma, \tau_{in}, \tau_{out}, \tau_{close}$. This is a finite dimensional vector subspace which is indeed Chebyshev, thus falling under

the assumption of guaranteeing theoretical existence and uniqueness of a minimizing pair.

The data-driven component (F_d) of the framework was implemented via a neural network. The choice of a neural network depends on the application problem and the dimension of the data. In this work, we used a ResNet network [He, 2016], because it could accurately reproduce complex cardiac EP dynamics [Ayed, 2019b; Kashtanova, 2021]. However, a simpler neural network could also be used for more rapid computations.

In APHYN-EP framework the physical (F_p) and the data-driven (F_d) components were trained simultaneously. This insured the finding of the best minimising pair for (3.1) determined by set of parameters $\theta = (\theta_p, \theta_d)$. The 'Loss function' (\mathcal{L}) in training phase consisted of 2 parts: trajectory-based loss (\mathcal{L}_{traj}) and loss on norm of F_d , being represented as following:

$$\mathcal{L}(\theta) = \lambda * \mathcal{L}_{traj}(\theta) + \|F_d^{\theta_d}\| = \lambda * \sum_{i=1}^N \sum_{h=1}^{T/\Delta t} \|X_{h\Delta t}^{(i)} - \tilde{X}_{h\Delta t}^{(i)}(\theta)\| + \|F_d^{\theta_d}\| \quad (3.2)$$

where each state

$$\tilde{X}_{h\Delta t}^{(i)}(\theta) = \int_{X_0^{(i)}}^{X_0^{(i)} + h\Delta t} (F_p^{\theta_p} + F_d^{\theta_d})(X_s) dX_s$$

was calculated from the initial state $X_0^{(i)}$ via a differentiable ODE solver [Chen, 2018; Chen, 2021].

The key role of the λ coefficient was to balance the two parts of the loss. During training, we used λ in a dynamic state, as:

$$\lambda_{j+1} = \lambda_j + \gamma \mathcal{L}_{traj}(\theta_{j+1}) \quad (j \text{ is an epoch number}),$$

in order to artificially increase the importance of \mathcal{L}_{traj} at the beginning of training and then to gradually decrease it, changing the focus of optimisation on the norm of F_d . This training algorithm, adapted from Yin *et al.* [Yin, 2021], is presented below.

Algorithm 1 APHYN-EP training

```

Initialization:  $\theta_0, \lambda_0 \geq 0, \gamma > 0$ ;
for epoch = 1 :  $N_{epochs}$  do
  for batch in 1 :  $B$  do
     $\mathcal{L}_{traj}(\theta_j) = \sum_{i=1}^N \sum_{h=1}^{T/\Delta t} \|X_{h\Delta t}^{(i)} - \tilde{X}_{h\Delta t}^{(i)}(\theta_j)\|$ 
     $\theta_{j+1} = \theta_j - \nabla \left[ \lambda_j \mathcal{L}_{traj}(\theta_j) + \|F_d^{\theta_{d_j}}\| \right]$ 
  end for
   $\lambda_{j+1} = \lambda_j + \gamma \mathcal{L}_{traj}(\theta_{j+1})$ 
end for

```

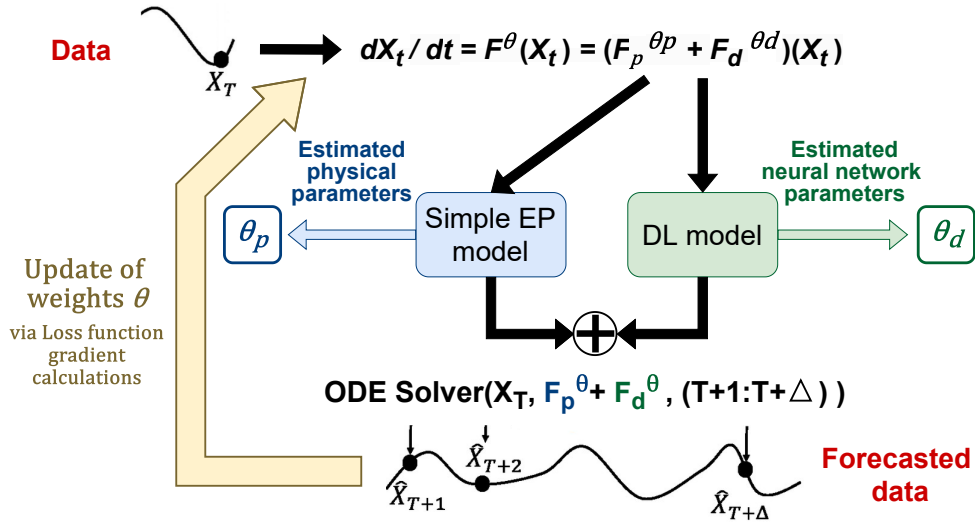


Fig. 3.2.1: General APHYN-EP framework training scheme.

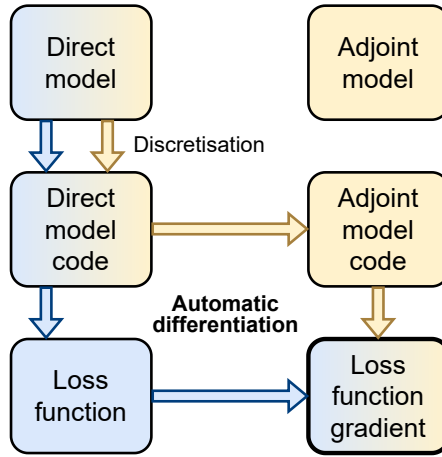


Fig. 3.2.2: Calculation of the Loss function gradient via Pytorch framework (blue part) and Adjoint model (yellow part).

Figure 3.2 is a schematic representation of the APHYN-EP training algorithm 1. This algorithm is similar to the traditional numerical method of Gradient Descent for solving Inverse and Ill-posed Problems, which is based on the solution of an adjoint problem to calculate the gradient of the Loss function [Kabanikhin, 2011] (Fig. 3.2.2(yellow part)). However, thanks to the automatic differentiation tools provided by the Pytorch library [Paszke, 2019], the gradient of (3.2) is calculated automatically inside of PyTorch framework (Fig. 3.2.2(blue part)).

Additionally, in order to train simultaneously the physical and the data-driven components of APHYN-EP, we implemented the Laplace operator in (2.1) with a simple finite-difference scheme. Lastly, in order to avoid potential difficulties associated with the high time resolution required in this numerical scheme, we used two different time

steps in the integration schemes for the physical component computations and for the computations of the final forecast given by the framework, respectively.

3.3 Experimental Settings

3.3.1 In silico datasets creation

To evaluate our framework we used synthetic datasets of transmembrane potential activations, which were simulated by two detailed ionic models [Ten Tusscher, 2004; Ten Tusscher, 2006].

First dataset was simulated with a monodomain reaction-diffusion equation and the Ten Tusscher – Noble – Noble – Panfilov ionic model [Ten Tusscher, 2004], which represents 12 different transmembrane ionic currents. The simulations were performed with a recent version of the propag-5 software [Krause, 2012; Potse, 2018], using a spatial step of 0.2 mm and a time step of 1 ms, similar with those used by Ten Tusscher et al. [Ten Tusscher, 2004]. This dataset was used for experiments on 2D data.

Second dataset was simulated by employing a monodomain reaction-diffusion equation and the Ten Tusscher – Panfilov ionic model [Ten Tusscher, 2006], which improved the detail of the previous model Ten Tusscher et al. [Ten Tusscher, 2004]. This dataset was used for experiments on OD data, to show in detail the capabilities and versatility of our physics-based framework. The simulations were performed using the same spatial and time steps, using the open-source Finetwave software¹ described in detail in the second chapter of [Nezlobinsky, 2021].

Both models used for dataset generation, included many hidden variables which are impossible to measure, making the models nearly impossible to personalise to a new (patient) condition. That was another reason for which we chose to use the simple Mitchell-Schaeffer model [Mitchell, 2003] as the incomplete physical model for our framework. Figure 3.3.1 graphically illustrates the difference between the three presented EP models.

For the generation of both synthetic datasets, the computational domain was chosen to represent a 2D slab of cardiac tissue (isotropic), with 24×24 elements in size. For one data sample, in order to activate the transmembrane potential, an excitation pulse delivered via a stimulus was applied for 1 ms on a selected area. Each simulation represented 300 ms of a heart beat (350 ms for the Ten Tusscher – Panfilov model, due to limitations in the Finetwave software), and was intended to achieve a full depolarization-

¹<https://github.com/TiNezlobinsky/Finetwave>




Fig. 3.3.1: OD signal example simulated via 3 EP models. Legend: Ten Tusscher – Panfilov model [Ten Tusscher, 2006] (TTP), Ten Tusscher – Noble – Noble – Panfilov ionic model [Ten Tusscher, 2004] (TTNNP), Mitchell-Schaeffer model [Mitchell, 2003] (M-Sch).

repolarization cycle. This required approximately 40 seconds of computation on a 12-core Intel Xeon W-2133 CPU, per simulation.

3.3.1.1 OD data

In order to generate OD databases, we selected the data samples with the pacing stimulus applied as follows: on the left top corner for training datasets, and near the center for validation datasets, as shown in Figure 3.3.2. Next, we saved a time sequence for each pixel of the cardiac slab in separate files, creating two databases: one for training and one for validation, respectively. To simplify the workflow for our framework, we removed the section of time sequences where the transmembrane potential was equal to zero, thus obtaining a time sequence of 350 ms per data sample, as shown in Figure 3.3.2(c). To test the ability of the framework to operate with noisy data, we added to each time sequence a 5 percent random noise characterized by a normal Gaussian distribution.

The data simulated via the Ten Tusscher – Panfilov model (with added noise) were considered here as the ground truth. The objective was then to learn the complex dynamics generated via this model using the APHYN-EP framework, by combining a simplified physics description with a deep learning component. We hypothesised that this approach will result in a low computational cost surrogate model of the computationally intensive, biophysically detailed Ten Tusscher – Panfilov model.




Fig. 3.3.2: (a,b) Example of a selected 2D myocardial tissue slab with the transmembrane potential activation (yellow) and resting phase (dark blue), for train and validation dataset, respectively. (c) Typical temporal sequence for the simulation experiment (without noise), with the normalised amplitude of the transmembrane potential being represented on the Y-axis and the time (ms) on the X-axis.

3.3.1.2 2D data

In order to generate 2D databases, we applied an excitation pulse at each point of 2D cardiac tissue slab. Notably, since the grid was symmetric (under 90-degree rotations), the stimulation was done only on each grid point of the first quarter of the cardiac slab. We then applied data-augmentation techniques to translate the simulation of activation wave onto the three remaining quarters. In the pursuit of conducting experiments with higher complexity, we added a 5 percent random noise with normal Gaussian distribution, directly on simulated data. Finally, we obtained a 2D database of around 500 training samples and 100 validation samples. The data simulated via the Ten Tusscher – Noble – Noble – Panfilov model (without/with added noise) were considered here as the ground truth.

3.3.2 Training settings

The physical model (F_p) described by Eq. 2.1 was implemented with a standard finite-difference scheme for the Laplace operator, using a spatial resolution of 1 mm^2 pixels (absent for 0D data experiments) and an inner time resolution of 0.1 ms. We estimated only σ , τ_{in} , τ_{out} and τ_{close} as unknown parameters in (2.1), since these control the major part of the model's dynamics (i.e., the main difference between the Mitchell–Schaeffer and the Ten Tusscher – Panfilov, the Ten Tusscher – Noble – Noble – Panfilov models in our simulations). The other Mitchell-Schaeffer model parameters were taken from the original paper [Mitchell, 2003], as follows: $\tau_{open} = 120$, $v_{gate} = 0.13$ and $t_{stim} = 1$. For the deep learning component (F_d) of the framework, we used: a ResNet [He, 2016] with 4 input/output channels (assimilating the first 4 ms of dynamics); 8 filters at the initial stage; and, 3 intermediary blocks, starting with a re-weighted orthogonal initialisation for its parameters.

We used a time resolution of 1 ms to compute the forecast given by APHYN-EP framework. The training was performed using a horizon of 350 ms. We trained the framework until a full model convergence (about 100 epochs for 0D data, 200 epochs for 2D data) was reached using an ADAM optimiser [Kingma, 2014] with initial learning rate of 10^{-3} .

It is important to emphasise that while the framework took a few hours to train (e.g. approximately 3h on Nvidia Quadro M2200 GPU in case of 2D data), once this was done, the inference step was rapidly computed (i.e., less than 10 sec to compute 350 ms of 2D forecasting) and did not require any re-calibration.

3.3.2.1 0D data

We excluded the convolution layers from the ResNet to be able to process the 0D data. The hyper-parameters λ_0 and γ of the algorithm were set to 1 and 10, respectively.

3.3.2.2 2D data

In the case the training horizon was larger than 6 ms, the training was performed with horizon augmentation technique, where we started with 6 time frames (equivalent to 6 ms horizon) for first 10 epochs and increased it gradually (by 6 ms) each 10 epochs, leading to more stable results.

The hyper-parameters λ_0 and γ of the algorithm were set to 1 and 1000, respectively.

The code and data used for APHYN-EP training are freely available at [the official github project page](#).

3.4 Results

3.4.1 0D synthetic data

The results demonstrated that our proposed APHYN-EP framework was able to accurately reproduce several key features, the wave morphology and the electrical conduction properties of the transmembrane potential solution generated by the Ten Tusscher – Panfilov model (see Fig. 3.4.1), even in the presence of noise in data.

Figure 3.4.1(a) shows that in the absence of learning of parameter τ_{close} (controlling the repolarisation), the data-driven component (ResNet model) completed the dynamics generated by physical component. The value of τ_{close} was fixed at 150 [Mitchell, 2003].

In the presence of learning of τ_{close} , which controls the action potential duration, the error of dynamics corrected by data-driven component is minimal (Fig. 3.4.1(b)).

Note that the predicted dynamic was generated via an Euler integration scheme, by assimilating only one first measurement of the transmembrane potential dynamics. Overall, the framework demonstrated robustness and was not sensitive to the noise in data, and, as a result, rapid changes in transmembrane potential activation were neglected, as observed during the first 40 ms of the action potential duration presented in Figure 3.4.1.

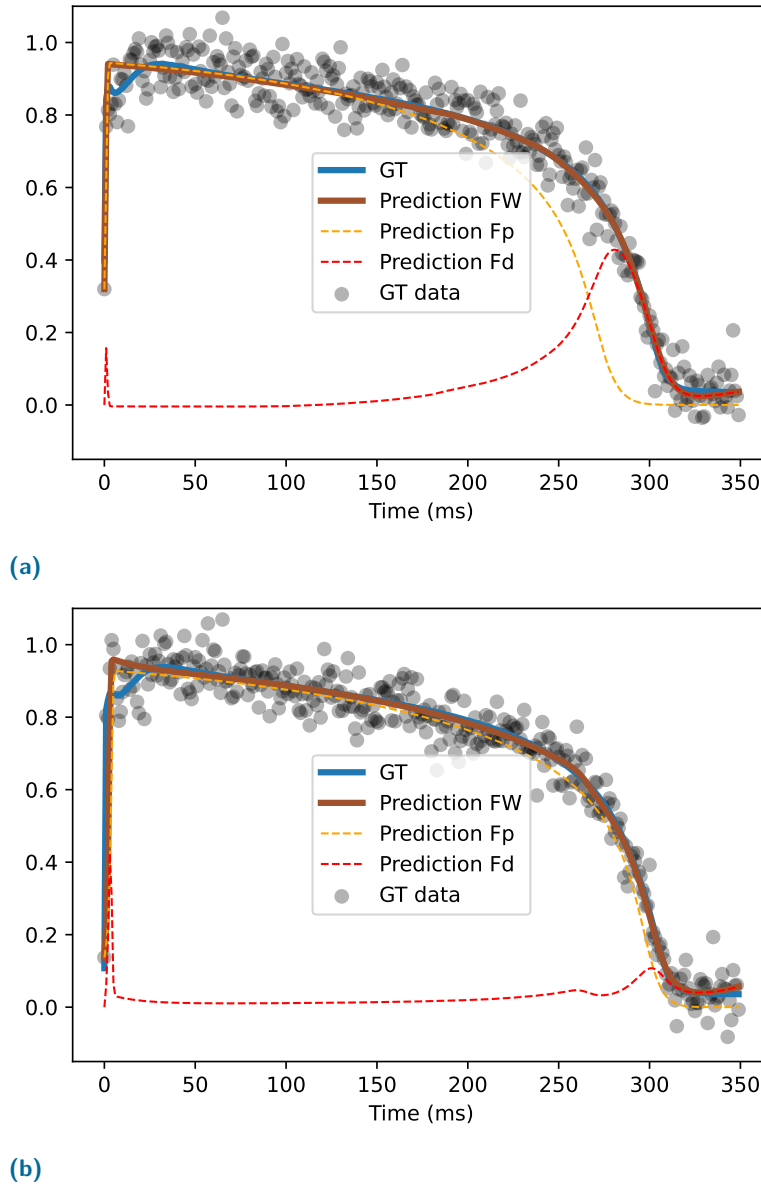


Fig. 3.4.1: Validation results of the trained framework with learning of: (a) 2 (τ_{in} and τ_{out}) and (b) 3 (τ_{in} , τ_{out} and τ_{close}) physical parameters. Legend: ground truth (GT), prediction of the framework (Prediction FW), decomposition of prediction on physical (F_p) and DL (F_d) parts.

3.4.2 2D synthetic data

In this section we demonstrate the capacity of the APHYN-EP framework to learn the 2D data of a full action potential cycle in presence of noise and we explored its generalisation ability on out-of-domain data. We compared our framework performance with two baseline models: the “incomplete” physical model (F_p , trained alone) and a fully data-driven model (EP-Net 2.0 [Kashtanova, 2021], presented in Chapter 2) trained on the same dataset as APHYN-EP.

Figure 3.4.2 presents qualitative results obtained for the forecast over 8 ms, after assimilating only one first frame of dynamics. These first 8 ms (i.e., the action potential upstroke) represent an important part of the cardiac dynamics, ranging from the earliest depolarisation phase to the full depolarisation phase. We can observe a very good agreement between the ground truth and the forecast transmembrane potentials generated by APHYN-EP in Figure 3.4.3. The effect of the correction term introduced by F_d is clearly visible.

Moreover, in Figure 3.4.3(b), one can also observe that for the 2D data (which includes now a diffusion part) APHYN-EP framework achieves a good precision in transmembrane potential forecasting even when noise is present in the data.




Fig. 3.4.2: APHYN-EP predicted dynamics for the transmembrane potential diffusion. Figure shows a period of 8 ms of the forecast. Red point is the reference point for Figure 3.4.3.

Figures 3.4.3 and 3.4.4 present the performance of different components of APHYN-EP and associated contribution to the final result. One can observe which part of the generated transmembrane potential was created by the physical component of the framework (as seen in the second row of Fig. 3.4.4). The data-driven component was used only to correct the difference between the ground-truth dynamics and the physical part (Fig. 3.4.4 (third and fourth rows)).

Table 3.4.1 presents the MSE results for our framework on the training and validation data samples. Note that to calculate this error, for each data sample, we fed the model with only one initial test measurement, then allowed the model to predict 300 ms forward without any additional input of information. For comparison, we added two baseline

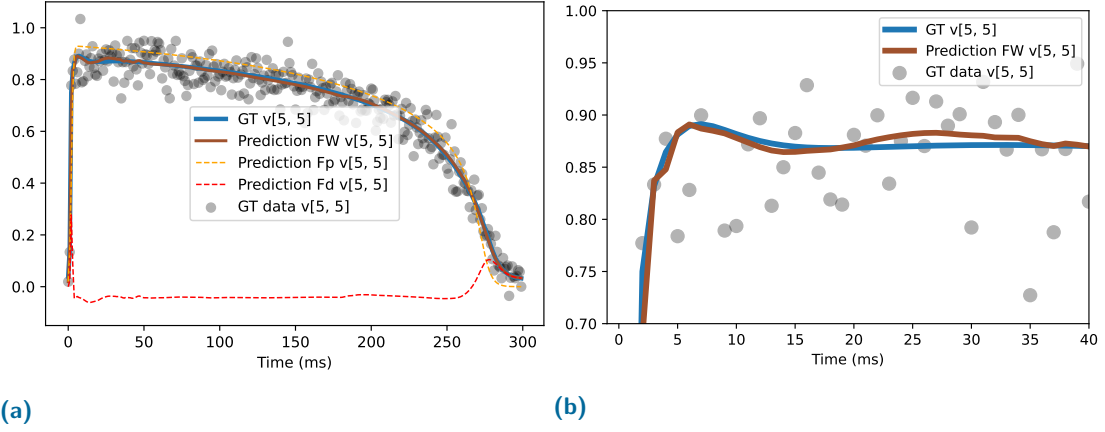


Fig. 3.4.3: Transmembrane potential at point (5,5) in the cardiac slab (red point, see Fig. 3.4.2): (a) Original, (b) Zoom on first 40 ms. Legend: ground truth (GT), APHYN-EP, physical (F_p) and data-driven (F_d) component of APHYN-EP.




Fig. 3.4.4: Exemplary results illustrating: the dynamics of the transmembrane potential diffusion predicted by the APHYN-EP physical component (second row); the error with ground-truth diffusion for this physical component of APHYN-EP (third row); and, the trained APHYN-EP data-driven component contribution (bottom row).

models: the “incomplete” physical model (F_p from APHYN-EP framework, trained alone) and a fully data-driven model (EP-Net 2.0 [Kashtanova, 2021]) trained on the same dataset as APHYN-EP described in 3.3.1 (2D). We clearly noticed that APHYN-EP captured the observed dynamics with good precision for a large time horizon (300 ms) and also outperformed the physical model for every dataset. At the same time, the pure data-driven model encountered difficulties to learn the proper dynamics.

3.4.2.1 Generalisation ability of APHYN-EP: Planar wave

Since our objective was to train a model able to generalise to new conditions, outside of the training environment, we performed a test on out-of-domain data represented by planar wave dynamics (see Fig. 3.4.5) simulated by Ten Tusscher – Panfilov model via the Finetwave software. One can observe that APHYN-EP (trained on data from 3.3.1, 2D) has successfully generated the forecast of the new transmembrane potential wave dynamics.

