



HAL
open science

Simultaneous data assimilation and cardiac electrophysiology model correction using differentiable physics and deep learning

Victoriya Kashtanova, Mihaela Pop, Ibrahim Ayed, Patrick Gallinari, Maxime Sermesant

► To cite this version:

Victoriya Kashtanova, Mihaela Pop, Ibrahim Ayed, Patrick Gallinari, Maxime Sermesant. Simultaneous data assimilation and cardiac electrophysiology model correction using differentiable physics and deep learning. *Interface Focus*, 2023, 13 (6), <10.1098/rsfs.2023.0043>. <hal-04359753>

HAL Id: hal-04359753

<https://hal.science/hal-04359753v1>

Submitted on 21 Dec 2023

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.



HAL Authorization

Simultaneous Data Assimilation and Cardiac Electrophysiology Model Correction Using Differentiable Physics and Deep Learning

Victoriya Kashtanova^{1,2} Mihaela Pop^{1,3} Ibrahim Ayed^{4,5}
Patrick Gallinari^{4,6} Maxime Sermesant^{1,2}

¹Inria Université Côte d’Azur, Nice, France

²3IA Côte d’Azur, Sophia Antipolis, France

³Sunnybrook Research Institute, Toronto, Canada

⁴Sorbonne University, Paris, France

⁵Therisis lab, Paris, France

⁶Criteo AI Lab, Paris, France

{victoriya.kashtanova@inria.fr, maxime.sermesant@inria.fr}

Abstract

Modelling complex systems, like the human heart, has made great progress over the last decades. Patient-specific models, called "digital twins", can aid in diagnosing arrhythmias and personalising treatments. However, building highly accurate predictive heart model requires a delicate balance between mathematical complexity, parameterisation from measurements, and validation of predictions. Cardiac electrophysiology models range from complex biophysical models to simplified phenomenological models. Complex models are accurate but computationally intensive and challenging to parameterise, while simplified models are computationally efficient but less realistic. In this paper, we propose a hybrid approach by leveraging deep learning to complete a simplified cardiac model from data. Our novel framework has two components, decomposing the dynamics into a physics-based and a data-driven term. This construction allows our framework to learn from data of different complexity, while simultaneously estimating model parameters. First, using *in silico* data, we demonstrate that this framework can reproduce the complex dynamics of cardiac transmembrane potential even in presence of noise in the data. Second, using *ex vivo* optical data of action potentials, we demonstrate that our framework can identify key physical parameters for anatomical zones with different electrical properties, as well as to reproduce the action potential wave characteristics obtained from various pacing locations. Our physics-based data-driven approach may improve cardiac electrophysiology modeling by providing a robust biophysical tool for predictions.

Keywords— Physics-based learning, Deep Learning, Cardiac electrophysiology, Simulations

1 Introduction

Computational cardiology is a multi-disciplinary field that has seen extensive progress in the past decade [23, 8, 4]. 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 routinely employed to diagnose arrhythmia, as well as to help planning the best therapy on an individual basis [10]. However, in order to build such accurate predictive heart models, one needs to select the most suitable theoretical framework, balancing the degree of mathematical complexity needed for the specific problem studied, the correct parameterisation of model from measurements, and the validation of predictions.

Physiological and multi-physics phenomena characterising the heart function in normal and pathological conditions can be mechanistically described by mathematical models of different complexity [5]. Notably, several electrophysiology (EP) models are able to accurately reproduce the electrical behaviour

of the heart at different scales (i.e. cell, tissue, organ). For instance, detailed biophysical models have been proposed in order to describe the dynamics of transmembrane potential, the flowing currents as well as the different ionic concentrations inside and outside the cardiac cell, such as the complex Ten Tusscher-Panfilov model [34, 35]. On the other hand, ionic models are not only intricate but also computationally expensive, requiring costly resources. In addition, these sophisticated models also employ numerous hidden variables that are impossible to be all measured, making difficult to accurately identify all model parameters.

An attractive alternative to complex computational models, is to use phenomenological models involving descriptions derived from simple biophysical models adapted to cardiac EP. Examples of such models are the FitzHugh-Nagumo [11, 25], Aliev-Panfilov [1, 26], and Mitchell-Schaeffer [24] models, which only make use of a few variables and parameters. These models are used for rapid simulations of the action potential wave propagation at tissue level (e.g. on a 2D tissue layer or a 3D tissue slab), and organ level (through the entire heart). Nevertheless, being less realistic, these simplified models are also less accurate. Given the increasing availability of synthetically-generated data and observation data, an important research issue is to consider how machine learning (ML), and more specifically, deep learning (DL), could help complementing simple physical models in order to improve their accuracy.