Method	Training data	Validation data	Out-of-domain test
APHYN-EP framework ($\ F_d\ ^2$)	2.54 (0.47)	2.54 (0.472)	4.2 (0.4)
Physical model	5.7	5.6	4.3
Data-driven model	10	10	100

Tab. 3.4.1: Mean-squared error, MSE ($\times 10^{-3}$) of the normalised transmembrane potential (adimensional) forecasting (forecasting horizon of 300 ms). Baseline models: the Physical model (2.1) and a fully data-driven model (EP-Net 2.0 [Kashanova, 2021]) trained on the same dataset as APHYN-EP. Out-of-domain tests: Planar wave.




Fig. 3.4.5: APHYN-EP predicted dynamics for the transmembrane potential diffusion of planar wave. The frames show a period of 8 ms of forecast obtained without re-training the APHYN-EP framework.

3.5 Conclusion

We have presented the APHYN-EP framework for modeling complex cardiac electrophysiology dynamics via a surrogate model combining simplified physics and deep neural networks. We demonstrated that this framework is able to reproduce with good precision the dynamics simulated by the Ten Tusscher – Noble – Noble – Panfilov and Ten Tusscher – Panfilov ionic models, even using a simplified electrophysiology model as a physical component of the framework.

Such framework opens up many possibilities in order to introduce prior knowledge in DL approaches through explicit equations, as well as to correct model errors from data.

Finally, real-data applications of this framework will be shown in the next chapters.

Appendix

3.A Learning and Forecasting from Small Time Sequences

In this Appendix, using 2D data without noise from section 3.3.1.2, we explored learning and forecasting abilities of APHYN-EP framework in conditions of training on small time sequences (about 6 ms of dynamics) representing only a depolarisation phase (i.e. as for the Ep-Net 2.0 framework in the Chapter 2).

We also compared APHYN-EP framework performance with the performances of each components of the framework used alone: only the “incomplete” physical model (F_p) and only the data-driven model (F_d , ResNet) trained on the same dataset and in the same conditions as APHYN-EP.

3.A.1 Results

Fig. 3.A.1 presents qualitative results on the forecast over 9 ms after assimilating only one first frame of dynamics. Importantly, one can observe a very good agreement between the ground truth and the forecast transmembrane potentials generated by APHYN-EP, as illustrated in Fig. 3.A.2, even if for training we used only the first 6 ms. The correction term brought by F_d is clearly visible.

Figures 3.A.2 and 3.A.3 represent the performance of different components of APHYN-EP along with their contribution to the final result. We can observe which part of the generated transmembrane potential was created by the physical component of the framework (see Fig. 3.A.3 (second row)). The data-driven component was used only to correct the difference between the ground-truth dynamics and the physical part (see Fig. 3.A.3 (third and fourth rows)).




Fig. 3.A.1: APHYN-EP predicted dynamics for the transmembrane potential diffusion. The figure shows a 9 ms of forecast).

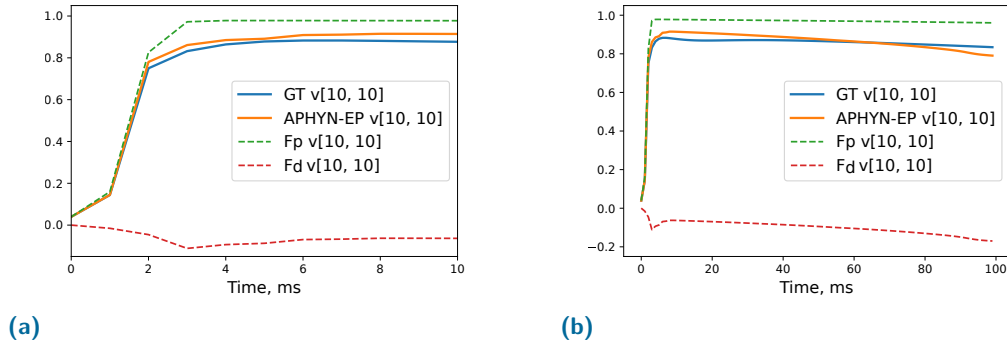


Fig. 3.A.2: (a,b) Transmembrane potential at point (10,10) in the cardiac slab (red point, see Fig. 3.A.1) with different forecasting horizons. Ground truth (GT), APHYN-EP, physical (F_p) and data-driven (F_d) component of APHYN-EP.




Fig. 3.A.3: Exemplary results illustrating: the dynamics of the transmembrane potential diffusion predicted by the APHYN-EP physical component (second row); the error with ground-truth diffusion for this physical component of APHYN-EP (third row); and, the trained APHYN-EP data-driven component contribution (bottom row).

Table 3.A.1 shows the mean squared error (MSE) results for our framework, for different forecasting horizons on validation data samples. To calculate this error, for each data sample, we fed the model with only one initial test measurement, then let it predict several steps forward without any additional information. We also added for comparison

two baseline models corresponding to the two components of our model, each used alone: only the “incomplete” physical model (F_p) and only the data-driven model (F_d , ResNet) trained on the same dataset as APHYN-EP, as described in 3.3.1. As we can notice, APHYN-EP captured the observed dynamics with good precision for different time horizons, even if for training we used only the first 6 ms. In the same time, the pure physical and the data-driven models when used alone struggle to learn the proper dynamics. Figure 3.A.4 visually confirms those numerical results.

The physical model (as well as APHYN-EP) correctly identifies the velocity and the activation time of transmembrane potential, but not the transmembrane potential values, due to its physical construction and limitations. The data-driven model can have a good precision, but it reproduces only average dynamics and is very sensitive to self-generated noise, which is crucial when forecasting.

Method	MSE (6 ms)	MSE (12 ms)	MSE (24 ms)	MSE (50 ms)
APHYN-EP	0.0057	0.0037	0.0029	0.002
Physical model only	0.0093	0.0111	0.0096	0.0085
Resnet model only	0.0195	0.0220	0.1593	9.9212

Tab. 3.A.1: Mean-squared error (MSE) of normalised transmembrane potential forecasting per time-step for different forecasting horizons.

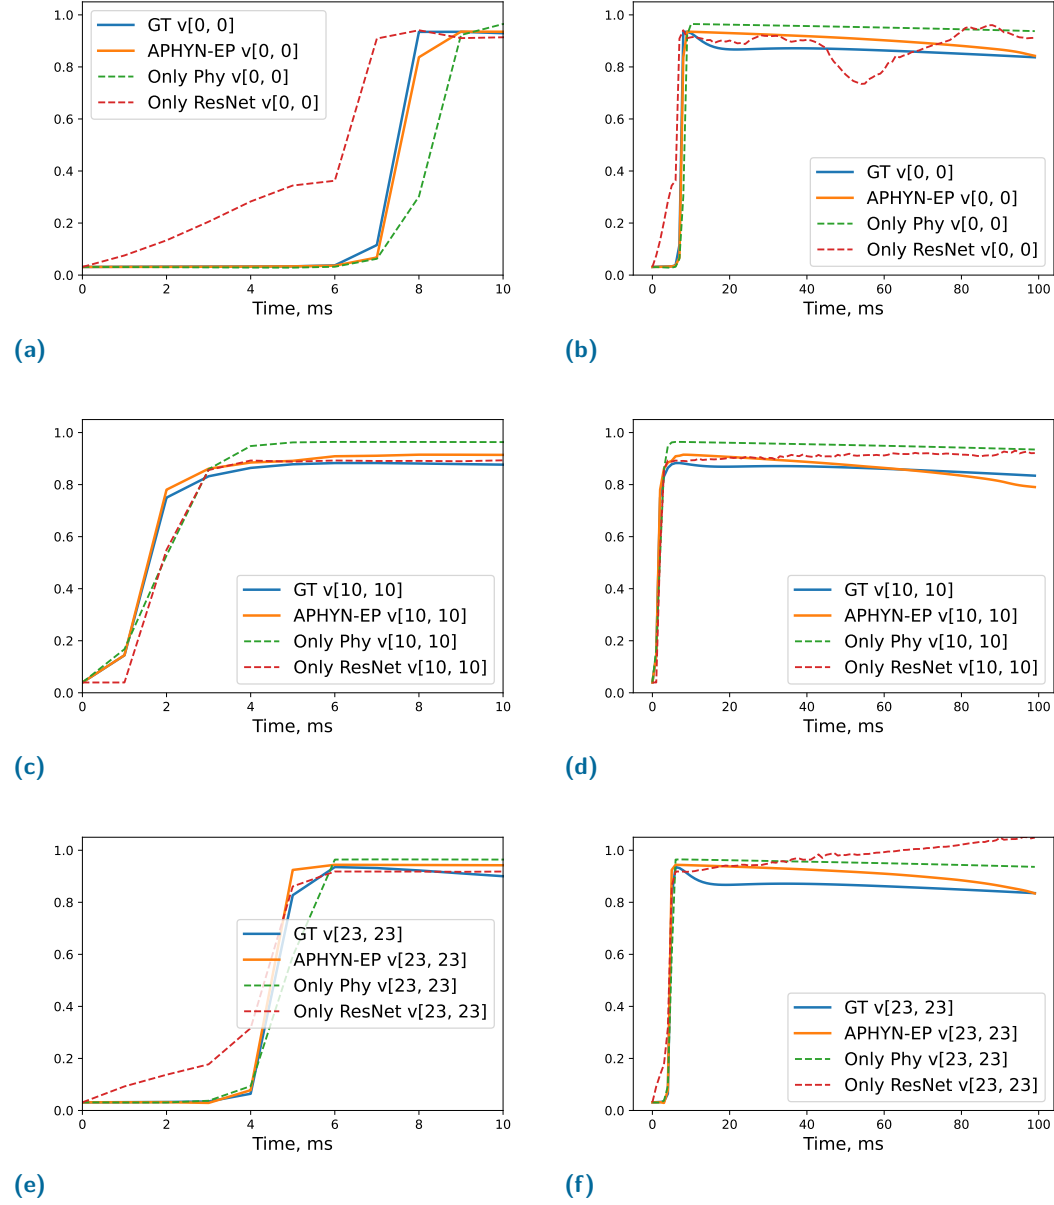


Fig. 3.A.4: Transmembrane potential ground truth (GT), generated by APHYN-EP, by Physical model and ResNet model at the leftmost upper point (1,1) (a,b), at point (10,10) (c,d) and the rightmost bottom point (23,23) (e,f) in the cardiac slab with different forecasting horizons, the same GT dynamics as in Fig. 3.A.1-3.A.2.

Personalisation of Cardiac Electrophysiology Model on Experimental Signals

Contents

4.1	Introduction	48
4.2	Materials and Methods	49
4.2.1	Data collection	49
4.2.2	Full personalisation of cardiac EP model	51
4.2.3	Fast personalisation of cardiac EP model	53
4.2.4	Training settings	56
4.3	Experiments and Results	56
4.3.1	Heart with an ischaemic region	56
4.3.2	Different pacing frequencies	63
4.4	Discussion and Conclusion	66

Abstract In chapter 3, we presented a fast physics-based deep learning framework to learn cardiac EP dynamics from data. This novel framework has two components, decomposing the dynamics into a physical term and a data-driven term, respectively. This construction allows the framework to learn from data of different complexity. Using in silico synthetic data, we demonstrated that this framework can reproduce the complex dynamics of transmembrane potential generated by two detailed ionic models, even in presence of noise in the data.

In this chapter, using ex vivo 0D epicardial optical fluorescence recordings of action potential, we demonstrated the ability of our APHYN-EP framework to identify the key physical parameters for different anatomical zones of the heart. Additionally, we proposed a method for our framework's fast adaptation to new physiological conditions (e.g. change in tissue functional properties, response to pacing at various stimulation rates).

The results of this chapter were published in part in Lecture Notes in Computer Science Proceedings of the 13th International Conference on Statistical Atlases and Computational Models of the Heart (STACOM) [[Kashtanova, 2022b](#)] (section 4.3.1.1).

Our main contributions in this chapter are:

- the evaluation of APHYN-EP framework's ability to learn the real OD cardiac EP data from ex vivo optical recordings (section 4.3);
- the identification via APHYN-EP framework of the key physical parameters for different anatomical zones of hearts (i.e., ischemic vs normal) (sections 4.3.1.1 and 4.3.2.1);
- the development and evaluation of a method for faster personalisation of cardiac EP model from cardiac OD ex vivo data (sections 4.2.3, 4.3.1.2 and 4.3.2.2).

4.1 Introduction

Multi-physics phenomena involved in cardiac function can be studied using mathematical models of different complexity. Among them, several electrophysiology (EP) models can accurately reproduce the electrical behaviour of the heart at cellular, tissue or organ level. However, biophysically detailed models are always computationally expensive and have numerous hidden variables which are nearly impossible to be measured all, making the model parameters difficult to personalise. Alternatively, simplified phenomenological models that have fewer variables and parameters, can be used for rapid computational modelling of wave propagation at tissue and organ level; however, these are less realistic and, consequently, need a complementary mechanism to enable their fitting to measured data.

Machine learning and in particular deep learning approaches could help provide such correction mechanisms. Thus, the combination of rapid phenomenological models and machine learning components may allow the development of rapid and accurate models of transmembrane dynamics.

In the previous chapter we introduced the novel APHYN-EP framework, a DL framework based on a fast low-fidelity (or incomplete) physical model, inspired by [Yin, 2021]. This framework has two components, decomposing the dynamics into a physical term and a data-driven term, respectively. Notably, the data-driven deep learning component is specifically designed to capture only the information that cannot be modeled by the incomplete physical model. In Chapter 3, we showed that this framework is able to reproduce with good precision the synthetic dynamics of full cardiac action potential

cycle simulated by the Ten Tusscher – Noble – Noble – Panfilov and the Ten Tusscher – Panfilov ionic models.

In this chapter and in chapter 5, we focus on testing the performance of the APHYN-EP framework using a real ex-vivo database obtained from optical fluorescence imaging of cardiac action potential (AP) presented in section 1.1.3.2.

Particularly, in this chapter, using ex vivo OD epicardial optical fluorescence recordings of action potential, we demonstrated the ability of our APHYN-EP framework to identify the key physical parameters for different anatomical zones of the heart. Additionally, we proposed a method of our framework's fast adaptation to new physiological conditions (e.g. change in tissue functional properties, response to pacing at various stimulation rates).

This combined model-based and data-driven approach could improve some of the current cardiac EP modeling methodologies and may provide a robust predictive tool.

4.2 Materials and Methods

4.2.1 Data collection

In this chapter, in order to test the performance of our APHYN-EP framework and to show its capability to reproduce real OD transmembrane potential dynamics from sources with different physiological properties, we selected two hearts from optical mapping database (presented in the section 1.1.3.2): the heart with an ischaemic region paced at one frequency, as well as the healthy heart with optical recordings of AP performed at different pacing frequencies resulting in different cycle lengths (i.e., heart rates). Using the optical images, we selected on the epicardium of each heart several rectangular regions of interest (ROIs) having different dynamics of AP signal across time. From all ROIs, we extracted a dataset of OD AP signals (note: details specific for each ROI are presented below). Then, on each dataset, we applied the methods described in sections 4.2.2 and 4.2.3.

4.2.1.1 Different zones of unhealthy heart: Heart with an ischaemic region

For this part, we used from our optical mapping database the representative pathologic heart having an ischaemic region. We manually selected two rectangular regions of interest (ROIs) with different action potential dynamics across time (see Figure 4.2.1). Next, we normalised the intensity of the optical signal in order to obtain a $[0, 1]$ min/max interval for the action potential, while keeping the noise in the data. We took a first full

cardiac cycle and removed the parts with zero potential, keeping only time-sequences of 300 ms per experiment. Then, we saved a time sequence for each pixel from each ROI in separate files, creating two databases (ROI A and ROI B, respectively), each containing 10 and 5 time-sequences for training and validation, respectively.

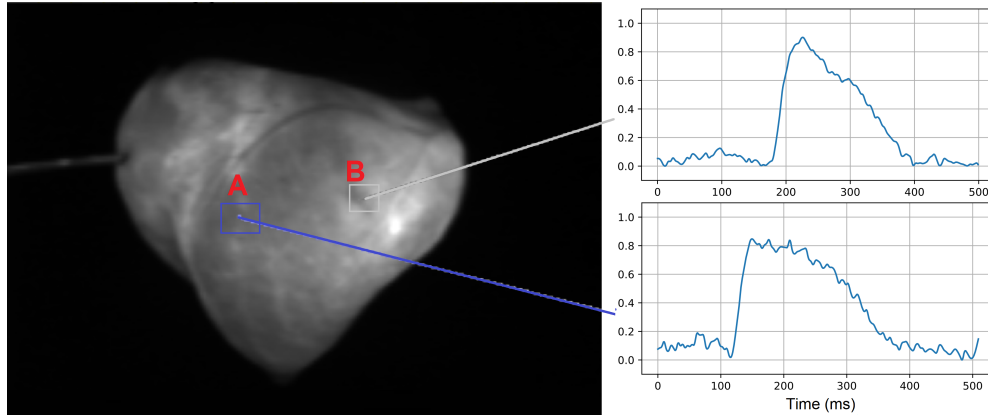


Fig. 4.2.1: Example of optical mapping data (tracings of denoised AP waves) recorded ex vivo in a porcine heart. ROI B represents an ischaemic region characterised by a shorter action potential duration (APD) compared to the normal APD recorded in ROI A.

The optical data were considered here as the ground truth (GT). Our specific objective was to learn the complex dynamics of measured action potential wave, and then to identify the relevant physical parameters for different regions of the heart.

4.2.1.2 Multiple pacing rates

For this part, we used a representative healthy heart from our optical mapping database, in which the AP recordings were performed by pacing the heart at 4 different frequencies. From each recording, we manually selected the same rectangular regions of interest (ROI) but with different AP dynamics across time (see Figure 4.2.2). Next, we normalised the intensity of the optical signal, to obtain a $[0, 1]$ min/max interval for the transmembrane potential, while keeping the noise in the data.

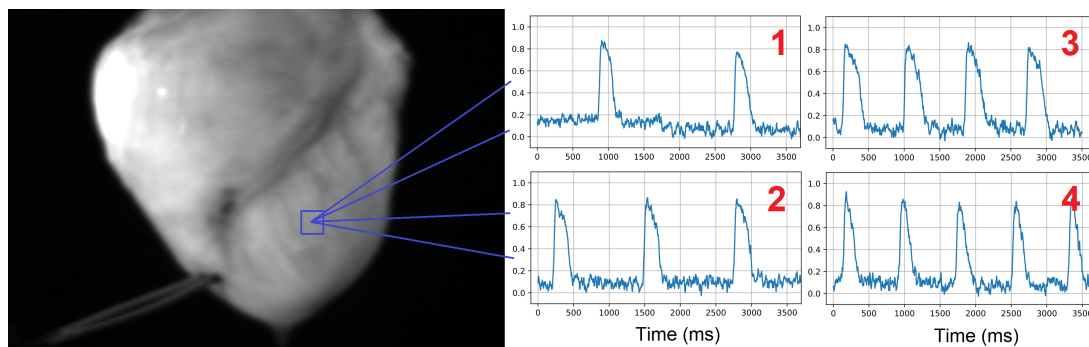


Fig. 4.2.2: Example of ex vivo optical mapping data with tracings of denoised AP waves recorded at multiple pacing rates in a healthy porcine heart.

For the personalisation experiment, we took the 2nd recording (with 3 AP waves), as the most representative and the least noisy example. From this recording, we selected the first full action potential cycle and removed the parts with zero potential, keeping only time-sequences of 300 ms per experiment. Then, we saved a time sequence for each pixel from the ROI in separate files, creating two databases containing about 10 and 5 time-sequences for training and validation, respectively.

The optical data were considered here as the ground truth. Our specific objective was to learn the complex dynamics of the measured action potential, and then to identify the relevant physical parameters for this zone.

From other optical recordings (including the full 2nd recording of 3 AP waves), we extracted a one time sequence corresponding to the central pixel of ROI. Then we saved these time sequences for testing phase, obtaining test recordings of approximately 3500 ms each.

4.2.2 Full personalisation of cardiac EP model

In this chapter, we explore the personalisation capacities of our APHYN-EP framework presented in section 3.2, by applying it to the above ex vivo optical mapping data.

As described earlier, this framework combines two components: i) a physical model (F_p) representing an incomplete description of the underlying phenomenon; and, ii) a neural network (F_d) which complements the physical model by capturing the information that cannot be modeled by the physical component.

The choice of each component is very important, as well as a good equilibrium between them. Should we choose to use a "weak" physical model which does not model enough the valuable information in the data, then the data-driven component would take all the learning. Thus, at the end of the APHYN-EP framework's training phase, we would still get a "black-box" model, giving sufficient but uninterpretable results. Alternatively, if we use a very detailed physical model, then its calculation and personalisation will be computationally demanding and therefore not suitable for training in the APHYN-EP framework.

For the experiments using the ex-vivo datasets, here we will continue to use the Mitchell-Schaeffer model [Mitchell, 2003] (presented in section 2.2) as the F_p component of the framework. This is a good candidate model because of the following reasons:

- it is a simplified biophysical model that still captures well the activation-recovery phases of AP;

- it provides a good analytical understanding of the transmembrane dynamics;
- it contains a limited number of parameters to estimate;
- each parameter has a simple physical interpretation;
- it has explicit analytical formula to express most of the measured features as well as the restitution properties using the model parameters [Mitchell, 2003];
- it models dynamics close to the ones obtained via optical fluorescence mapping of action potential [Relan, 2011b]

Figure 4.2.3 graphically illustrates the role of different biophysical parameters in the Mitchell-Schaeffer model.




Fig. 4.2.3: Impact of different Mitchell-Schaeffer model parameters on modelling of cardiac action potential cycles. The variable v represents a normalised ($v \in [0, 1]$) dimensionless transmembrane potential, variable h is “gating” variable controlling the repolarisation phase. From [Cedilnik, 2020].

Parameter τ_{in} controls the velocity of the inward current that acts to depolarise the membrane, controlling the steepness of the phase 0 slope of the action potential wave.

Parameter τ_{out} controls the slope of the repolarisation phase of the action potential.

Parameter τ_{close} is a governing parameter controlling the action potential duration.

In figure 4.2.3 we can also observe the contribution of J_{stim} component: when the stimulation current is not strong enough to raise v above a certain threshold v_{gate} , it is inefficient (as indicated in red). However, if the gating variable is sufficiently recovered, the depolarisation can happen even from a small stimulation (last cycle).

Additionally, the Mitchell-Schaeffer model successfully describes another phenomenon of the cardiomyocyte membrane. When the tissue is stimulated before being fully recovered (i.e., diastolic interval (DI) is shortened), the action potential duration (APD) of the following cardiac cycle is shortened (Fig. 4.2.4(left image)). This physiological response is achieved in the model through the construction of a second gating variable

h (see Fig. 4.2.4(right image)). This physical property, as well as the new APD, can be derived from the so-called restitution curve (Fig. 4.2.5), which has an explicit formulation and depends on the Mitchell-Schaeffer model parameters [Mitchell, 2003]. Note that APD adaptation to different pacing frequencies is essential in the genesis of arrhythmia; thus, accurate cardiac EP models must satisfy the restitution property.




Fig. 4.2.4: Transmembrane potential (v , left) and gating variable (h , right) curves. Legend: action potential duration (APD), diastolic interval (DI). From [Mitchell, 2003].




Fig. 4.2.5: The restitution curve derived from the 2-current model 2.1, including the threshold behavior. From [Mitchell, 2003].

This physical model will help us in the personalisation of APHYN-EP framework and provide a comprehensive knowledge about the estimated physical properties of the signal. Figure 4.2.6 is a schematic representation of the proposed full personalisation method.

4.2.3 Fast personalisation of cardiac EP model

Relan et al. [Relan, 2011b] demonstrated that the Mitchell-Schaeffer model, with physical parameters personalised for each heart region, can accurately model the ex vivo recorded data for a healthy heart.

To simplify the problem, we assumed that, with a "precise" F_p component, the F_d component will mainly capture the missing dynamics from ex vivo data which do not depend on the physical properties of heart zones (including recording inaccuracies due to light fluctuations and noise). In this case, we hypothesised that it is possible to obtain

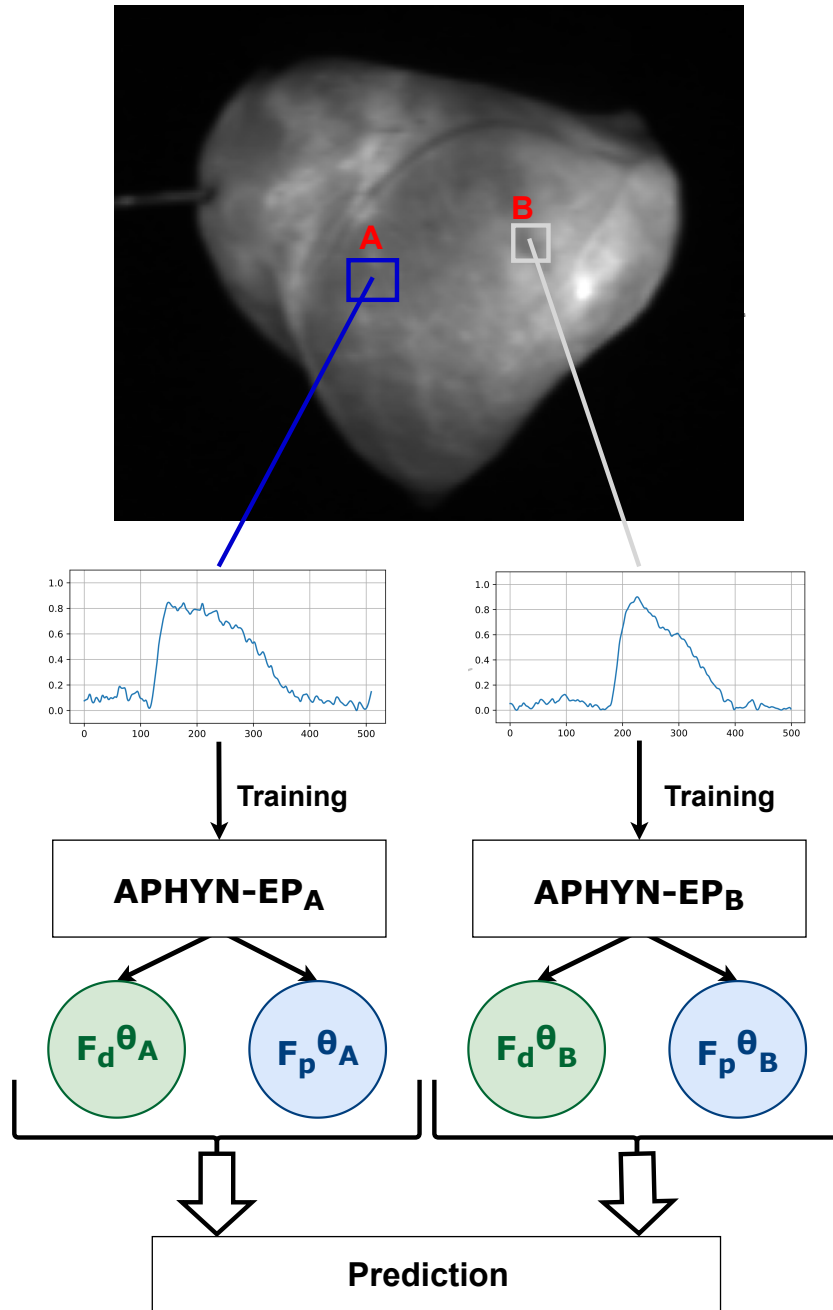


Fig. 4.2.6: Schematic diagram of the **Full personalisation method** for the APHYN-EP framework.

a single DL correction term (F_d) suitable for each selected region of interest (ROI) within the whole heart.

The general algorithm of this method is:

1. train only once a full APHYN-EP framework, using one ROI selected on the heart;
2. then, fit pre-trained APHYN-EP framework to a new ROI of the heart, by training only a physical component (F_p) of the framework (or using an external information

about physical parameters specific to this region), having a suitable prediction of AP dynamics specific for this ROI.

This approach will drastically reduce the training time and help the personalisation process. Figure 4.2.7 is a schematic representation of the proposed fast personalisation method. This method could also be applied using the full APHYN-EP framework trained on another ROI (e.g. ROI B).

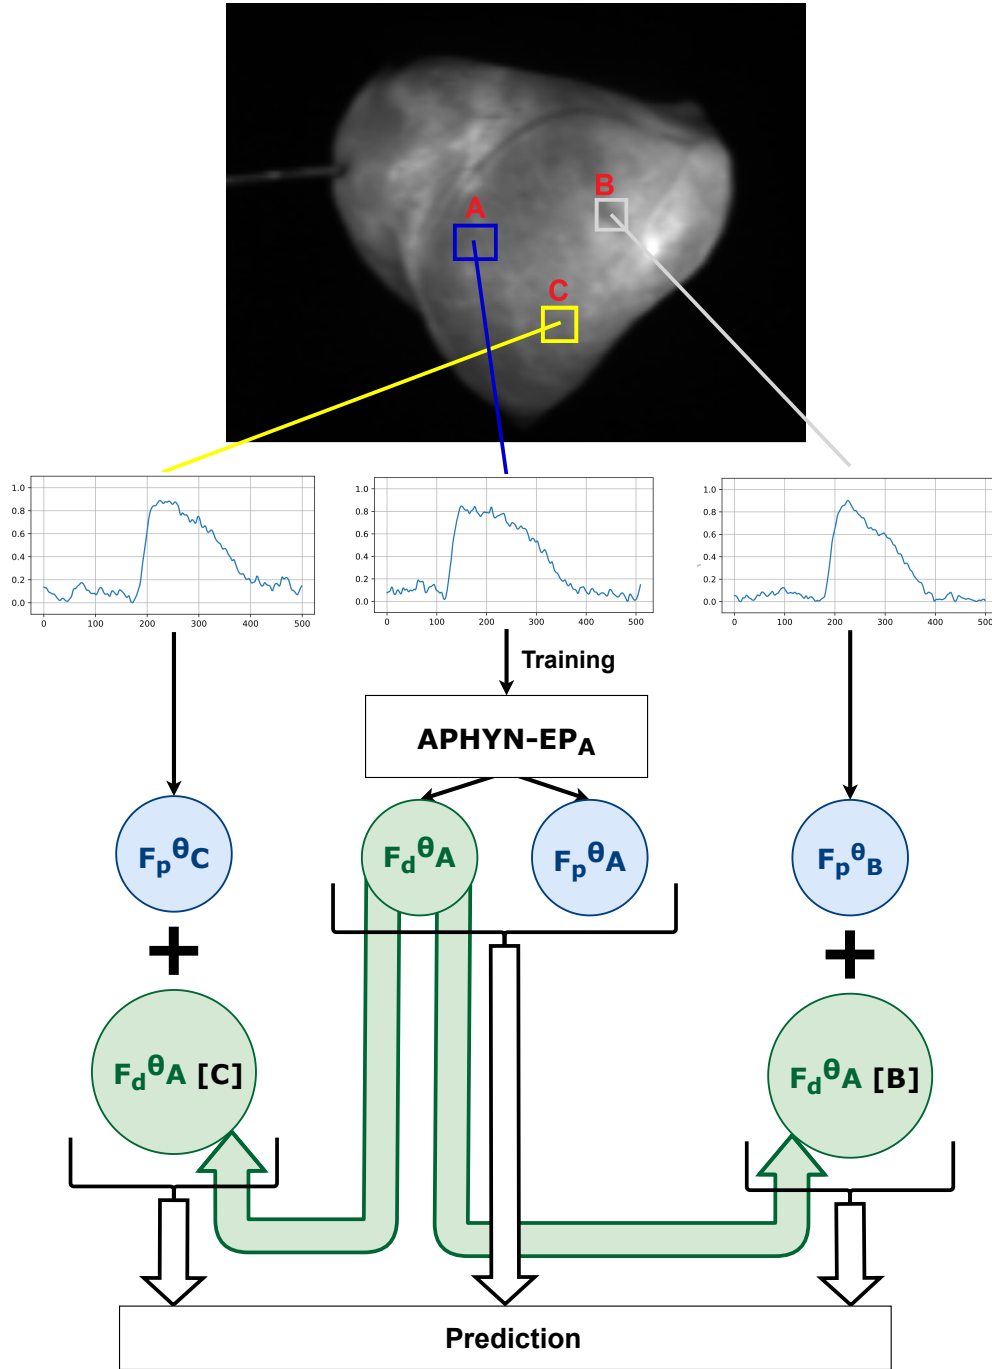


Fig. 4.2.7: Schematic diagram of the **Fast personalisation method** for the APHYN-EP framework.

To obtain more reproducible results, we performed the training of full APHYN-EP framework on the first ROI in two stages. First, we trained the F_p term of the framework to get an estimation of the physical parameters θ_p specific to the region. Second, we fixed the estimated θ_p parameters in the APHYN-EP framework and trained only the F_d term.

The results confirming the validity and reliability of this approach are presented in sections 4.3.1.2 and 4.3.2.2.

4.2.4 Training settings

The training settings for the experiments in this chapter are similar to those described in section 3.3.2 (OD data), except for the training horizon (i.e., 300 ms) and the parameter γ which was set to 1 for better equilibrium.

We also tested an MLP network (with 2 hidden layers of 200 neurons) as a deep learning component (F_d) of the framework, in order to compare the performance of our APHYN-EP framework with a simpler neural network.

Each training was conducted for 100 epochs, taking about half-hour by training on Nvidia Quadro M2200 GPU.

4.3 Experiments and Results

4.3.1 Heart with an ischaemic region

4.3.1.1 Parameter estimation with APHYN-EP using Full personalisation method

Using the optical imaging mapping data of the heart with ischaemic region, we showed that our APHYN-EP framework is able to reproduce the observed action potential dynamics for ROI A and ROI B within the heart, identifying the 3 major physical dynamics parameters (τ_{in} , τ_{out} and τ_{close}). For this experiments we trained APHYN-EP framework separately on both ROI, obtaining two trained versions of the framework specific for each ROI.

Figure 4.3.1 demonstrates that the framework was able to correctly estimate the differences in value for the parameter τ_{close} , which either increased the APD or shortened it, respectively.

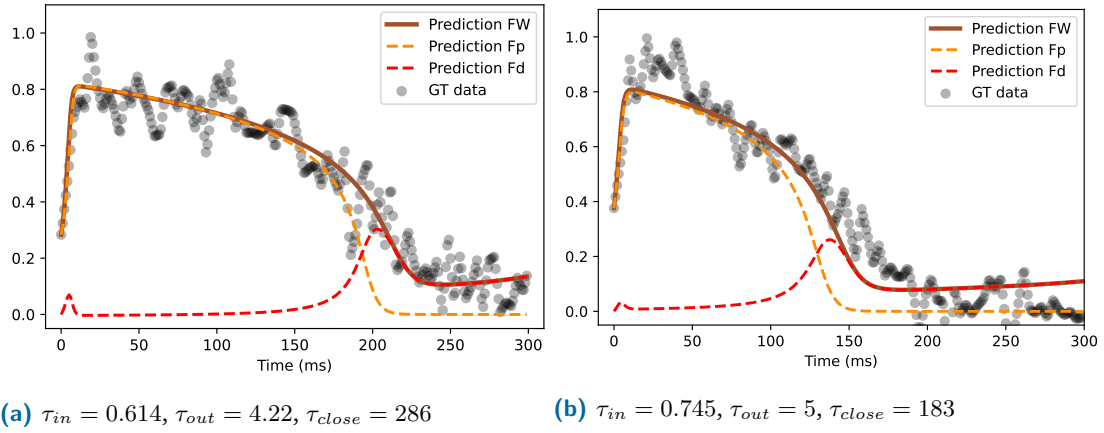


Fig. 4.3.1: Validation results of the framework trained on: (a) ROI A data and (b) ROI B data, respectively. Ground truth (GT) data, prediction of the framework (Prediction FW), the decomposition of prediction into physical (F_p) and DL (F_d) components.

Dataset	Method	Training data	Validation data
ROI A	APHYN-EP framework with ResNet ($\ F_d\ ^2$)	9.12 (0.16)	5.72 (0.08)
	APHYN-EP framework with MLP ($\ F_d\ ^2$)	9 (0.0785)	5.37 (0.077)
	Physical model	14	10
	Data-driven model	20	9.78
ROI B	APHYN-EP framework with ResNet ($\ F_d\ ^2$)	10 (0.08)	8 (0.07)
	APHYN-EP framework with MLP ($\ F_d\ ^2$)	8.79 (0.18)	7 (0.21)
	Physical model	14.5	9.3
	Data-driven model	7.78	6.79

Tab. 4.3.1: Mean-squared error, MSE ($\times 10^{-3}$) of the normalised, adimensional transmembrane potential, forecasting with forecasting horizon of 300 ms. Baseline models: the Physical model (2.1) and a fully data-driven model (EP-Net 2.0 [Kashtanova, 2021]) trained on the same dataset as APHYN-EP.

Table 4.3.1 summarises the quantitative results for our framework forecasting on training and validation data samples, in comparison to baseline methods trained on the same data. To calculate the associated error, for each data sample we fed each model with only one initial test measurement, and then allowed the model to predict 300 ms forward without any additional input of information.

The obtained MSE is relatively small for both ROIs, and, despite the use of a limited dataset for training, the APHYN-EP framework achieved forecasting the dynamics with good accuracy for new data samples from the validation dataset. Furthermore, our framework clearly outperformed the physical model for every dataset, while the contribution of F_d component was still minimal. Despite having good results on ROI B, the pure data-driven model encountered difficulties to learn the dynamics from ROI A data.

Figures 4.3.2, 4.3.3 and 4.3.4 visually confirm the quantitative results summarized in Table 4.3.1. One can also observe that the APHYN-EP framework with a MLP as F_d component has a similar performance with the APHYN-EP framework with a ResNet as F_d component; however, it is more sensitive to noise in the data and tends to overfit (see Fig. 4.3.2(b)).

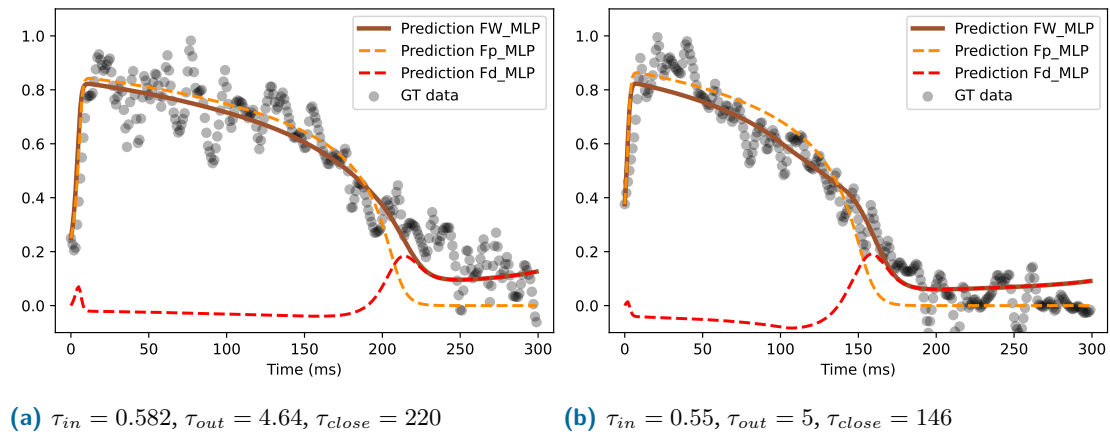


Fig. 4.3.2: Validation results of the framework with simpler NN (MLP) as F_d component trained on: (a) ROI A data and (b) ROI B data, respectively. Ground truth (GT) data, prediction of the framework (Prediction FW_MLP), decomposition of prediction on physical (F_p) and DL (F_d) components.

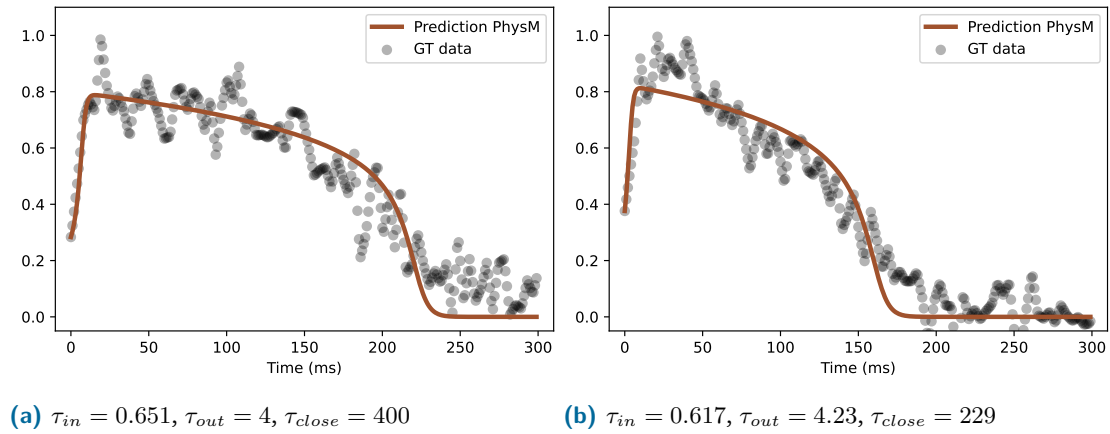


Fig. 4.3.3: Validation results of the Physical model (2.1) trained alone on: (a) ROI A data and (b) ROI B data, respectively. Ground truth (GT) data, prediction of the physical model (Prediction PhysM).

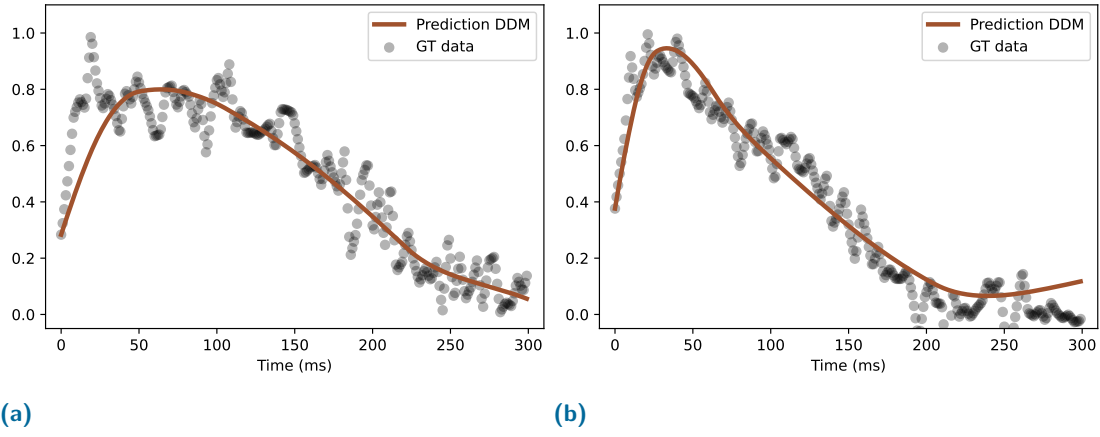


Fig. 4.3.4: Validation results of a fully data-driven model (EP-Net 2.0 [Kashtanova, 2021]) trained on: (a) ROI A data and (b) ROI B data, respectively. Ground truth (GT) data, prediction of the data-driven model (Prediction DDM).

4.3.1.2 Generalisation: Fast Personalisation to new region

Having noted that the obtained correction of the DL component is similar for both ROIs (Fig. 4.3.1), we assumed that it is possible to obtain a single DL correction term (F_d) suitable for each selected ROI in the whole heart.

To prove our hypothesis, we first performed a test where we simply replaced the F_d component of APHYN-EP trained on data in one ROI by the F_d component of APHYN-EP trained on data from another ROI (and vice versa), and obtained confirming results.

Second, we performed a series of personalisation experiments. Since we can obtain an estimation of physical parameters specific for each ROI (e.g. by training only the F_p component of the framework), we fixed the estimated θ_p parameters (for ROI A and ROI B, respectively) in the APHYN-EP framework and trained only the F_d component (separately for each ROI). Next, we replaced in the framework the F_d component trained on ROI A data ($F_{d_A}(A)$) by F_d component trained on ROI B data ($F_{d_B}(B)$) (and vice versa), and achieved similar results in forecasting (see Fig. 4.3.5-4.3.8). Figures 4.3.9 and 4.3.10 demonstrate the Fast Personalisation method application results on new region (ROI C).

Figures 4.3.6, 4.3.8 and 4.3.10 visually confirm our initial hypothesis, suggesting that in order to obtain an appropriate forecast of OD ex vivo dynamics for any region selected on the epicardium, it is sufficient to train the whole APHYN-EP framework using only one selected ROI.