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 [17, 39, 38]. This research topic motivates works from different disciplines, ranging from climate to aeronautics and biology, by encompassing diverse objectives including: accelerating numerical simulations; 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 [29, 30, 41] or as strict constraints in the NN architecture [15]; enforcing invariance or equivariance properties characteristic of physical laws in DL models [36, 37], combining DL and physical components into hybrid systems [42].

Pure DL methods have been also increasingly used in order to learn dynamical models from data and make intelligent decision without human intervention. Notably, neural networks are capable to learn and to model various complex (e.g. nonlinear) relationships between input and output data [40]. For example, a framework useful for automatic learning PDEs from data [22, 21] has been proposed. Another group used an adjoint method to learn differential equations parameterised with neural networks [7], while *Ayed et al.* [3] 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 [2, 19], these works suggested that data-driven models alone were not able to reproduce complex unseen dynamics (such as the repolarisation phase of an action potential) [19]; thus, the maximum forecasting horizon remains limited.

Subsequently, several researchers have started to use coupled physico-statistical approaches for cardiac EP simulations, in order to achieve a high precision at low cost. For example, one group designed a neural network that approximates the FitzHugh-Nagumo model [9], while others used a physics-informed neural networks to construct cardiac activation maps by accounting for the underlying wave propagation dynamics [33] and to estimate the cardiac fiber architecture of the human atria from catheter recordings of the electrical activation [32]. 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 [12]. 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 [14]. Unfortunately, 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 real data may not only be computationally expensive, but also difficult especially in order to properly deal with frequently observed large discrepancies between simulated and real data.

To address this critical limitation, here we propose a framework to Augment incomplete PHYSical models with a deep learning component for ideNtifying complex cardiac ElectroPhysiology dynamics (APHYN-EP) from data, based on a fast low-fidelity (or incomplete) physical model. This framework has two main components which decompose the model dynamics into a physical term and a data-driven term, respectively. The data-driven DL component is designed such that it captures only the information that cannot be modeled by the incomplete physical model. This construction allows our framework to learn from data of different complexity and to handle the different temporal and spatial resolution. The proposed framework closely follows the approach of *Yin et al.* [42]. However, in contrast to this previous work (which considers fully-observable dynamics and simple test use cases), cardiac EP dynamics have a

high complexity and represent simultaneously multiple underlying processes. Furthermore, most cardiac EP models lack measurements for some of the model variables, which requires inferring the dynamics from incomplete observations only, making a model partially-observable.

Figure 1 presents the general scheme of our approach. The observed data, which represent only part of cardiac EP dynamics, is decomposed by the framework into two "physics-based" and "data-driven" parts. The "data-driven" part acts as a complement (or a correction) to the "physics-based" part, and handles quantities and dynamics which were neglected. The parameters of the physical model (inverse problem) and of the neural network (direct problem) were learned simultaneously during the training phase (see Fig. 2). After the complete training, the APHYN-EP framework can be used to forecast the learned dynamics at multiple horizons.

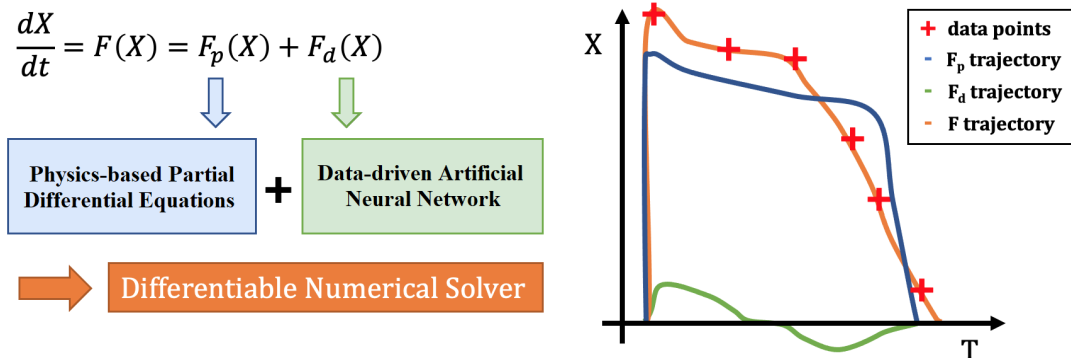


Figure 1: General APHYN-EP framework. Correction of a PDE-based model with an additional term learned from the data.