ROI A experiments




Fig. 4.3.5: (a) Validation results of the F_p trained alone on ROI A data to obtaining $\theta_p(A)$ parameters. (b) Validation results for APHYNITY framework (with fixed $\theta_p(A)$ parameters) trained on ROI A data. (c) Validation results of the APHYNITY framework (with fixed $\theta_p(A)$ parameters) with F_d component trained on ROI B data. Legend: ground truth (GT) data, prediction of the framework (Prediction FW), decomposition of prediction on physical (F_p) and DL (F_d) components.

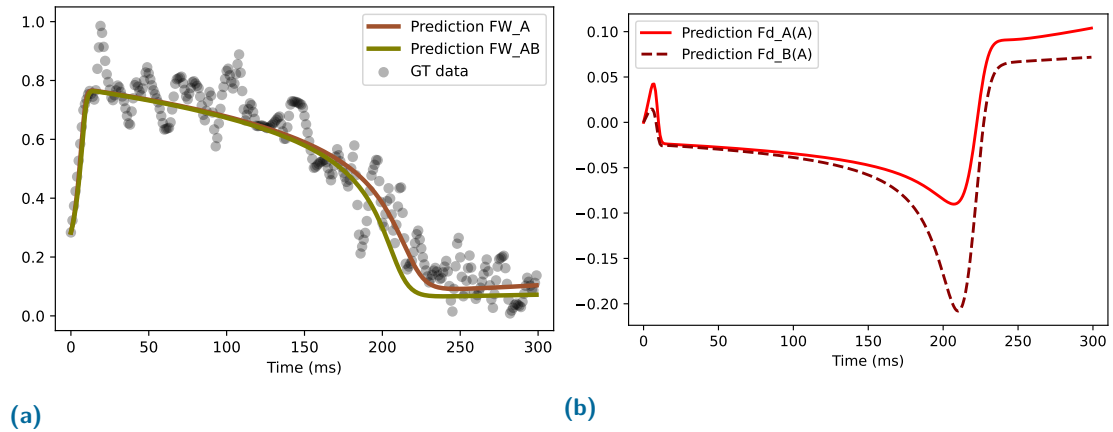


Fig. 4.3.6: (a) Validation results of the framework for ROI A, with the following legend: Ground truth (GT) data, prediction of the framework (Prediction FW_A) trained on only ROI A data, and prediction of the framework composed of (F_p) component trained on ROI A data and DL (F_d) component trained on ROI B data (Prediction FW_AB). (b) Comparison of DL component trained on ROI A data ($F_{d_A}(A)$) and on ROI B data ($F_{d_B}(A)$) applied on validation ROI A data.

ROI B experiments




Fig. 4.3.7: (a) Validation results of the F_p trained alone on ROI B data to obtain $\theta_p(B)$ parameters. (b) Validation results of the APHYNITY framework (with fixed $\theta_p(B)$ parameters) trained on ROI B data. (c) Validation results of the APHYNITY framework (with fixed $\theta_p(B)$ parameters) with F_d component trained on ROI A data. Legend: ground truth (GT) data, prediction of the framework (Prediction FW), decomposition of prediction on physical (F_p) and DL (F_d) components.

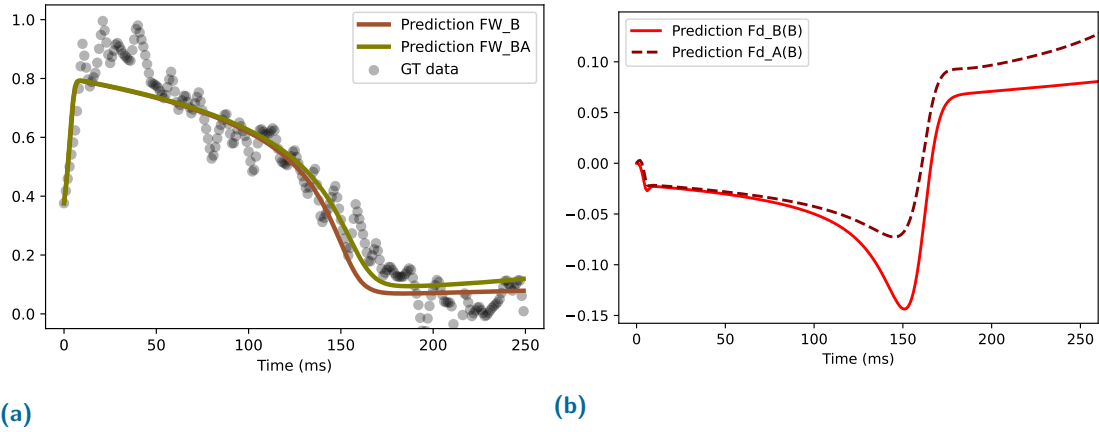


Fig. 4.3.8: (a) Validation results of the framework for ROI B, with the following legend: Ground truth (GT) data, prediction of the framework (Prediction FW_B) trained on only ROI B data, and prediction of the framework composed of (F_p) component trained on ROI B data and DL (F_d) component trained on ROI A data (Prediction FW_BA). (b) Comparison of DL components trained on ROI B data ($F_{dB}(B)$) and on ROI A data ($F_{dA}(B)$) applied on validation ROI B data.

ROI C experiments




Fig. 4.3.9: (a) Validation results of the F_p trained alone on ROI C data to obtain $\theta_p(C)$ parameters. (b) Validation results of the APHYNITY framework (with fixed $\theta_p(C)$ parameters) with F_d component trained on ROI A data. (c) Validation results of the APHYNITY framework (with fixed $\theta_p(C)$ parameters) with F_d component trained on ROI B data. Legend: ground truth (GT) data, prediction of the framework (Prediction FW), decomposition of prediction on physical (F_p) and DL (F_d) components.

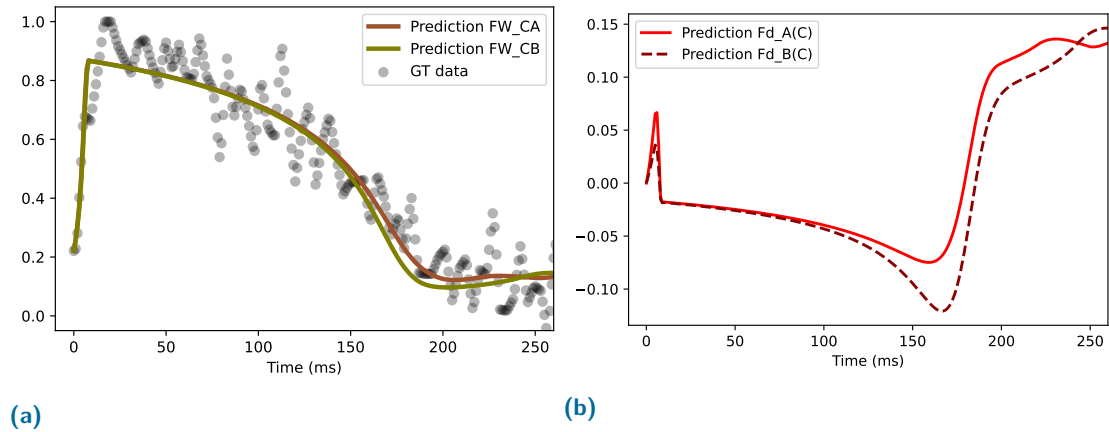


Fig. 4.3.10: (a) Validation results of the framework for ROI C, with the following legend: Ground truth (GT) data, prediction of the framework composed of (F_p) component trained on ROI C data and DL (F_d) component trained on ROI A data (Prediction FW_CA), and prediction of the framework composed of (F_p) component trained on ROI C data and DL (F_d) component trained on ROI B data (Prediction FW_CB). (b) Comparison of DL component trained on ROI A data ($F_{d_A}(C)$) and on ROI B data ($F_{d_B}(C)$) applied on ROI C data.

4.3.2 Different pacing frequencies

4.3.2.1 Parameter estimation with APHYN-EP using Full personalisation method

In this section, using optical imaging mapping data recorded on a healthy heart, we show that our APHYN-EP framework is able to reproduce the observed AP dynamics.

Having proved the capability of our framework to correctly estimate relevant physical parameters (see section 4.3.1.1), and to test further the robustness of the model, we decided to reduce the constraints on the norm of F_d in the Loss (3.2) during the training, allowing the framework to auto-regulate itself. Along with that, our purpose was to demonstrate two main points:

- even in the absence of constraints, F_d component does not erase F_p component completely;
- data-driven F_d component does only the correction functions to the physical F_p component, which is demonstrated in the next section.

One can observe that the dynamics predicted via APHYN-EP framework are similar to the those obtained for ROI A in the section 4.3.1.1; however, the correction from the F_d term is larger.




Fig. 4.3.11: Validation results of the framework trained on data described in section 4.2.1.2, $\tau_{in} = 0.4$, $\tau_{out} = 3$, $\tau_{close} = 350$. Ground truth (GT) data, prediction of the framework (Prediction FW), decomposition of prediction into physical (F_p) and DL (F_d) components, respectively.

4.3.2.2 Generalisation: Fast Personalisation to new frequency

To perform the testing experiments in this section, we used the trained APHYN-EP framework from 4.3.2.1 and applied it on full sequences of registered ex vivo AP dynamics.

Note that we applied the APHYN-EP framework without any re-training, except adding additional stimulations at the time when stimulation occurs in the recorded test data.

In Figures 4.3.12-4.3.15, one can observe that the resulting dynamics is in close resemblance with the ground truth dynamics, except when the additional noise occurs in the data. The contribution of F_d component is visible and triggered only by a sufficient stimulation of AP with F_p component.

Figures 4.3.15 demonstrate the effect of shortened APD due to early stimulation, discussed in the section 4.2. It is important to notice that it is necessary to simulate several cardiac cycles with a new stimulation frequency (pacing rate) in F_p to allow for tissue adaptation and to reproduce this phenomenon repetitively (as in the optical recordings).

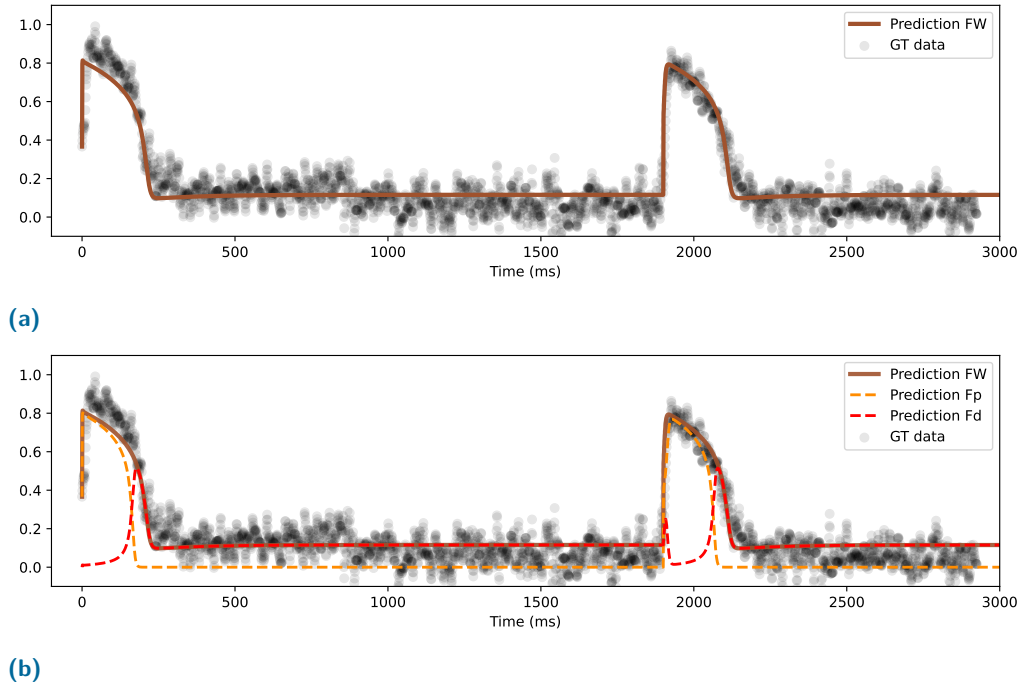




Fig. 4.3.12: Test results of the framework trained on data described in section 4.2.1.2 and its decomposition on physical and DL parts. Ground truth (GT) data, prediction of the framework (Prediction FW), decomposition of prediction on physical (F_p) and DL (F_d) parts.




(a)




(b)

Fig. 4.3.13: Test results of the framework trained on data described in section 4.2.1.2 and its decomposition on physical and DL parts. Ground truth (GT) data, prediction of the framework (Prediction FW), decomposition of prediction on physical (F_p) and DL (F_d) parts.



(a)



(b)

Fig. 4.3.14: Test results of the framework trained on data described in section 4.2.1.2 and its decomposition on physical and DL parts. Ground truth (GT) data, prediction of the framework (Prediction FW), decomposition of prediction on physical (F_p) and DL (F_d) parts.

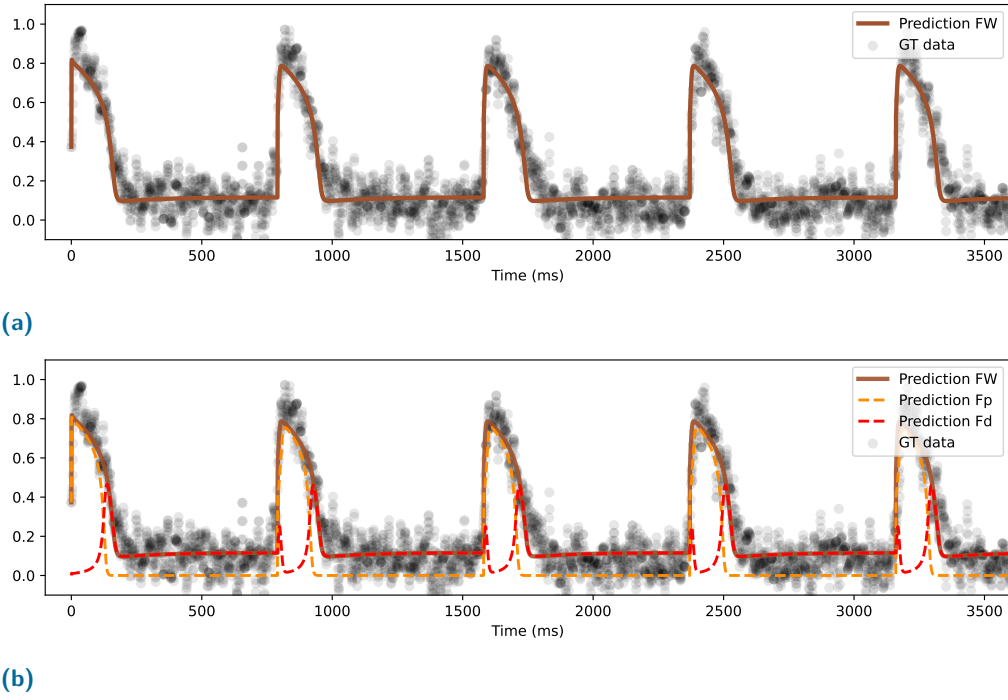


Fig. 4.3.15: Test results of the framework trained on data described in section 4.2.1.2 and its decomposition on physical and DL parts. Ground truth (GT) data, prediction of the framework (Prediction FW), decomposition of prediction on physical (F_p) and DL (F_d) parts.

4.4 Discussion and Conclusion

In this chapter, we demonstrated that the learning framework presented in chapter 3 is able to learn the considered real OD dynamics. Overall, our results suggest that automated learning of cardiac EP signals is feasible and has great potential for such applications. This framework may be useful for applications concerning fast parameter estimation of computational heart models.

In section 4.3, we demonstrated that APHYN-EP can identify the major biophysical parameters for different heart regions. Alongside, we successfully demonstrated examples of our framework's fast adaptation to new physiological conditions (e.g. change in tissue functional properties, response to pacing at various stimulation rates).

However, we acknowledge that some limitations exist in our model. For instance, due to the data-driven architecture of the framework, its training is not regular. This could lead to a local minimum for the parameters of the physical component and additional involvement of the deep learning component.

It is also important to notice that an extensive noise level in the training data could be learned by the F_d component and could induce its additional value enhancement (see

in Fig. 4.3.1 the tail of F_d curves). This can provoke an additional stimulation of AP even without the F_p component activation. Figure 4.4.1 presents an example of such spontaneous AP activation produced by the F_d component alone.

The potential solutions for these problems may be given by:

- A more advanced training protocol, that can use for example separate sequential training of F_p and F_d components, as used for Fast personalisation method experiments (see section 4.2.3).
- Adding rigid boundaries on the physical model parameters.
- Performing ablation studies on the DL component of the framework, including its architectural changes.
- Data cleaning from noise artifacts between cardiac cycles.

This could be addressed in future work.




Fig. 4.4.1: Example of spontaneous AP activation produced by F_d component.

Deep Personalisation of Cardiac Electrophysiology Model on 2D Optical Mapping Data.

Contents

5.1	Introduction	70
5.2	Materials and Methods	72
5.2.1	Data collection	72
5.2.2	Diffusion tensor	72
5.2.3	Deep personalisation of cardiac EP model	75
5.2.4	Training settings	79
5.3	Experiments and Results	80
5.3.1	Isotropic diffusion	80
5.3.2	Anisotropic diffusion	80
5.3.3	Generalisation: prediction extension to the whole heart	83
5.4	Discussion and Conclusion	90

Abstract In chapter 3, we presented a fast physics-based deep learning framework (APHYN-EP) able to learn cardiac EP dynamics from data. Using *in silico* data, we demonstrated that this framework can reproduce the complex dynamics of transmembrane potential generated by two detailed ionic models, even in presence of noise in the data. In chapter 4, using *ex vivo* 0D optical fluorescence mapping data of action potential, we demonstrated the ability of the APHYN-EP framework to identify the key physical parameters for different anatomical zones of the heart and to adapt to changes in its physical part F_p (i.e., as a change of physical parameters values and different stimulation rates).

In this chapter, using *ex vivo* 2D optical mapping data of action potential, we investigated further the capacities of our APHYN-EP framework to learn and reproduce complex cardiac EP data, as well as to generalise to new conditions. We presented the framework's potential to automatic learning of AP wave local conduction velocity, as imposed by the anisotropic arrangement of cardiac fibers, and its extrapolation to the entire organ.

The results of this chapter were in part submitted to a journal [Kashtanova, 2023] (section 5.3.1).

Our main contributions in this chapter are the following:

- the evaluation of APHYN-EP framework’s ability to learn the real 2D ex vivo recorded cardiac EP data (section 5.3);
- the implementation and evaluation of various types of diffusion tensors to model the depolarisation wave propagation (section 5.2.2);
- the development and evaluation of a method of cardiac EP model deep personalisation on 2D ex vivo data (section 5.2.3);
- the development and evaluation of a robust method to predict extension from small patches to the entire organ (sections 5.2.3.2 and 5.3.3).

5.1 Introduction

In silico modelling of cardiac electrophysiology has been an important research topic for the last decade (see section 1.2), holding the potential to be a very efficient tool for better mechanistic understanding of arrhythmia genesis. Personalisation of such models to experimental data is needed in order to test their realism and predictive power; however, this remains a difficult task to be accomplished at the whole organ scale. Typically, personalisation is defined as the estimation of model parameters which best fit simulations to data, but it is not less significant for the model to be able to automatically adapt to new patient data and to give a suitable prediction of new dynamics forecasting.

In the previous chapter, we demonstrated the ability of our framework to identify the key physical parameters for different anatomical zones of the heart using 0D real data acquired ex vivo. Additionally, we proposed a method for the fast framework adaptation to new conditions, such as changes of physical parameters’ values and different stimulation rates.

In this chapter, we work with 2D real ex vivo optical fluorescence mapping data of action potential; therefore, we have to include a depolarisation wave propagation through the cardiac tissue in overall learned real dynamics.

Traditionally, the modelling of this wave propagation employs an EP model calculation using a computational approach on specific grids (e.g. tetrahedral mesh for the Finite Element Method [Desrues, 2021; Cedilnik, 2018; Relan, 2011b] or a regular grid for the Lattice Boltzmann Method [Rapaka, 2012a; Campos, 2016]). Specifically, it mostly consists in determining the local conduction velocities for each grid point in order to compute a full wave propagation. The estimation of a local conduction velocity is typically obtained directly from invasive catheter-based electro-anatomical mapping, or indirectly from non-invasive imaging data (such as MRI or CT). Figure 5.1.1 presents an example of such indirect 3D estimation associated to a patient heart thickness map obtained from a CT scan. However, the depolarisation wave propagation modelling in this state cannot overcome the discrepancy between the idealised mathematical model and the measured data.

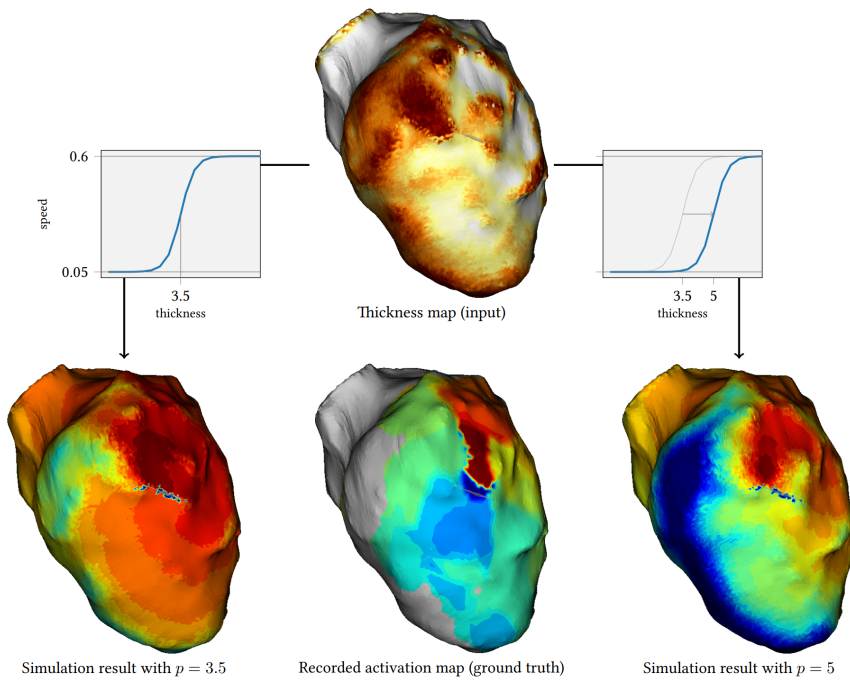


Fig. 5.1.1: Comparison of measured and simulated 3D activation maps obtained by changing the parameter value for the wall thickness, which is further correlated to the velocity transfer function. From [Cedilnik, 2018].

In our APHYN-EP two-component framework, the learning of depolarisation wave propagation is conducted automatically from data and simultaneously by two components. During the training, F_p component learns the general conductivity parameters which directly influence the speed of the depolarisation wave. The F_d component learns the additional local spread and directions of the observed depolarisation wave.

This construction allows our framework to learn and forecast the real 2D cardiac ex vivo dynamics, even when the available training data is limited. Furthermore, it will help the framework with the generalisation to new conditions (e.g. pacing from different

locations) and with extrapolation of obtained results to the entire 2D surface of the heart.

5.2 Materials and Methods

5.2.1 Data collection

We tested the APHYN-EP framework performance using ex vivo experimental datasets from optical fluorescence imaging of action potential (described in details in the section 1.1.3.2). For the experiments in this chapter, we chose one healthy heart stimulated via a pacing electrode, which was placed onto two different zones of the heart (see Fig. 5.2.1), having two recording of depolarisation wave propagation.

For one recording, we normalised the optical signal intensity in order to obtain a $[0, 1]$ min/max interval for the transmembrane potential, slightly denoising the 2D data. We manually selected a rectangular region of interest ROI (see Fig. 5.2.1), where no changes in the tissue physical properties were present. Next, we divided this ROI into smaller 2D squares (patches) of 10×10 pixels. For each patch, we took a first full cardiac cycle and removed the frames with zero potential, keeping only 2D time-sequences of 300 ms per experiment. Then, we saved each obtained 2D time sequence, having about 120 and 30 samples for training and validation respectively. The optical data recorded in the case of pacing from the right ventricle, RV (see Fig. 5.2.1(a)), were considered here as the ground truth for training and validation, while the optical data recorded with the pacing electrode placed onto the left ventricle, LV (see Fig. 5.2.1(b)), were used for testing.


5.2.2 Diffusion tensor

In order to model the 2D cardiac electrophysiology dynamics, the Mitchell-Schaeffer model [Mitchell, 2003] (presented in section 2.2) was used as a F_p component of the framework, and here it included a diffusion propagation term. Notably, this term was neglected in the previous chapter describing the 0D experiments.


The diffusion term in the model:

$$\text{div}(\Psi \nabla v)$$

(i.e., the first component of 2.1) is controlled by the diffusion tensor Ψ , which can be isotropic or anisotropic, as per the user choice to model an isotropic or an anisotropic depolarisation wave propagation, respectively (see Fig. 5.2.2). The anisotropic case is more realistic, as it considers the muscle arrangement in the myocardial tissue, with the AP wave propagating faster in a direction parallel to the fibers.



(a) RV pacing



(b) LV pacing

Fig. 5.2.1: Example of depolarisation maps selected from the optical data recorded ex vivo in a porcine heart, with the stimulating electrode (used for pacing the heart) placed onto: (a) the RV, and (b) the LV, respectively. Red areas correspond to early activation times (i.e., where the excitation pulse was delivered) while the late depolarisation times are depicted in green-blue.




Fig. 5.2.2: Example of simulated transmembrane potential (yellow) propagation in a 2D cardiac tissue slab modelled using: (a) isotropic diffusion tensor; and, (b) anisotropic diffusion tensor (note that the stimulation started in the centre of the square).

In this chapter we tested both types of diffusion tensor forms, demonstrating the advantages and disadvantages of each approach.

5.2.2.1 Isotropic diffusion

For initial experiments applied on in silico 2D data (see section 3.3.1.2), we used an isotropic diffusion tensor, because of the simplicity of observed data. The isotropic diffusion tensor was represented in this form:

$$\Psi = \sigma I$$

where σ is a general conductivity parameter.

The F_p component of the APHYN-EP framework with an isotropic tensor is able to model only a simplified propagation of EP dynamics (Fig. 5.2.2(a)) and does not take into account the specific properties of cardiac tissue. However, it allows the F_d component of the framework to independently learn the correction of missing tissue specific dynamics.

The results of our APHYN-EP framework with the F_p component included (based on an isotropic diffusion tensor) on ex vivo data, are demonstrated in the section 5.3.1.

5.2.2.2 Anisotropic diffusion

Historically, there has been many studies of the ventricular tissue structure, composition and architecture [Thomas, 1957; Streeter, 1979; Nielsen, 1991; Fenton, 1998]. It has been proved that:

- the myocardial cells are shaped as flattened tubes;
- they are arranged in sheets roughly parallel to the outer/inner surfaces (i.e., epicardium and endocardium, respectively) of the heart;
- the fiber axis (i.e., long axis of the cell) rotates continuously between the top and the bottom sheets by a certain angle $\Delta\alpha$, and the rotation is counter-clockwise from epicardium to endocardium as viewed from the top of the heart.

Taking into account these characteristics, we can use an anisotropic diffusion tensor, which considers the fibre orientations, in the following 2D form [Rodríguez-Padilla, 2018] (excluding the impact of transmural propagation):

$$\Psi = \begin{pmatrix} \Psi_{11} & \Psi_{12} \\ \Psi_{21} & \Psi_{22} \end{pmatrix}$$

where the coefficients are:

$$\begin{aligned}\Psi_{11} &= \sigma_{||} \cos^2(\alpha) + \sigma_{\perp} \sin^2(\alpha) \\ \Psi_{22} &= \sigma_{||} \sin^2(\alpha) + \sigma_{\perp} \cos^2(\alpha) \\ \Psi_{12} = \Psi_{21} &= (\sigma_{||} - \sigma_{\perp}) \cos(\alpha) \sin(\alpha)\end{aligned}$$

The $\sigma_{||}$ and σ_{\perp} represent conductivity parameters in parallel and perpendicular direction to the fiber axis.

In section 5.3.2, using ex vivo data, we present the performance of APHYN-EP with a F_p component able to model anisotropic dynamics (Fig. 5.2.2(b)).

5.2.3 Deep personalisation of cardiac EP model

One of the main goals of developing the APHYN-EP framework is to be able to apply this framework on new unseen data and get a suitable forecast. In this chapter we specifically explored the capability of our framework to reproduce the action potential wave characteristics obtained at different pacing locations.

Due to the heterogeneity in the heart shape and different physical characteristics among the hearts (and within each heart) in our Optical mapping database (presented in section 1.1.3.2), as well as to the variability in noise and quality of recordings, we were not able to create a larger database with whole hearts that would be sufficient for training.

Therefore, instead of using the full recordings obtained on whole hearts, we decided to use a set of small patches cut out from initial whole heart images. Furthermore, this significantly reduced the training time and the needs in computational resources.

Additionally, considering the data-driven construction of our framework, we were able to extrapolate the locally learned dynamics to the entire organ.

The general steps of our Deep Personalisation method are:

1. Cut small 2D patches containing various AP dynamics from a normalised whole heart data recording.
2. Train the APHYN-EP framework on these patches of dynamics, obtaining output of the framework θ_p^{patch} and θ_d^{patch} specific for each patch.

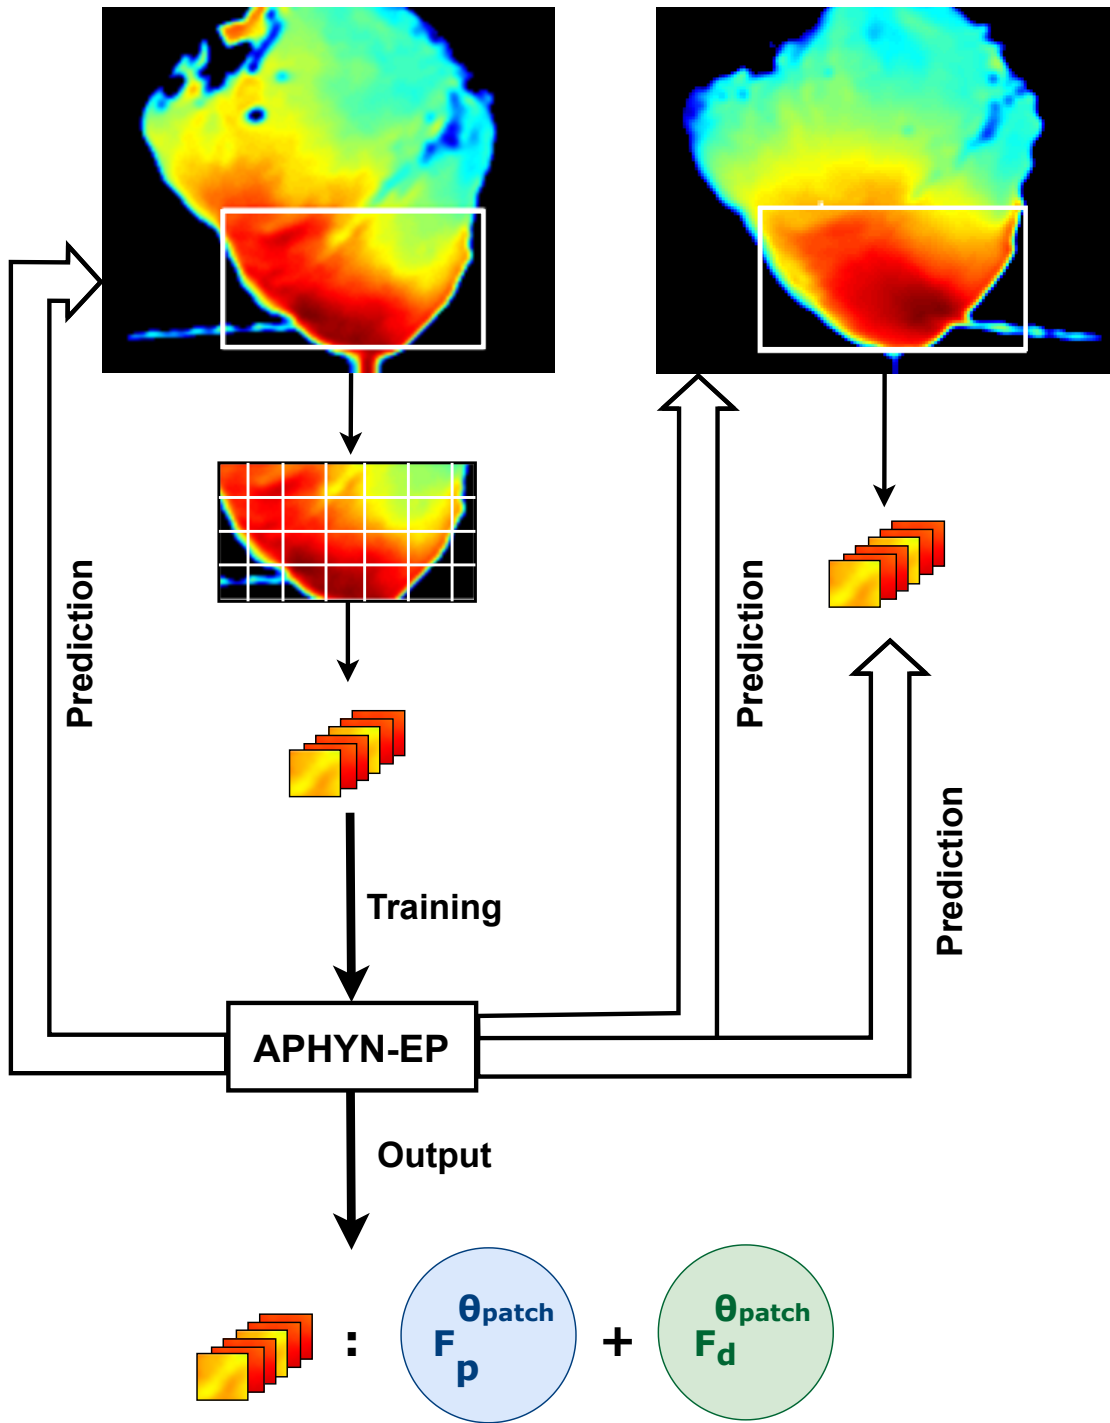


Fig. 5.2.3: Schematic diagram of the APHYN-EP framework using **Deep Personalisation method**. Personalisation on one heart with two pacing locations indicated by the position of the electrode (left vs. right).

3. Make a prediction of dynamics on patches with the same physical properties, but containing new dynamics (for example, pacing applied from a different location).

3*. Or, extend learned dynamics on a whole heart data recording.

Figure 5.2.3 presents a schematic representation of the proposed approach.

To simplify the problem, we trained the APHYN-EP framework on small patches cut out from the region where there were no significant changes in tissue EP properties, having the $F_p^{\theta^{patch}}$ and $F_d^{\theta^{patch}}$ common for all cut patches used for training. This way the APHYN-EP framework can learn the local 2D dynamics without over-fitting on specific cases (see section 5.3.1.1).

5.2.3.1 Extension of prediction via convolutions

Convolution-based neural networks (NNs) (such as CNN, ResNet etc.) are specifically designed to process image data of different sizes. In contrast to fully-connected NNs, they do not require to use a large number of neurons when they learn data of high-resolution consisting of large number of pixels. On the contrary, they try to identify a general feature map from data, thanks to the shared-weight architecture of the convolution kernels and filters that slide along an input picture.

Due to their convolution layers, these networks have a potential to extend the acquired knowledge to images of a larger size than of those used during the training.

The results of applying the APHYN-EP framework (trained on patches) on the ex vivo cardiac EP dynamics of the entire heart, are demonstrated in sections 5.3.3.1 and 5.3.3.2.

5.2.3.2 Extension of prediction from patches to the whole heart

The extension of locally learned dynamics from small image patches (used for training) to a larger area (or the entire exposed surface of the optically imaged heart) can be also achieved via patch-by-patch reconstruction, following next algorithm:

1. cut this area (large image) into patches of the same size as used during the training;
2. apply the framework separately on each image patch;
3. restore the initial image from patches, obtaining the forecast of dynamics on full image.

However, a complicating factor arises when an image patch does not contain any appearing dynamics, which could be a case for full-heart image reconstruction.

Since the initial idea of our framework usage was to assimilate only several ms of the initial dynamics and to generate the forecast without any further information, here we

developed a dynamics-based algorithm of patch reconstruction, improving the previously presented one.

Dynamics based algorithm for patch-by-patch reconstruction (forecasting time horizon of T ms):

1. Cut this area (large image) into **overlapping** patches having the same size as used during the training.
2. Apply the framework separately on every image patch, following a specific order:
 - 0) first, apply the framework on one of the patches with initiated dynamics and get T ms of dynamics forecasting for this patch;
 - 1) next, apply the framework on the first order neighboring patches of the initial patch (see Fig. 5.2.4(b)), getting T ms of dynamics forecasting for each new patch;
 - i) then, apply the framework on the i -th - order neighboring patches of the initial patch (see Fig. 5.2.4(c,d)), getting T ms of dynamics forecasting for each new patch; continuing until process all the patches.
3. Lastly, restore the initial image from patches, obtaining the forecast of dynamics on the full image.

The dynamics predicted inside the areas of patch overlapping is taken as average.




Fig. 5.2.4: Example of simulating various neighboring patches (10x10 pixels) orders from an initial patch in image of 50x50 pixels (without overlapping). (a) Initial patch, (b) first order neighboring patches, (c) second order neighboring patches and (d) third order neighboring patches.

This algorithm solves the problem of non-activated patches, but can not prevent erroneous propagation of dynamics through the patches generated by the APHYN-EP framework (if that is the case).

The results of reconstructing the ex vivo cardiac EP dynamics on the entire heart by the APHYN-EP framework trained on patches, are demonstrated in the sections 5.3.3.3 and 5.3.3.4.

5.2.4 Training settings

The training settings for the experiments in this chapter are similar to the ones described in section 3.3.2 (2D data). The training was performed using a horizon augmentation technique, where we started with a 20 ms horizon for the first 10 epochs and increased it gradually (by 20 ms) every 10 epochs, leading to more stable results.

We also tested a ConvNet (with 3 basic convolutional layers) as a simpler deep learning component (F_d) of the framework, for comparison.

After several trials, the hyper-parameters λ_0 and γ of the algorithm were set to 1 and 100, respectively, for better equilibrium.

Data augmentation

To force our framework to learn the propagation with respect to different fiber directions under the conditions of limited data, we use a simple data-augmentation technique, where we flip horizontally and vertically the observed dynamics (see Fig. 5.2.5), thus generating 4 data samples from one dataset with recordings of 2D time sequences.




Fig. 5.2.5: Data-augmentation example of one patch of true ex vivo recorded 2D optical mapping data of action potential. The first frame is the true patch, and the next frames are its flipped versions.

5.2.4.1 Isotropic diffusion

We estimated σ , τ_{in} , τ_{out} and τ_{close} as unknown parameters in F_p .

5.2.4.2 Anisotropic diffusion

We estimated $\sigma_{||}$, σ_{\perp} , α , τ_{in} , τ_{out} and τ_{close} as unknown parameters in F_p .

5.3 Experiments and Results

5.3.1 Isotropic diffusion

We observed that the APHYN-EP framework with F_p component based on an isotropic diffusion tensor was able to reproduce the features of action potential wave from 2D optical mapping data (see Fig. 5.3.1 and Fig. 5.3.2). Table 5.3.1 summarizes the quantitative results for our framework forecasting on train, validation and test data samples, in comparison to baseline methods trained on the same data.

One can notice that the Physical model and a simple Data-driven model surpass the APHYN-EP framework in terms of training and validation datasets; however, they have worse performance on a test dataset, which indicates either their possible over-fitting or disability to generalise to new conditions. It is also important to note that the Physical and Data-driven models have better performance because they managed to predict an average dynamics of AP in time for all data samples. While this is useful to predict the depolarisation/repolarisation phases in time for 2D ex vivo recordings (which vary a lot for every data sample), it is hardly applicable to correctly predict the heterogeneous spatial depolarisation wave propagation.

Nevertheless, APHYN-EP had the best results of forecasting for first 150 ms (see Table 5.3.1), which exclude the repolarisation phase, and retained persistent performance on the test data.

5.3.1.1 Generalisation ability of APHYN-EP: LV pacing

As shown in Figure 5.3.2, the framework keeps the capability to generalise to unseen conditions (e.g. LV pacing). It is also important to notice that results were obtained without any re-training the APHYNEP framework.

The absolute error was slightly larger than that on images of RV pacing (used for training); however, this error was still acceptable. The quantitative results are provided in the table 5.3.1.

5.3.2 Anisotropic diffusion

Figure 5.3.3 demonstrates the validation results obtained at different training epochs using the APHYN-EP framework, with the F_p component based on an anisotropic diffusion tensor. These results clearly show that the reproduced AP dynamics are very close to the ground truth dynamics.

Dataset	Method	Training data	Validation data	Out-of-domain test
RV pacing (300 ms)	APHYN-EP framework with ResNet ($\ F_d\ ^2$)	54.5 (92.6)	53.26 (94.2)	49.9 (134.6)
	APHYN-EP framework with ConvNet ($\ F_d\ ^2$)	60.3 (105.6)	58.9 (108.1)	42 (128.6)
	Physical model	39.2	37.5	106.4
	Data-driven model	33.89	31.7	93
RV pacing (first 150 ms)	APHYN-EP framework with ResNet ($\ F_d\ ^2$)	8.9 (110.85)	7.66 (109.9)	8.93 (135.8)
	APHYN-EP framework with ConvNet ($\ F_d\ ^2$)	8.26 (101.16)	6.93 (107.5)	7.47 (124.4)
	Physical model	30.3	33	32.5
	Data-driven model	13.25	12.6	14.9

Tab. 5.3.1: Mean-squared error, MSE ($\times 10^{-3}$) of the normalised transmembrane potential (adi-dimensional) forecasting (forecasting horizons of 300 ms and 150 ms). Baseline models: the Physical model (2.1) and a fully data-driven model (EP-Net 2.0 [Kashanova, 2021]) trained on the same dataset as APHYN-EP. Out-of-domain tests: LV pacing.







Fig. 5.3.1: APHYN-EP predicted dynamics for the transmembrane potential diffusion, RV pacing (data used for training). The frames show results obtained over a 30ms period of forecast.



(a)



(b)



(c)

Fig. 5.3.2: APHYN-EP predicted dynamics for the transmembrane potential diffusion, IV pacing (test data). The frames show a period of 30 ms of forecast obtained without re-training the APHYN-EP framework.

However, it is also noticeable that even within the 3 examples of dynamics presented in Fig. 5.3.3, we have three different types of anisotropy forcing drastic changes in the values of the physical parameters that control the anisotropic diffusion ($\sigma_{||}$, σ_{\perp} and α) each epoch.

Taking into account the heterogeneity of the fiber directions within one heart (i.e., the fibers' rotation within the myocardial wall), the usage of data-augmentation technique described in section 5.2.4, and the potential application of the trained framework on new unseen dynamics, we decided to discard the usage of APHYN-EP with F_p component based on an anisotropic diffusion tensor.

Nevertheless, the anisotropic F_p component could be used for tuning in the pre-trained framework, when we need more sophisticated results.

5.3.3 Generalisation: prediction extension to the whole heart

5.3.3.1 Extension of prediction via convolutions: RV pacing

Figure 5.3.4 demonstrates the results of application of APHYN-EP framework on whole heart image, with pacing catheter placed at the RV location.

It is important to notice that in order to generate these results we used the APHYN-EP framework trained on patches (from section 5.3.1), without any re-training or final tuning, and applied it directly on whole heart 2D image (with size of 100x100 pixels).

One can observe that APHYN-EP framework predicts the true dynamics (similar with the dynamics used for training) with good precision for several time steps (about 33 ms). Unfortunately, predicted velocity of wave propagation is gradually decreasing in time (see Fig. 5.3.4(b,c)). This occurred as a result of potential F_d component overfitting to the specific velocity of spatial AP wave propagation for a given patch size (10x10 pixels), and encountered some difficulties to extend this dynamics to larger images for larger time horizon. The same effect was observed by applying the trained APHYN-EP framework on images of smaller size (but still bigger than training patches), for instance 30x30 pixels.

5.3.3.2 Extension of prediction via convolutions: LV pacing

Figure 5.3.5 demonstrates the results of application of APHYN-EP framework on whole heart image, with the heart pacing initiated at the LV location.

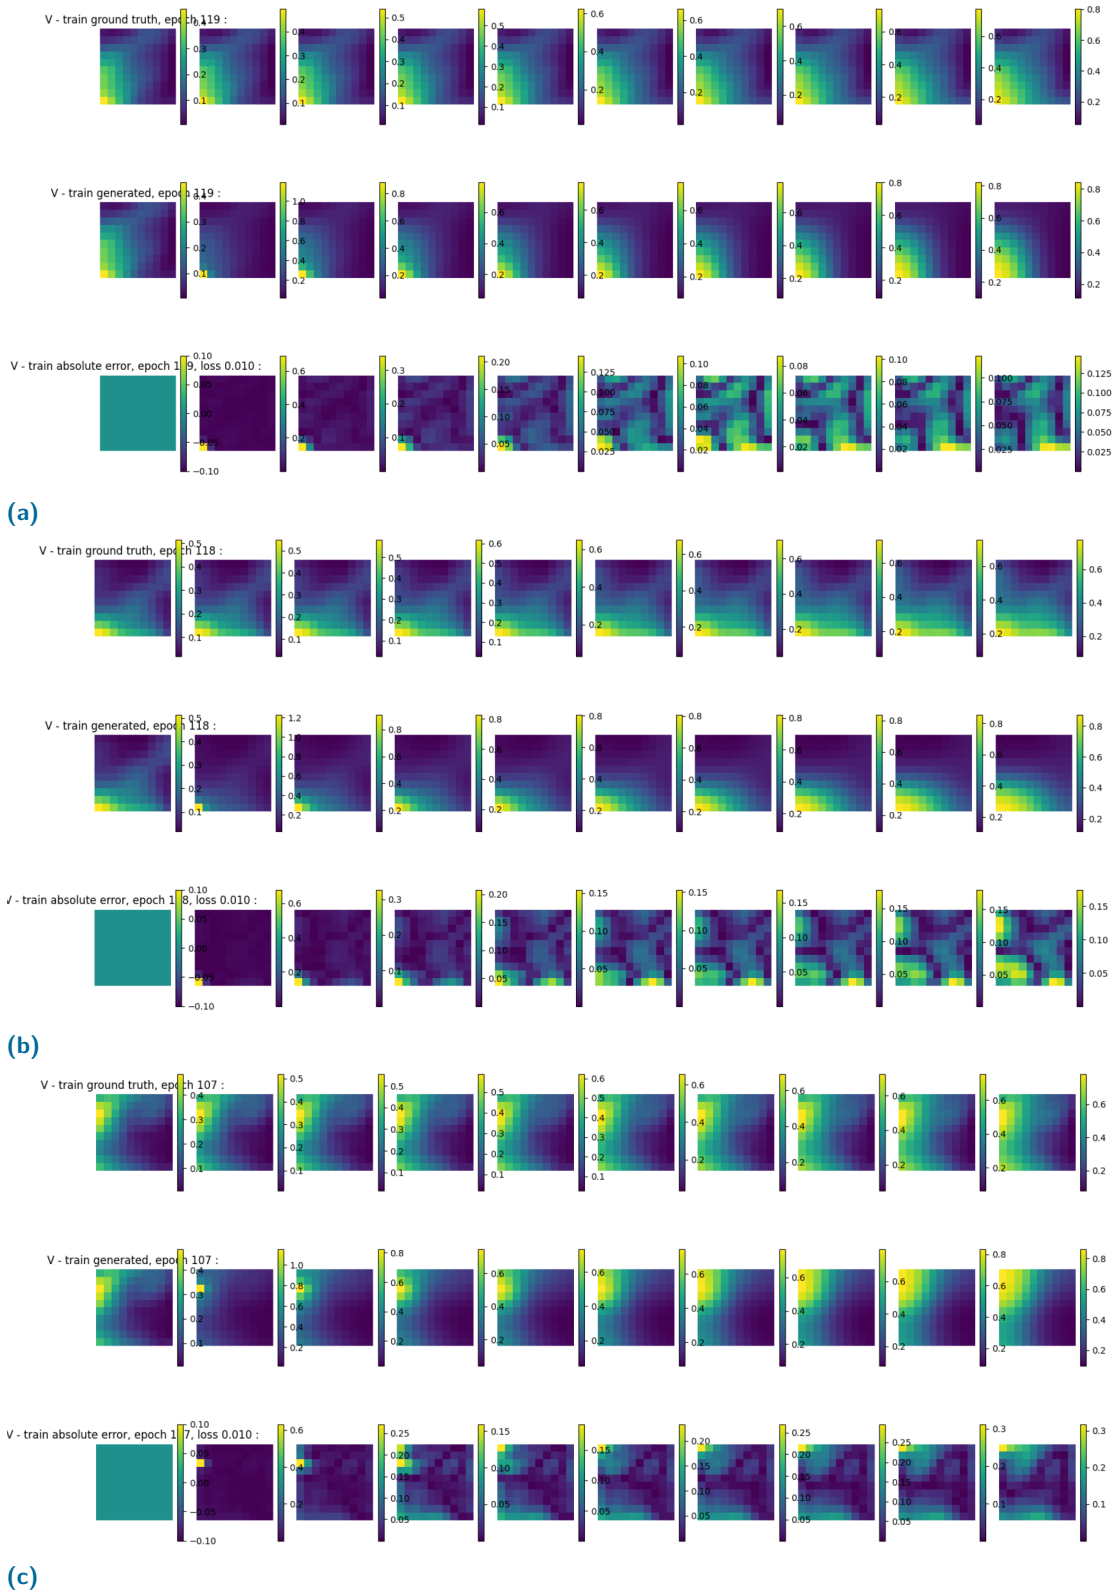


Fig. 5.3.3: Snapshots of APHYN-EP predicted dynamics for the transmembrane potential diffusion (validation dataset) at different training epochs. The frames show a period of 10 ms of forecast.




Fig. 5.3.4: APHYN-EP predicted dynamics for the transmembrane potential diffusion on the entire heart with RV pacing. The frames show a period of 107 ms of forecast obtained without re-training the APHYN-EP framework.

It is important to notice that in order to generate these results we used the APHYN-EP framework trained on patches (from section 5.3.1), without any re-training or final tuning, and applied it directly on whole heart 2D image (with size of 100x100 pixels).

One can observe that APHYN-EP framework can predict with sufficient precision the true dynamics for first 30 ms, but encounters some difficulties to further produce a proper forecast.

Nevertheless, convolution-based F_d component enables a training (or a fine-tuning) of the APHYN-EP framework on images of any size without any framework architectural changes, that could be used to improve the obtained results.




Fig. 5.3.5: APHYN-EP predicted dynamics for the transmembrane potential diffusion on the entire heart with LV pacing. The frames show a period of 107 ms of forecast obtained without re-training the APHYN-EP framework.

5.3.3.3 Reconstruction of dynamics from patches: RV pacing

Figures 5.3.6-5.3.7 present the results of the translation of real ex vivo recorded cardiac EP dynamics from patches to the whole heart, with pacing at the RV location.

It is important to notice that in order to generate these results we used the APHYN-EP framework trained on patches (from section 5.3.1), without any re-training or final tuning, using only the "Dynamics based algorithm for patch-by-patch reconstruction" described in section 5.2.3.2.

One can observe a good precision of the predicted dynamics on the lower part of the heart, where the patches producing dynamics are very close to the patches from the

training dataset. One can also observe a depolarisation delay at the upper part the heart (i.e. towards its base), where the patches producing dynamics were not seen during the training. However, the framework did not encounter any problems when completing the AP wave propagation in the right form. Furthermore, it is important to notice that we cut the diffusion propagated to the background of the heart.




Fig. 5.3.6: APHYN-EP predicted dynamics for the transmembrane potential diffusion extrapolated to the entire heart with RV pacing. The frames show a period of 107 ms of forecast obtained without re-training the APHYN-EP framework.

5.3.3.4 Reconstruction of dynamics from patches: LV pacing

Figure 5.3.8 present the results of translation of real ex vivo recorded cardiac EP dynamics from patches to the whole heart, with the pacing catheter placed on the LV.

It is important to notice that to generate these results we used the APHYN-EP framework trained on patches (from section 5.3.1), without any re-training or final tuning, using






Fig. 5.3.7: APHYN-EP predicted dynamics for the transmembrane potential diffusion extrapolated to the entire heart with RV pacing. The frames show a period from 111 ms to 181 ms of forecast obtained without re-training the APHYN-EP framework.

only the "Dynamics based algorithm for patch-by-patch reconstruction" described in the section 5.2.3.2.

Unfortunately, the results of dynamics reconstruction were worse than those obtained for the heart paced from RV. One can observe that even having a good generalisation results of local dynamics (see section 5.3.1.1), the framework learned at "macro" level the specific directions of diffusion propagation (faster propagation to the top-right inherent to the RV pacing), which is different from the one in the new data (see the faster propagation to the top-left).



(a)



(b)

Fig. 5.3.8: APHYN-EP predicted dynamics for the transmembrane potential diffusion extrapolated to the entire heart with LV pacing. The frames show a period of 70 ms of forecast obtained without re-training the APHYN-EP framework.

5.4 Discussion and Conclusion

In this chapter, we demonstrated that the learning framework (presented in chapter 3) is able to learn the real 2D dynamics recorded ex vivo. Overall, our results indicate that automated learning and translation of cardiac EP dynamics is feasible and has great potential for such biophysical applications.

In sections 5.3.1 and 5.3.2, we visually and quantitatively confirmed that the APHYN-EP framework can reproduce the complex 2D dynamics of AP wave propagation, even if this contains anisotropic properties. Additionally, we demonstrated that the framework can locally (i.e., on small patches) reproduce the new unseen 2D dynamics (see section 5.3.1.1).

Alongside, we also demonstrated examples of translation of real 2D cardiac EP dynamics from patches to the whole heart, presenting promising results (see section 5.3.3).

However, we acknowledge several limitations in the presented results.

- Due to high heterogeneity of the training database, it is inefficient to use an anisotropic diffusion tensor inside the F_d component of the APHYN-EP framework (see section 5.3.2), even if this has a great potential for realistic modelling of clinical cases, where the patient-specific anisotropy is unknown.
- Due to limited training data, the F_d component learned the specific velocity of spatial AP wave propagation for a given patch size and encountered some difficulties to extend this dynamics to larger images (see sections 5.3.3.1 and 5.3.3.2). Specific adaptation of the integration schemes could help, for instance leveraging different spatial discretisation steps.
- The framework has some inaccuracies in predicting the AP wave propagation velocity on data representing the heart zones not used during training (see section 5.3.3.3). However, it has been previously proved in [Relan, 2011b] for similar optical mapping recordings that the physical parameters could significantly vary from one heart zone to another one (as per AHA partition) as can be seen in Fig. 5.4.1.
- The framework's learned locally (on small patches) dynamics contains the information of AP wave diffusion specific for one pacing location, which can lead to errors during "to whole heart" prediction extension (see section 5.3.3.4).

The majority of these limitations can be improved by adding additional information in the framework (such as the parameter maps, etc.), by completing the training dataset with patches representing the pacing from other locations, or by fine-tuning the pre-trained framework on new data. These potential refinements could be addressed in future work.

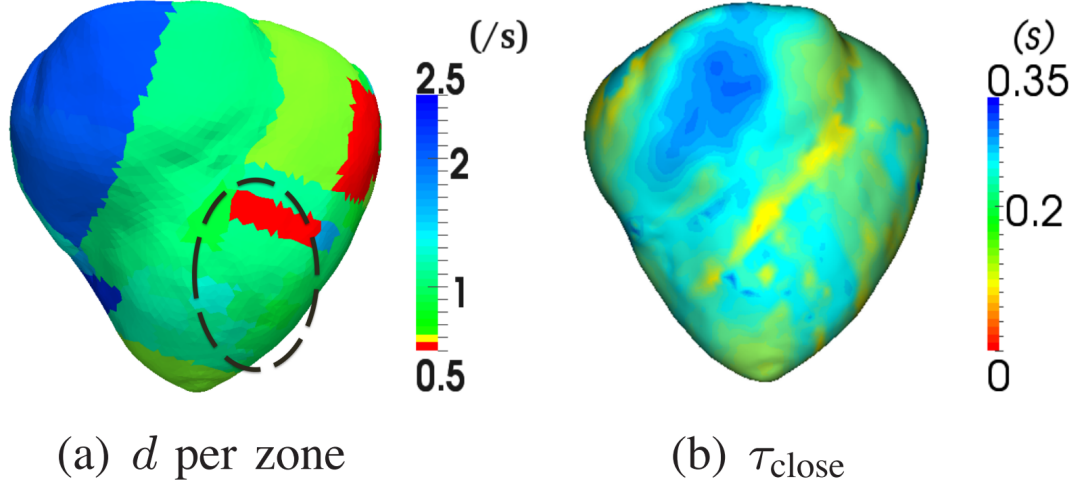


Fig. 5.4.1: Parameter maps for a healthy heart with LV pacing location. Estimated d (σ in our F_p model) and τ_{close} values per zone. From [Relan, 2011b]

Conclusion

Contents

6.1	Contributions Summary	93
6.1.1	Fully data-driven framework	93
6.1.2	Physics-based deep learning framework	94
6.2	Publications	96
6.3	Discussion and Perspectives	96

In this thesis, we explored the potential of data-driven frameworks to improve cardiac electrophysiology modeling approaches and model personalisation.

In the following sections we summarise the main contributions of our novel research work, discuss the main drawbacks of our frameworks, and propose some potential perspectives for future work.

6.1 Contributions Summary

6.1.1 Fully data-driven framework

In chapter 2 we presented a newly improved version of a fully data-driven framework EP-Net 2.0, which is able to automatically learn and forecast the 2D cardiac EP dynamics in the presence of unexcitable myocardial scars integrated into a virtually-defined slab of cardiac tissue.

- In section 2.2, we explored in detail the framework's architecture and training stages.
- In section 2.3.1, we presented the generation of in silico data that was specifically designed for the training and evaluations of our framework.
- In section 2.4.1, we evaluated the framework's ability to learn and forecast simple cardiac EP dynamics in conditions similar to the training ones (i.e., scars of rectangular shape and one stimulation onset in a cardiac tissue slab).

- In sections 2.4.2 and 2.4.3, we evaluated the framework’s ability to generalise to diverse unseen complex conditions such as: scars of various shapes; multiple onsets; and, various conduction velocities in a cardiac tissue slab.
- In section 2.5, we discussed the main drawbacks of our framework and proposed possible improvements.

6.1.2 Physics-based deep learning framework

In chapter 3 we presented a novel physics-based deep learning framework (APHYN-EP) that can automatically learn and personalise cardiac EP dynamics from data of different physical complexities. With this respect, in chapters 3, 4 and 5, we explored and evaluated various applications of this framework using synthetic data as well as real data.

- In section 3.2, we introduced in detail the two component architecture of the framework and specific training configurations.
- In section 3.3.1, we presented the generation of silico data specifically created for the training and evaluations of our framework using two biophysically detailed ionic models.
- In section 1.1.3.2, we presented an ex vivo database obtained from the optical fluorescence signals of action potential (recorded in explanted pig hearts), which was further used for the framework’s applications on real data.

In silico data:

- In section 3.4, we evaluated the framework’s ability to learn and forecast the cardiac EP dynamics from synthetic data of different complexities, as well as to both generalise to new conditions and to forecast cardiac EP dynamics outside of its training domain. Additionally:
 - In section 3.4.1, we explored the framework performance depending on the involvement of its physical component using 0D synthetic data.
 - In section 3.4.2, using 2D synthetic data, we demonstrated the contribution of each of the framework’s components (physical and data-driven, respectively).
 - In section 3.4.2, we also provided a qualitative performance comparison between our APHYN-EP framework and baseline methods.

Ex vivo 0D data:

- In section 4.3, we evaluated the framework's ability to learn the full cardiac cycle of action potential from real 0D ex vivo recorded EP data. In addition, we provided a qualitative performance comparison between the APHYN-EP framework and baseline methods.
- In sections 4.3.1.1 and 4.3.2.1, we demonstrated that the APHYN-EP framework can identify the key physical parameters for different anatomical zones of porcine hearts;
- In section 4.2.3, we presented a method of faster personalisation of cardiac EP model, using a pre-trained APHYN-EP framework.
- In sections 4.3.1.2 and 4.3.2.2, we evaluated our fast personalisation method on data samples with physiological changes in tissue properties (e.g. changes in electrical properties), and pacing at various stimulation rates.

Ex vivo 2D data:

- In section 5.3.1, we evaluated the framework's ability to learn the real 2D ex vivo recorded EP data.
- In the section 5.2.2, we developed and evaluated the various types of diffusion tensors to model a depolarisation wave propagation in physical component of the framework.
- In section 5.3.1, we provided qualitative performance comparison of APHYN-EP framework and baseline methods.
- In section 5.2.3, we proposed a method of cardiac EP model deep personalisation on data of 2D ex vivo dynamics, training APHYN-EP framework on small data patches.
- In section 5.3.1.1, we evaluated the proposed method on test patches selected from 2D ex vivo data.
- In sections 5.2.3.2 and 5.3.3, we introduced and evaluated a method of learned dynamics extension from small patches to the entire heart.
- In sections 4.4 and 5.4, we discussed the main drawbacks of our framework and proposed potential improvements.

6.2 Publications

The contributions described above led to the following peer-reviewed publications:

Journal Articles

- [[Kashtanova, 2023](#)] **Kashtanova, V.**, Pop, M., Ayed, I., Gallinari, P., Sermesant, M. Simultaneous Data Assimilation and Model Correction Using Differentiable Physics and Deep Learning. *In preparation*.

Conference Papers

- [[Kashtanova, 2021](#)] **Kashtanova, V.**, Ayed, I., Cedilnik, N., Gallinari, P., Sermesant, M. EP-Net 2.0: Out-of-Domain Generalisation for Deep Learning Models of Cardiac Electrophysiology. In *Int. Conf. FIMH*, volume 12738 of *Lecture Notes in Computer Science*, pages 482–492. Springer International Publishing, 2021.
- [[Kashtanova, 2022a](#)] **Kashtanova, V.**, Ayed, I., Arrieula, A., Potse, M., Gallinari, P., Sermesant, M. Deep Learning for Model Correction in Cardiac Electrophysiological Imaging. In *Int. Conf. MIDL*, volume 172 of *Proceedings of Machine Learning Research*, pages 665–675. PMLR, 2022.
- [[Kashtanova, 2022b](#)] **Kashtanova, V.**, Pop, M., Ayed, I., Gallinari, P., Sermesant, M. APHYN-EP: Physics-Based Deep Learning Framework to Learn and Forecast Cardiac Electrophysiology Dynamics. In *Int. Conf. STACOM*, volume 13593 of *Lecture Notes in Computer Science*, pages 190–199. Springer, Cham., 2022.

6.3 Discussion and Perspectives

In Chapter 3 we presented a novel physics-based Deep Learning framework which can learn the cardiac EP dynamics from data of different complexity. This framework performance was validated using *in silico* data (see Chapter 3) as well as *ex vivo* data (see Chapters 4 and 5), demonstrating promising results.

However, we suggest several opportunities for further improvement of the framework.

We acknowledge that the performance of the framework depends on many hyper-parameters, including the choice of framework’s components and training strategy.

For example, we observed that the usage of simple neural networks for the DL component of our APHYN-EP framework, is sometimes able to improve the results (see Tables 4.3.1 and 5.3.1). However, the choice of such models depends entirely on the learning data configurations such as data dimension, presence/absence of noise, etc. In this manuscript we have focused mostly on results generated by the framework with the DL part represented by a ResNet network, that is because of its robustness and capability of resting stable during different simulations. Additionally, thanks to its residual connections, ResNet can accurately reproduce complex cardiac EP dynamics [Ayed, 2019b; Kashtanova, 2021].

Another encountered difficulty in the framework training was linked to the lack of training data or its insufficient heterogeneity, especially for real 2D experiments (see section 5.4). That leads to the framework overfitting and poorer generalisation of dynamics in new conditions (such as the change of pacing location etc).

However, there is no straightforward solution to this problem. Each real heart is unique, having a specific morphology and containing tissular components with different physical properties. This makes the physical parameters (e.g. tissue conductivity) to significantly vary from one heart zone/segment to another [Relan, 2011b]. Additionally, the optical signals recorded in real hearts are usually sparse and noisy, and contain anatomical components and structures (e.g. blood vessels, fat) that do not propagate the depolarisation wave.

There are several strategies to improve these limitations, which will be addressed in the future work:

- Giving additional information to the framework (e.g. physical parameter maps).
- Final tuning of the framework, trained on poor quality real data, on new data, or tuning of only one of its component (with restrained number of parameters).
- Usage of the framework, pre-trained on good quality synthetic data, and its re-training on poor quality real data.

Our future work will also focus on the adaptation of APHYN-EP framework for applications on more advanced types of real data, such as 3D catheter-based in vivo EP recordings or various types of non-invasive data (ECG, BSPM, CT and MRI imaging). Accomplishing that kind of adaptability, would enable the usage of our framework for fast EP modelling in the "Ventricular tachycardia simulation from non-invasive data" pipeline (see Fig. 6.3.1), which could improve cardiac intervention planning.

The framework's adaptation to 3D real data would demand several changes in the framework architecture, such as the implementation of a more rapid method for physical

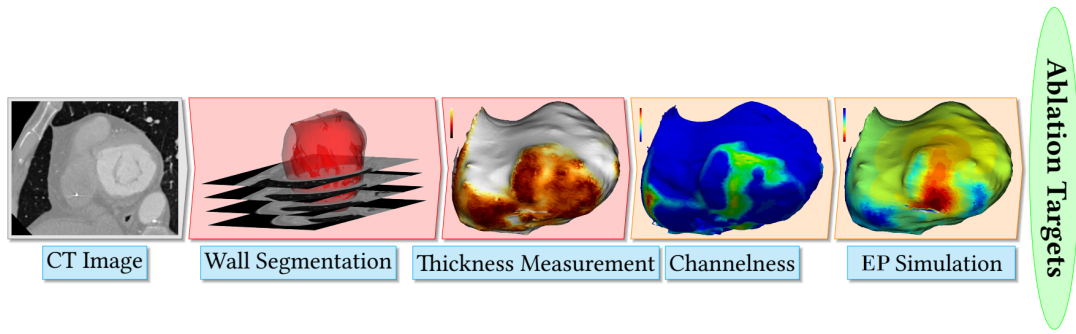


Fig. 6.3.1: Ventricular Tachycardia (VT) modelling pipeline using non-invasive CT data to define the ablation targets for VT Ablation Planning. From [Cedilnik, 2018].

model computation on 3D grid, as well as adding 3D convolutions and an interpolation operator to reduce the impact of data sparsity. As for the cardiac EP personalisation from non-invasive data, this could require an additional data-driven correction of an observation operator, in addition to the physical model correction.

For instance, in case of translating the AP dynamics from the heart surface (Fig. 6.3.2(a)) to the body surface dynamics (Fig. 6.3.2(b)) which could be captured via non-invasive body surface potential mapping (BSPM) methods, the Current Dipole Formulation can be employed [Giffard-Roisin, 2017b]. However, this formulation is not very accurate and has discrepancies with measured data.




Fig. 6.3.2: (a) Myocardial mesh with measured pacing locations (red), (b) 64 BSPM electrodes (black), myocardial mesh (pink). From [Giffard-Roisin, 2017b].

Another possible future research direction may focus on exploring the data-driven estimation of the physical model (or of the observation operator) equations using various data-driven symbolic regression approaches [Rudy, 2017; Long, 2018; Long, 2019; Petersen, 2021; Holt, 2023], which we did not investigate in this thesis.

To sum-up, there are many different ways to combine the strengths of machine learning and scientific computing for electrophysiology modelling. We envision that the integration of such AI-based methodologies will be critical to the rapid translation of in silico computational techniques into clinical platforms, where they could improve the currently invasive diagnostic procedures of scar-related arrhythmias and the success of cardiac interventions aiming to treat these potentially lethal heart conditions.

Bibliography

- [Aliev, 1996] Rubin R. Aliev and Alexander V. Panfilov. “A simple two-variable model of cardiac excitation”. In: *Chaos, Solitons & Fractals* 7.3 (1996), pp. 293–301 (cit. on pp. [10](#), [31](#)).
- [Aliot, 2009] Etienne M. Aliot, William G. Stevenson, Jesus Ma Almendral-Garrote, Frank Bogun, C. Hugh Calkins, Etienne Delacretaz, Paolo Della Bella, Gerhard Hindricks, Pierre Jaïs, Mark E. Josephson, Josef Kautzner, G. Neal Kay, Karl-Heinz Kuck, Bruce B. Lerman, Francis Marchlinski, Vivek Reddy, Martin-Jan Schalij, Richard Schilling, Kyoko Soejima, and David Wilber. “EHRA/HRS Expert Consensus on Catheter Ablation of Ventricular Arrhythmias”. In: *EP Europace* 11.6 (2009), pp. 771–817 (cit. on p. [6](#)).
- [Alvarez, 2013] Mauricio A Alvarez, David Luengo, and Neil D Lawrence. “Linear latent force models using Gaussian processes”. In: *IEEE PAMI* 35.11 (2013), pp. 2693–2705 (cit. on pp. [11](#), [17](#)).
- [Ayed, 2019a] Ibrahim Ayed, Emmanuel de Bézenac, Arthur Pajot, Julien Brajard, and Patrick Gallinari. “Learning dynamical systems from partial observations”. In: *arXiv preprint:1902.11136* (2019) (cit. on pp. [11](#), [17](#)).
- [Ayed, 2019b] Ibrahim Ayed, Nicolas Cedilnik, Patrick Gallinari, and Maxime Sermesant. “EP-Net: Learning Cardiac Electrophysiology Models for Physiology-based Constraints in Data-Driven Predictions”. In: *Int. Conf. FIMH*. Springer. 2019, pp. 55–63 (cit. on pp. [11](#), [17](#), [19](#), [33](#), [97](#)).
- [Bengio, 2015] Samy Bengio, Oriol Vinyals, Navdeep Jaitly, and Noam Shazeer. “Scheduled sampling for sequence prediction with recurrent neural networks”. In: *arXiv preprint:1506.03099* (2015) (cit. on p. [21](#)).
- [Camara, 2011] O. Camara, M. Sermesant, P. Lamata, L. Wang, M. Pop, J. Relan, M. De Craene, H. Delingette, H. Liu, S. Niederer, A. Pashaei, G. Plank, D. Romero, R. Sebastian, K.C.L. Wong, H. Zhang, N. Ayache, A.F. Frangi, P. Shi, N.P. Smith, and G.A. Wright. “Inter-model consistency and complementarity: Learning from ex-vivo imaging and electrophysiological data towards an integrated understanding of cardiac physiology”. In: *Progress in Biophysics and Molecular Biology* 107.1 (2011). Experimental and Computational Model Interactions in Bio-Research: State of the Art, pp. 122–133 (cit. on p. [9](#)).

- [Campos, 2016] J.O. Campos, R.S. Oliveira, R.W. dos Santos, and B.M. Rocha. “Lattice Boltzmann method for parallel simulations of cardiac electrophysiology using GPUs”. In: *Journal of Computational and Applied Mathematics* 295 (2016). VIII Pan-American Workshop in Applied and Computational Mathematics, pp. 70–82 (cit. on p. 71).
- [Cedilnik, 2018] Nicolas Cedilnik, Josselin Duchateau, Rémi Dubois, Frédéric Sacher, Pierre Jaïs, Hubert Cochet, and Maxime Sermesant. “Fast Personalized Electrophysiological Models from CT Images for Ventricular Tachycardia Ablation Planning”. In: *EP-Europace* 20 (2018) (cit. on pp. 27, 71, 98).
- [Cedilnik, 2020] Nicolas Cedilnik. “Image-based Personalised Models of Cardiac Electrophysiology for Ventricular Tachycardia Therapy Planning”. PhD thesis. 2020 (cit. on pp. 4, 52).
- [Chen, 2018] Ricky T. Q. Chen, Yulia Rubanova, Jesse Bettencourt, and David Duvenaud. “Neural Ordinary Differential Equations”. In: *Adv Neural Inform Process Syst* (2018) (cit. on pp. 11, 17, 33).
- [Chen, 2021] Ricky T. Q. Chen, Brandon Amos, and Maximilian Nickel. “Learning Neural Event Functions for Ordinary Differential Equations”. In: *ICRL* (2021) (cit. on p. 33).
- [Court, 2021] Sebastien Court and Karl Kunisch. “Design of the monodomain model by artificial neural networks”. In: *arXiv preprint:2107.03136* (2021) (cit. on p. 12).
- [Crutchfield, 1987] James P Crutchfield and BS McNamara. “Equations of motion from a data series”. In: *Complex systems* 1.417-452 (1987), p. 121 (cit. on pp. 11, 17).
- [Desrues, 2021] Gaëtan Desrues, Delphine Feuerstein, Thierry Legay, Serge Cazeau, and Maxime Sermesant. “Personal-by-design: a 3D Electromechanical Model of the Heart Tailored for Personalisation”. In: *FIMH 2021 - 11th International Conference on Functional Imaging and Modeling of the Heart*. Stanford, CA, United States, 2021 (cit. on p. 71).
- [Durrer, 1970] Dirk Durrer, R. Th. van Dam, G. E. Freud, Michiel J. Janse, Frits L. Meijler, and Robert C. Arzbaecher. “Total Excitation of the Isolated Human Heart”. In: *Circulation* 41 (1970), pp. 899–912 (cit. on p. 3).
- [Fenton, 1998] Flavio Fenton and Alain Karma. “Vortex dynamics in three-dimensional continuous myocardium with fiber rotation: Filament instability and fibrillation”. In: *Chaos* 8.1 (1998), pp. 20–47 (cit. on p. 74).
- [FitzHugh, 1961] Richard FitzHugh. “Impulses and Physiological States in Theoretical Models of Nerve Membrane”. In: *Biophysical Journal* 1.6 (1961), pp. 445–466 (cit. on pp. 10, 31).
- [Fresca, 2020] Stefania Fresca, Andrea Manzoni, Luca Dedè, and Alfio Quarteroni. “Deep learning-based reduced order models in cardiac electrophysiology”. In: *PLOS ONE* 15.10 (2020), e0239416 (cit. on p. 17).

- [Fresca, 2021] Stefania Fresca, Andrea Manzoni, Luca Dedè, and Alfio Quarteroni. “POD-Enhanced Deep Learning-Based Reduced Order Models for the Real-Time Simulation of Cardiac Electrophysiology in the Left Atrium”. In: *Front. Physiol.* 12 (2021) (cit. on p. 12).
- [Ghanbari, 2014] Hamid Ghanbari, Kazim Baser, Miki Yokokawa, William Stevenson, Paolo Della Bella, Pasquale Vergara, Thomas Deneke, Karl-Heinz Kuck, Hans Kottkamp, She Fei, Fred Morady, and Frank Bogun. “Noninducibility in postinfarction ventricular tachycardia as an end point for ventricular tachycardia ablation and its effects on outcomes: a meta-analysis”. In: *Circulation: Arrhythmia and Electrophysiology* 7.4 (2014), pp. 677–683 (cit. on p. 5).
- [Giffard-Roisin, 2017a] Sophie Giffard-Roisin, Thomas Jackson, Lauren Fovargue, Jack Lee, Herve Delingette, Reza Razavi, Nicholas Ayache, and Maxime Serme-sant. “Non-Invasive Personalization of a Cardiac Electrophysiology Model From Body Surface Potential Mapping”. In: *IEEE Trans. Biomed. Eng.* 64.9 (2017), pp. 2206–2218 (cit. on p. 30).
- [Giffard-Roisin, 2017b] Sophie Giffard-Roisin, Thomas Jackson, Lauren Fovargue, Jack Lee, Hervé Delingette, Reza Razavi, Nicholas Ayache, and Maxime Serme-sant. “Non-Invasive Personalisation of a Cardiac Electrophysiology Model from Body Surface Potential Mapping”. In: *IEEE Transactions on Biomedical Engineering* 64.9 (2017), pp. 2206–2218 (cit. on p. 98).
- [He, 2016] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. “Deep residual learning for image recognition”. In: *IEEE conf. CVPR.* 2016, pp. 770–778 (cit. on pp. 20, 33, 37).
- [Herrero Martin, 2022] Clara Herrero Martin, Alon Oved, Rasheda A. Chowdhury, Elisabeth Ullmann, Nicholas S. Peters, Anil A. Bharath, and Marta Varela. “EP-PINNs: Cardiac Electrophysiology Characterisation Using Physics-Informed Neural Networks”. In: *Front. Cardiovasc. Med.* 8 (2022) (cit. on p. 12).
- [Holt, 2023] Samuel Holt, Zhaozhi Qian, and Mihaela van der Schaar. “Deep Generative Symbolic Regression”. In: *Int. Conf. ICLR 2023.* 2023 (cit. on p. 98).
- [Jiang, 2019] Chiyu ”Max” Jiang, Karthik Kashinath, Prabhat, and Philip Marcus. “Enforcing Physical Constraints in CNNs through Differentiable PDE Layer”. In: *ICLR 2020 Workshop on Integration of Deep Neural Models and Differential Equations.* 2019 (cit. on p. 12).
- [Kabanikhin, 2011] Sergey I. Kabanikhin. *Inverse and Ill-posed Problems: Theory and Applications.* Berlin, Boston: De Gruyter, 2011 (cit. on p. 34).
- [Karniadakis, 2021] George Em Karniadakis, Ioannis G. Kevrekidis, Lu Lu, Paris Perdikaris, Sifan Wang, and Liu Yang. “Physics-informed machine learning”. In: *Nature Reviews Physics* 3 (2021), pp. 422–440 (cit. on p. 11).
- [Karoui, 2021] Amel Karoui, Mostafa Bendahmane, and Nejib Zemzemi. “Cardiac activation maps reconstruction: a comparative study between data-driven and physics-based methods”. In: *Front. Physiol.* 12 (2021), p. 1265 (cit. on p. 30).

- [Kashtanova, 2021] Victoriya Kashtanova, Ibrahim Ayed, Nicolas Cedilnik, Patrick Gallinari, and Maxime Sermesant. “EP-Net 2.0: Out-of-Domain Generalisation for Deep Learning Models of Cardiac Electrophysiology”. In: *Int. Conf. FIMH*. Vol. 12738. Lecture Notes in Computer Science. Springer International Publishing, 2021, pp. 482–492 (cit. on pp. [11](#), [15](#), [17](#), [33](#), [40–42](#), [57](#), [59](#), [81](#), [96](#), [97](#)).
- [Kashtanova, 2022a] Victoriya Kashtanova, Ibrahim Ayed, Andony Arrieula, Mark Potse, Patrick Gallinari, and Maxime Sermesant. “Deep Learning for Model Correction in Cardiac Electrophysiological Imaging”. In: *Proceedings of The 5th International Conference on Medical Imaging with Deep Learning*. Vol. 172. Proceedings of Machine Learning Research. PMLR, 2022, pp. 665–675 (cit. on pp. [30](#), [96](#)).
- [Kashtanova, 2022b] Victoriya Kashtanova, Mihaela Pop, Ibrahim Ayed, Patrick Gallinari, and Maxime Sermesant. “APHYN-EP: Physics-Based Deep Learning Framework to Learn and Forecast Cardiac Electrophysiology Dynamics”. In: *Statistical Atlases and Computational Models of the Heart. Regular and CMRxMotion Challenge Papers*. Cham: Springer Nature Switzerland, 2022, pp. 190–199 (cit. on pp. [30](#), [47](#), [96](#)).
- [Kashtanova, 2023] Victoriya Kashtanova, Mihaela Pop, Ibrahim Ayed, Patrick Gallinari, and Maxime Sermesant. “Simultaneous Data Assimilation and Model Correction Using Differentiable Physics and Deep Learning”. preprint submitted to a journal. 2023 (cit. on pp. [70](#), [96](#)).
- [Keener, 1991] James P. Keener. “The effects of discrete gap junction coupling on propagation in myocardium”. In: *Journal of Theoretical Biology* 148.1 (1991), pp. 49–82 (cit. on p. [10](#)).
- [Kingma, 2014] Diederik P Kingma and Jimmy Ba. “Adam: A method for stochastic optimization”. In: *arXiv preprint:1412.6980* (2014) (cit. on pp. [20](#), [38](#)).
- [Krause, 2012] Dorian Krause, Mark Potse, Thomas Dickopf, Rolf Krause, Angelo Auricchio, and Frits W. Prinzen. “Hybrid Parallelization of a Large-Scale Heart Model”. In: *Facing the Multicore-Challenge II*. Vol. 7174. Lecture Notes in Computer Science. Springer, 2012, pp. 120–132 (cit. on p. [35](#)).
- [Long, 2018] Zichao Long, Yiping Lu, Xianzhong Ma, and Bin Dong. “PDE-net: Learning PDEs from data”. In: *Int. Conf. ICML*. PMLR. 2018, pp. 3208–3216 (cit. on pp. [11](#), [17](#), [98](#)).
- [Long, 2019] Zichao Long, Yiping Lu, and Bin Dong. “PDE-Net 2.0: Learning PDEs from data with a numeric-symbolic hybrid deep network”. In: *J. Comput. Phys.* 399 (2019), p. 108925 (cit. on pp. [11](#), [17](#), [98](#)).
- [Mansi, 2020] Tommaso Mansi, Tiziano Passerini, and Dorin Comaniciu. *Artificial Intelligence for Computational Modeling of the Heart*. Elsevier, 2020 (cit. on pp. [17](#), [30](#)).
- [Mitchell, 2003] Colleen C Mitchell and David G Schaeffer. “A two-current model for the dynamics of cardiac membrane”. In: *Bull. Math. Biol* 65.5 (2003), pp. 767–793 (cit. on pp. [10](#), [17](#), [19](#), [21](#), [31](#), [32](#), [35–38](#), [51–53](#), [72](#)).

- [Nagumo, 1962] J. Nagumo, S. Arimoto, and S. Yoshizawa. “An Active Pulse Transmission Line Simulating Nerve Axon”. In: *Proceedings of the IRE* 50.10 (1962), pp. 2061–2070 (cit. on pp. 10, 31).
- [Nash, 2004] Martyn P. Nash and Alexander V. Panfilov. “Electromechanical model of excitable tissue to study reentrant cardiac arrhythmias”. In: *Prog. Biophys. Mol. Biol.* 85.2 (2004), pp. 501–522 (cit. on pp. 10, 31).
- [Nelles, 2001] Oliver Nelles. *Nonlinear System Identification*. Springer Berlin Heidelberg, 2001 (cit. on pp. 11, 17).
- [Nezlobinsky, 2021] T. Nezlobinsky. “Software for cardiac modelling and patch clamp experiments, and its application to study conduction in the presence of fibrosis”. PhD thesis. Ghent University, 2021 (cit. on p. 35).
- [Nielsen, 1991] P M Nielsen, I J Le Grice, B H Smaill, and P J Hunter. “Mathematical model of geometry and fibrous structure of the heart”. In: *Am J Physiol* 260.4 Pt 2 (1991), H1365–78 (cit. on p. 74).
- [Paszke, 2019] Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, et al. “Pytorch: An imperative style, high-performance deep learning library”. In: *Adv. Neural Inf. Process. Syst.* 32 (2019) (cit. on p. 34).
- [Pernod, 2011] Erik Pernod, Maxime Sermesant, Ender Konukoglu, Jatin Relan, Hervé Delingette, and Nicholas Ayache. “A Multi-Front Eikonal Model of Cardiac Electrophysiology for Interactive Simulation of Radio-Frequency Ablation”. In: *Computers and Graphics* 35 (2011), pp. 431–440 (cit. on p. 10).
- [Petersen, 2021] Brenden K Petersen, Mikel Landajuela Larma, Terrell N. Mundhenk, Claudio Prata Santiago, Soo Kyung Kim, and Joanne Taery Kim. “Deep symbolic regression: Recovering mathematical expressions from data via risk-seeking policy gradients”. In: *Int. Conf. ICLR 2021*. 2021 (cit. on p. 98).
- [Pop, 2009] Mihaela Pop, Maxime Sermesant, Damien Lepiller, Michael V. Truong, Elliot R. McVeigh, Eugene Crystal, Alexander Dick, Herve Delingette, Nicholas Ayache, and Graham A. Wright. “Fusion of optical imaging and MRI for the evaluation and adjustment of macroscopic models of cardiac electrophysiology: A feasibility study”. In: *Medical Image Analysis* 13.2 (2009). Includes Special Section on Functional Imaging and Modelling of the Heart, pp. 370–380 (cit. on p. 8).
- [Potse, 2018] Mark Potse. “Scalable and accurate ECG simulation for reaction-diffusion models of the human heart”. In: *Front. Physiol.* 9 (2018), p. 370 (cit. on p. 35).
- [Raissi, 2017] Maziar Raissi, Paris Perdikaris, and George Em Karniadakis. “Machine learning of linear differential equations using Gaussian processes”. In: *J. Comput. Phys.* 348 (2017), pp. 683–693 (cit. on pp. 11, 17).
- [Raissi, 2018a] Maziar Raissi. “Deep hidden physics models: Deep learning of non-linear partial differential equations”. In: *Journal of Machine Learning Research* 19.1 (2018), pp. 932–955 (cit. on p. 11).

- [Raissi, 2018b] Maziar Raissi and George Em Karniadakis. “Hidden physics models: Machine learning of nonlinear partial differential equations”. In: *Journal of Computational Physics* 357 (2018), pp. 125–141 (cit. on pp. [12](#), [17](#)).
- [Raissi, 2019] M. Raissi, P. Perdikaris, and G.E. Karniadakis. “Physics-informed neural networks: A deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations”. In: *Journal of Computational Physics* 378 (2019), pp. 686–707 (cit. on p. [12](#)).
- [Rapaka, 2012a] S. Rapaka, T. Mansi, B. Georgescu, M. Pop, G. A. Wright, A. Kamen, and Dorin Comaniciu. “LBM-EP: Lattice-Boltzmann Method for Fast Cardiac Electrophysiology Simulation from 3D Images”. In: *Medical Image Computing and Computer-Assisted Intervention – MICCAI 2012*. Ed. by Nicholas Ayache, Hervé Delingette, Polina Golland, and Kensaku Mori. Berlin, Heidelberg: Springer Berlin Heidelberg, 2012, pp. 33–40 (cit. on p. [71](#)).
- [Rapaka, 2012b] Saikiran Rapaka, Tommaso Mansi, Bogdan Georgescu, Mihaela Pop, Graham A Wright, Ali Kamen, and Dorin Comaniciu. “LBM-EP: Lattice-Boltzmann method for fast cardiac electrophysiology simulation from 3D images”. In: *Int. Conf. MICCAI*. Springer. 2012, pp. 33–40 (cit. on p. [21](#)).
- [Relan, 2011a] Jatin Relan, Phani Chinchapatnam, Maxime Sermesant, Kawal Rhode, Matt Ginks, Hervé Delingette, C Aldo Rinaldi, Reza Razavi, and Nicholas Ayache. “Coupled personalization of cardiac electrophysiology models for prediction of ischaemic ventricular tachycardia”. In: *Interface Focus* 1.3 (2011), pp. 396–407 (cit. on pp. [19](#), [32](#)).
- [Relan, 2011b] Jatin Relan, Mihaela Pop, Hervé Delingette, Graham Wright, Nicholas Ayache, and Maxime Sermesant. “Personalisation of a Cardiac Electrophysiology Model using Optical Mapping and MRI for Prediction of Changes with Pacing”. In: *IEEE Transactions on Bio-Medical Engineering* 12 (2011), pp. 3339–3349 (cit. on pp. [52](#), [53](#), [71](#), [90](#), [91](#), [97](#)).
- [Rodríguez-Padilla, 2018] Jairo Rodríguez-Padilla and Daniel Olmos-Liceaga. “Chebyshev multidomain pseudospectral method to solve cardiac wave equations with rotational anisotropy”. In: *International Journal of Modeling, Simulation, and Scientific Computing* 09.04 (2018), p. 1850025 (cit. on p. [74](#)).
- [Rudy, 2017] Samuel H Rudy, Steven L Brunton, Joshua L Proctor, and J Nathan Kutz. “Data-driven discovery of partial differential equations”. In: *Science Advances* 3.4 (2017), e1602614 (cit. on pp. [11](#), [17](#), [98](#)).
- [Ruiz Herrera, 2022] Carlos Ruiz Herrera, Thomas Grandits, Gernot Plank, Paris Perdikaris, Francisco Sahli Costabal, and Simone Pezzuto. “Physics-informed neural networks to learn cardiac fiber orientation from multiple electroanatomical maps”. In: *Engineering with Computers* 38.5 (2022), pp. 3957–3973 (cit. on p. [12](#)).

- [Sahli Costabal, 2020] Francisco Sahli Costabal, Yibo Yang, Paris Perdikaris, Daniel E. Hurtado, and Ellen Kuhl. “Physics-Informed Neural Networks for Cardiac Activation Mapping”. In: *Front. Phys.* 8 (2020), p. 42 (cit. on p. 12).
- [Sirignano, 2018] Justin Sirignano and Konstantinos Spiliopoulos. “DGM: A deep learning algorithm for solving partial differential equations”. In: *J. Comput. Phys.* 375 (2018), pp. 1339–1364 (cit. on pp. 11, 17).
- [Streeter, 1979] Dd. Streeter. “Gross morphology and fiber geometry of the heart”. In: 1979 (cit. on p. 74).
- [Ten Tusscher, 2004] K. H. W. J. Ten Tusscher, D. Noble, P. J. Noble, and A. V. Panfilov. “A model for human ventricular tissue”. In: *Am. J. Physiol. - Heart Circ. Physiol.* 286 (2004), H1573–H1589 (cit. on pp. 10, 31, 35, 36).
- [Ten Tusscher, 2006] K. H. W. J. Ten Tusscher and A. V. Panfilov. “Alternans and spiral breakup in a human ventricular tissue model”. In: *Am. J. Physiol. - Heart Circ. Physiol.* 291.3 (2006), H1088–H1100 (cit. on pp. 10, 31, 35, 36).
- [Thomas, 1957] C E Thomas. “The muscular architecture of the ventricles of hog and dog hearts”. In: *Am J Anat* 101.1 (1957), pp. 17–57 (cit. on p. 74).
- [Trayanova, 2021] Natalia A. Trayanova, Dan M. Popescu, and Julie K. Shade. “Machine Learning in Arrhythmia and Electrophysiology”. In: *Circulation Research* 128 (2021), pp. 544–566 (cit. on p. 30).
- [Wang, 2021] Rui Wang, Robin Walters, and Rose Yu. “Incorporating Symmetry into Deep Dynamics Models for Improved Generalization”. In: *International Conference on Learning Representations*. 2021 (cit. on p. 12).
- [Wang, 2022a] Rui Wang, Robin Walters, and Rose Yu. “Approximately equivariant networks for imperfectly symmetric dynamics”. In: *International Conference on Machine Learning*. PMLR. 2022, pp. 23078–23091 (cit. on p. 12).
- [Wang, 2022b] Rui Wang and Rose Yu. “Physics-Guided Deep Learning for Dynamical Systems: A Survey”. In: 2022 (cit. on p. 11).
- [Willard, 2020] Jared D. Willard, Xiaowei Jia, Shaoming Xu, Michael Steinbach, and Vipin Kumar. “Integrating physics-based modeling with machine learning: A survey”. In: *arXiv preprint:2003.04919* (2020) (cit. on pp. 11, 17).
- [Willard, 2022] Jared Willard, Xiaowei Jia, Shaoming Xu, Michael Steinbach, and Vipin Kumar. “Integrating Scientific Knowledge with Machine Learning for Engineering and Environmental Systems”. In: *ACM Computing Surveys* 1.1 (2022), pp. 1–35. eprint: 2003.04919 (cit. on p. 11).
- [Yang, 2021] Liu Yang, Xuhui Meng, and George Em Karniadakis. “B-PINNs: Bayesian physics-informed neural networks for forward and inverse PDE problems with noisy data”. In: *Journal of Computational Physics* 425 (2021), p. 109913 (cit. on p. 12).

- [Yin, 2021] Yuan Yin, Vincent Le Guen, Jeremie Dona, Ibrahim Ayed, Emmanuel de Bezenac, Nicolas Thome, and Patrick Gallinari. “Augmenting physical models with deep networks for complex dynamics forecasting”. In: *Int. Conf. ICRL*. 2021 (cit. on pp. [12](#), [31–33](#), [48](#)).
- [Zhang, 2018] Sheng Zhang and Guang Lin. “Robust data-driven discovery of governing physical laws with error bars”. In: *Proc. Math. Phys. Eng. Sci.* 474.2217 (2018), p. 20180305 (cit. on pp. [11](#), [17](#)).

List of Figures

1.1	Anatomy and conduction system of the heart. Illustration from Open Stax Anatomy & Physiology , CC-By license.	2
1.2	Anatomy and conduction system of the heart. Illustration from Wikimedia commons , CC-By license.	3
1.3	Cardiac Muscle organisation: a) Schematic structure of cardiac muscle cell connections; b) Photomicrograph of cardiac muscle cells showing the intercalated discs; c) An intercalated disc connects cardiac muscle cells and consists of desmosomes and gap junctions (which allow the ionic flow and propagation of action potential from one myocyte to another). Illustration taken from Open Stax Anatomy & Physiology , CC-By license.	4
1.4	Ventricular action potential evolution in mV. The main ions involved in each phase are represented next to the curve. Arrows indicate whether the ions enter (up) or leave (down) the cytoplasm through the specific ionic channels located in the cellular membrane. From [Cedilnik, 2020].	4
1.5	Infarct : a) Overview of the coronary system; b) Cross-section of the coronary artery with details on the occlusion process. In this thesis, we use the term "infarct scar" instead of "dead heart muscle" to refer to the necrotic tissue deprived from oxygen. Illustration from https: National Heart, Lung and Blood Institute , public domain.	6
1.6	Illustration of the clinical set-up for an x ray image-guided procedure in the EP lab, using a 3-dimensional electro-anatomical mapping system. Illustration taken from Rhythmia HDx, Boston scientific, USA	7
1.7	Translation of captured electrical signals into a 3D visualisation system, with the position of catheter being clearly visible. Illustration taken from Rhythmia HDx, Boston scientific, USA	7
1.8	Optical fluorescence imaging of action potential in an explanted porcine heart perfused via a Langendorff system: (a) snapshot of the optical set-up using two CCD cameras; (b) 2D optical image showing the location of the pacing electrode; (c) example of depolarization map recorded (red represents the depolarization phase, while blue represents the repolarization phase); and (d) example of normalized action potential waves recorded from one pixel, at a pacing frequency of 1.1 Hz (shown over 3.7 s). From [Camara, 2011].	9

2.1	Example of transmembrane potential (yellow) propagation in the cardiac tissue slab in absence (a) and the presence (b) of scar tissue, through successive time steps.	18
2.2	General setting used throughout the manuscript. Once trained, the EP-Net 2.0 model takes as input a context consisting of a few (i.e., $n=4$ here) observations of transmembrane potential dynamics, plus an indication of the scar area (left), and forecasts transmembrane potential depolarisation wave propagation (right).	18
2.3	The ResNet architecture used in EP-Net 2.0. It has 5 input frames (mask frame plus 4 frames of the transmembrane potential) and 1 output frame of forecast.	20
2.4	(a) Representative results for the trained EP-Net 2.0 model (9 ms of forecast, conductivity of the cardiac slab tissue $\sigma = 2$). (b,c) Plots of the transmembrane potential at the leftmost upper point (0,0) and the rightmost bottom point (23,23) in the slab with different forecasting horizons.	22
2.5	Results of trained EP-Net 2.0 model on scar with circular (top three rows) and complex (bottom three rows) shape (9 ms of forecast, cardiac slab conductivity $\sigma = 2$).	24
2.6	Results of trained EP-Net 2.0 model with two stimulation currents applied on different pixels and at different times (9 ms of forecast, cardiac slab conductivity $\sigma = 2$).	24
2.7	Results of EP-Net 2.0 model using a scar of circular shape and cardiac slab conductivity $\sigma = 3.8$ (top three rows), and a scar of triangular shape, two onsets and cardiac slab conductivity $\sigma = 1.5$ (bottom three rows).	24
2.8	Results of trained EP-Net 2.0 model on thin scar with circular shape. . . .	25
2.A.1	Examples of masks containing 3 values of tissue conductivity parameter (0-conductivity, dark blue area, to model a scar tissue) used for training. .	27
2.A.2	Test results of EP-Net 2.0 trained on masks with multiple tissue conductivities. Forecast of 21 ms, first frame is an input mask.	28
3.1.1	General APHYN-EP framework scheme. During the training phase two-component framework learn the parameters for the physical (F_p) and the data-driven (F_a) components from data. Then, via an ODE solver, the framework can further forecast the learned dynamics.	32
3.2.1	General APHYN-EP framework training scheme.	34
3.2.2	Calculation of the Loss function gradient via Pytorch framework (blue part) and Adjoint model (yellow part).	34
3.3.1	0D signal example simulated via 3 EP models. Legend: Ten Tusscher – Panfilov model [Ten Tusscher, 2006] (TTP), Ten Tusscher – Noble – Noble – Panfilov ionic model [Ten Tusscher, 2004] (TTNNP), Mitchell-Schaeffer model [Mitchell, 2003] (M-Sch).	36

3.3.2	(a,b) Example of a selected 2D myocardial tissue slab with the transmembrane potential activation (yellow) and resting phase (dark blue), for train and validation dataset, respectively. (c) Typical temporal sequence for the simulation experiment (without noise), with the normalised amplitude of the transmembrane potential being represented on the Y-axis and the time (ms) on the X-axis.	37
3.4.1	Validation results of the trained framework with learning of: (a) 2 (τ_{in} and τ_{out}) and (b) 3 (τ_{in} , τ_{out} and τ_{close}) physical parameters. Legend: ground truth (GT), prediction of the framework (Prediction FW), decomposition of prediction on physical (F_p) and DL (F_d) parts.	39
3.4.2	APHYN-EP predicted dynamics for the transmembrane potential diffusion. Figure shows a period of 8 ms of the forecast. Red point is the reference point for Figure 3.4.3.	40
3.4.3	Transmembrane potential at point (5,5) in the cardiac slab (red point, see Fig. 3.4.2): (a) Original, (b) Zoom on first 40 ms. Legend: ground truth (GT), APHYN-EP, physical (F_p) and data-driven (F_d) component of APHYN-EP.	41
3.4.4	Exemplary results illustrating: the dynamics of the transmembrane potential diffusion predicted by the APHYN-EP physical component (second row); the error with ground-truth diffusion for this physical component of APHYN-EP (third row); and, the trained APHYN-EP data-driven component contribution (bottom row).	41
3.4.5	APHYN-EP predicted dynamics for the transmembrane potential diffusion of planar wave. The frames show a period of 8 ms of forecast obtained without re-training the APHYN-EP framework.	42
3.A.1	APHYN-EP predicted dynamics for the transmembrane potential diffusion. The figure shows a 9 ms of forecast).	44
3.A.2	(a,b) Transmembrane potential at point (10,10) in the cardiac slab (red point, see Fig. 3.A.1) with different forecasting horizons. Ground truth (GT), APHYN-EP, physical (F_p) and data-driven (F_d) component of APHYN-EP.	44
3.A.3	Exemplary results illustrating: the dynamics of the transmembrane potential diffusion predicted by the APHYN-EP physical component (second row); the error with ground-truth diffusion for this physical component of APHYN-EP (third row); and, the trained APHYN-EP data-driven component contribution (bottom row).	44
3.A.4	Transmembrane potential ground truth (GT), generated by APHYN-EP, by Physical model and ResNet model at the leftmost upper point (1,1) (a,b), at point (10,10) (c,d) and the rightmost bottom point (23,23) (e,f) in the cardiac slab with different forecasting horizons, the same GT dynamics as in Fig. 3.A.1-3.A.2.	46

4.2.1	Example of optical mapping data (tracings of denoised AP waves) recorded ex vivo in a porcine heart. ROI B represents an ischaemic region characterised by a shorter action potential duration (APD) compared to the normal APD recorded in ROI A.	50
4.2.2	Example of ex vivo optical mapping data with tracings of denoised AP waves recorded at multiple pacing rates in a healthy porcine heart.	50
4.2.3	Impact of different Mitchell-Schaeffer model parameters on modelling of cardiac action potential cycles. The variable v represents a normalised ($v \in [0, 1]$) dimensionless transmembrane potential, variable h is “gating” variable controlling the repolarisation phase. From [Cedilnik, 2020].	52
4.2.4	Transmembrane potential (v , left) and gating variable (h , right) curves. Legend: action potential duration (APD), diastolic interval (DI). From [Mitchell, 2003].	53
4.2.5	The restitution curve derived from the 2-current model 2.1, including the threshold behavior. From [Mitchell, 2003].	53
4.2.6	Schematic diagram of the Full personalisation method for the APHYN-EP framework.	54
4.2.7	Schematic diagram of the Fast personalisation method for the APHYN-EP framework.	55
4.3.1	Validation results of the framework trained on: (a) ROI A data and (b) ROI B data, respectively. Ground truth (GT) data, prediction of the framework (Prediction FW), the decomposition of prediction into physical (F_p) and DL (F_d) components.	57
4.3.2	Validation results of the framework with simpler NN (MLP) as F_d component trained on: (a) ROI A data and (b) ROI B data, respectively. Ground truth (GT) data, prediction of the framework (Prediction FW_MLP), decomposition of prediction on physical (F_p) and DL (F_d) components.	58
4.3.3	Validation results of the Physical model (2.1) trained alone on: (a) ROI A data and (b) ROI B data, respectively. Ground truth (GT) data, prediction of the physical model (Prediction PhysM).	58
4.3.4	Validation results of a fully data-driven model (EP-Net 2.0 [Kashtanova, 2021]) trained on: (a) ROI A data and (b) ROI B data, respectively. Ground truth (GT) data, prediction of the data-driven model (Prediction DDM).	59
4.3.5	(a) Validation results of the F_p trained alone on ROI A data to obtaining $\theta_p(A)$ parameters. (b) Validation results for APHYNITY framework (with fixed $\theta_p(A)$ parameters) trained on ROI A data. (c) Validation results of the APHYNITY framework (with fixed $\theta_p(A)$ parameters) with F_d component trained on ROI B data. Legend: ground truth (GT) data, prediction of the framework (Prediction FW), decomposition of prediction on physical (F_p) and DL (F_d) components.	60

4.3.6	(a) Validation results of the framework for ROI A, with the following legend: Ground truth (GT) data, prediction of the framework (Prediction FW_A) trained on only ROI A data, and prediction of the framework composed of (F_p) component trained on ROI A data and DL (F_d) component trained on ROI B data (Prediction FW_AB). (b) Comparison of DL component trained on ROI A data ($F_{d_A}(A)$) and on ROI B data ($F_{d_B}(A)$) applied on validation ROI A data.	60
4.3.7	(a) Validation results of the F_p trained alone on ROI B data to obtain $\theta_p(B)$ parameters. (b) Validation results of the APHYNITY framework (with fixed $\theta_p(B)$ parameters) trained on ROI B data. (c) Validation results of the APHYNITY framework (with fixed $\theta_p(B)$ parameters) with F_d component trained on ROI A data. Legend: ground truth (GT) data, prediction of the framework (Prediction FW), decomposition of prediction on physical (F_p) and DL (F_d) components.	61
4.3.8	(a) Validation results of the framework for ROI B, with the following legend: Ground truth (GT) data, prediction of the framework (Prediction FW_B) trained on only ROI B data, and prediction of the framework composed of (F_p) component trained on ROI B data and DL (F_d) component trained on ROI A data (Prediction FW_BA). (b) Comparison of DL components trained on ROI B data ($F_{d_B}(B)$) and on ROI A data ($F_{d_A}(B)$) applied on validation ROI B data.	61
4.3.9	(a) Validation results of the F_p trained alone on ROI C data to obtain $\theta_p(C)$ parameters. (b) Validation results of the APHYNITY framework (with fixed $\theta_p(C)$ parameters) with F_d component trained on ROI A data. (c) Validation results of the APHYNITY framework (with fixed $\theta_p(C)$ parameters) with F_d component trained on ROI B data. Legend: ground truth (GT) data, prediction of the framework (Prediction FW), decomposition of prediction on physical (F_p) and DL (F_d) components.	62
4.3.10	(a) Validation results of the framework for ROI C, with the following legend: Ground truth (GT) data, prediction of the framework composed of (F_p) component trained on ROI C data and DL (F_d) component trained on ROI A data (Prediction FW_CA), and prediction of the framework composed of (F_p) component trained on ROI C data and DL (F_d) component trained on ROI B data (Prediction FW_CB). (b) Comparison of DL component trained on ROI A data ($F_{d_A}(C)$) and on ROI B data ($F_{d_B}(C)$) applied on ROI C data.	62
4.3.11	Validation results of the framework trained on data described in section 4.2.1.2, $\tau_{in} = 0.4$, $\tau_{out} = 3$, $\tau_{close} = 350$. Ground truth (GT) data, prediction of the framework (Prediction FW), decomposition of prediction into physical (F_p) and DL (F_d) components, respectively.	63

4.3.12	Test results of the framework trained on data described in section 4.2.1.2 and its decomposition on physical and DL parts. Ground truth (GT) data, prediction of the framework (Prediction FW), decomposition of prediction on physical (F_p) and DL (F_d) parts.	64
4.3.13	Test results of the framework trained on data described in section 4.2.1.2 and its decomposition on physical and DL parts. Ground truth (GT) data, prediction of the framework (Prediction FW), decomposition of prediction on physical (F_p) and DL (F_d) parts.	65
4.3.14	Test results of the framework trained on data described in section 4.2.1.2 and its decomposition on physical and DL parts. Ground truth (GT) data, prediction of the framework (Prediction FW), decomposition of prediction on physical (F_p) and DL (F_d) parts.	65
4.3.15	Test results of the framework trained on data described in section 4.2.1.2 and its decomposition on physical and DL parts. Ground truth (GT) data, prediction of the framework (Prediction FW), decomposition of prediction on physical (F_p) and DL (F_d) parts.	66
4.4.1	Example of spontaneous AP activation produced by F_d component.	67
5.1.1	Comparison of measured and simulated 3D activation maps obtained by changing the parameter value for the wall thickness, which is further correlated to the velocity transfer function. From [Cedilnik, 2018].	71
5.2.1	Example of depolarisation maps selected from the optical data recorded ex vivo in a porcine heart, with the stimulating electrode (used for pacing the heart) placed onto: (a) the RV, and (b) the LV, respectively. Red areas correspond to early activation times (i.e., where the excitation pulse was delivered) while the late depolarisation times are depicted in green-blue.	73
5.2.2	Example of simulated transmembrane potential (yellow) propagation in a 2D cardiac tissue slab modelled using: (a) isotropic diffusion tensor; and, (b) anisotropic diffusion tensor (note that the stimulation started in the centre of the square).	73
5.2.3	Schematic diagram of the APHYN-EP framework using Deep Personalisation method . Personalisation on one heart with two pacing locations indicated by the position of the electrode (left vs. right).	76
5.2.4	Example of simulating various neighboring patches (10x10 pixels) orders from an initial patch in image of 50x50 pixels (without overlapping). (a) Initial patch, (b) first order neighboring patches, (c) second order neighboring patches and (d) third order neighboring patches.	78
5.2.5	Data-augmentation example of one patch of true ex vivo recorded 2D optical mapping data of action potential. The first frame is the true patch, and the next frames are its flipped versions.	79

5.3.1	APHYN-EP predicted dynamics for the transmembrane potential diffusion, RV pacing (data used for training). The frames show results obtained over a 30ms period of forecast.	82
5.3.2	APHYN-EP predicted dynamics for the transmembrane potential diffusion, LV pacing (test data). The frames show a period of 30 ms of forecast obtained without re-training the APHYN-EP framework.	82
5.3.3	Snapshots of APHYN-EP predicted dynamics for the transmembrane potential diffusion (validation dataset) at different training epochs. The frames show a period of 10 ms of forecast.	84
5.3.4	APHYN-EP predicted dynamics for the transmembrane potential diffusion on the entire heart with RV pacing. The frames show a period of 107 ms of forecast obtained without re-training the APHYN-EP framework.	85
5.3.5	APHYN-EP predicted dynamics for the transmembrane potential diffusion on the entire heart with LV pacing. The frames show a period of 107 ms of forecast obtained without re-training the APHYN-EP framework.	86
5.3.6	APHYN-EP predicted dynamics for the transmembrane potential diffusion extrapolated to the entire heart with RV pacing. The frames show a period of 107 ms of forecast obtained without re-training the APHYN-EP framework.	87
5.3.7	APHYN-EP predicted dynamics for the transmembrane potential diffusion extrapolated to the entire heart with RV pacing. The frames show a period from 111 ms to 181 ms of forecast obtained without re-training the APHYN-EP framework.	88
5.3.8	APHYN-EP predicted dynamics for the transmembrane potential diffusion extrapolated to the entire heart with LV pacing. The frames show a period of 70 ms of forecast obtained without re-training the APHYN-EP framework.	89
5.4.1	Parameter maps for a healthy heart with LV pacing location. Estimated d (σ in our F_p model) and τ_{close} values per zone. From [Relan, 2011b]	91
6.3.1	Ventricular Tachycardia (VT) modelling pipeline using non-invasive CT data to define the ablation targets for VT Ablation Planning. From [Cedilnik, 2018].	98
6.3.2	(a) Myocardial mesh with measured pacing locations (red), (b) 64 BSPM electrodes (black), myocardial mesh (pink). From [Giffard-Roisin, 2017b].	98

List of Tables

2.1	Relative mean-squared error (MSE) of transmembrane potential forecasting in the presence of scars of various forms, computed for different forecasting horizons (cardiac slab conductivity $\sigma = 2$).	23
2.2	Relative mean-squared error (MSE) of potential forecasting in presence of multiple onsets and the scar of rectangular shape for different forecasting horizons (cardiac slab conductivity $\sigma = 2$).	23
2.3	Relative mean-squared error (MSE) of potential forecasting in the presence of various conduction velocities of cardiac slab and scar of rectangular form for different forecasting horizons.	25
3.4.1	Mean-squared error, MSE ($\times 10^{-3}$) of the normalised transmembrane potential (adimensional) forecasting (forecasting horizon of 300 ms). Baseline models: the Physical model (2.1) and a fully data-driven model (EP-Net 2.0 [Kashtanova, 2021]) trained on the same dataset as APHYN-EP. Out-of-domain tests: Planar wave.	42
3.A.1	Mean-squared error (MSE) of normalised transmembrane potential forecasting per time-step for different forecasting horizons.	45
4.3.1	Mean-squared error, MSE ($\times 10^{-3}$) of the normalised, adimensional transmembrane potential, forecasting with forecasting horizon of 300 ms. Baseline models: the Physical model (2.1) and a fully data-driven model (EP-Net 2.0 [Kashtanova, 2021]) trained on the same dataset as APHYN-EP.	57
5.3.1	Mean-squared error, MSE ($\times 10^{-3}$) of the normalised transmembrane potential (adimensional) forecasting (forecasting horizons of 300 ms and 150 ms). Baseline models: the Physical model (2.1) and a fully data-driven model (EP-Net 2.0 [Kashtanova, 2021]) trained on the same dataset as APHYN-EP. Out-of-domain tests: LV pacing.	81