We showed in previous papers [18, 20] the preliminary results of the this framework application. In this current paper, we present in detail our novel framework, demonstrating results of its application on more complex settings. Furthermore, we also explore the generalisation capabilities of the framework.

In Section 2, we give a thorough overview of the method, with training details and a comparison with a traditional numerical method. In Section 3, we introduce the data used to test different abilities of the framework. In Section 4, we present numerical results obtained using our framework in comparison to the baseline methods. Finally, in Section 5, we discuss the possible framework improvements and future applications.

2 Learning Framework

The cardiac EP dynamics can be driven by an equation of the form:

$$\frac{dX_t}{dt} = F(X_t) \quad (1)$$

defined over a finite time interval $[0, T]$, where the state X_t is a spatio-temporal vector field over the domain $\Omega \in \mathbb{R}^k$, with $k \in \{0, 2\}$ in our experiments. We suppose that we have access to a set of observed trajectories:

$$\mathcal{D} = \{X : [0, T] \rightarrow A | \forall t \in [0, T], \frac{dX_t}{dt} = F(X_t)\},$$

where A is the set of vector field values X . In our case, the unknown F has A as domain and we only assume that $F \in \mathcal{F}$, with $(\mathcal{F}, \|\cdot\|)$ a normed vector space.

Since we consider that only part of the measured dynamics can be modeled via a family of PDEs characterised by their temporal evolution $F_p \in \mathcal{F}_p \subset \mathcal{F}$, APHYN-EP framework introduce a data-driven augmentation term $F_d \in \mathcal{F}$ complementing F_p . \mathcal{F} being a vector space, thus we can write:

$$F = F_p + F_d, \quad (2)$$

where F_p represents a physical model (which is an incomplete description of the underlying phenomenon) and F_d is a neural network complementing the physical model by capturing the information that cannot

be modeled by the physics-described component. The physical and data-driven parameters defining F_p and F_d are unknown and need to be estimated from data by fitting the trajectories from \mathcal{D} .

The decomposition (2) is not unique, as shown in [2, 19], all the measured dynamics could be captured by the F_d component alone. This decomposition is therefore ill-defined, which complicates the interpretability and extrapolation possibilities of the framework. However, we aim to reduce the action of F_d on the dynamics measured through its norm, and to find its optimal minimum.

Therefore, in order to learn the cardiac EP dynamics X_t , in this work, we solve following optimisation problem via our physics-based data-driven APHYN-EP framework:

$$\min_{F_p \in \mathcal{F}_p, F_d \in \mathcal{F}} \|F_d\| \quad \text{subject to } \forall X \in \mathcal{D}, \forall t \in [0, T], \frac{dX_t}{dt} = F(X_t) = F_p(X_t) + F_d(X_t). \quad (3)$$

Assuming that \mathcal{F}_p is a Chebyshev set, Propositions 1 and 2 from *Yin et al.* [42] guarantee the existence and uniqueness of a minimising pair for (3).

Specifically, in our experiments the incomplete physical model is derived ($F_p(X_t) = v_t$) from the two-variable (v, h) model [24] for cardiac EP simulations, as described by equations (4). The variable v 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):

$$\begin{aligned} \partial_t v &= \text{div}(\sigma \text{I}\nabla v) + \frac{hv^2(1-v)}{\tau_{in}} - \frac{v}{\tau_{out}} + J_{stim} \\ \partial_t h &= \begin{cases} \frac{1-h}{\tau_{open}} & \text{if } v < v_{gate} \\ \frac{-h}{\tau_{close}} & \text{if } v > v_{gate} \end{cases} \end{aligned} \quad (4)$$

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.

This physical model has been successfully used in patient-specific modelling [31], 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 set in the numerical analysis. Thus, assuming the initial conditions for this system (4) $v(t=0) = 0$ 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 (see Section 3), \mathcal{F}_p is the set of models spanned by the R.H.S. of the equations above 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 guaranteeing theoretical existence and uniqueness of a minimising 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 paper, we used a ResNet network [13], because it could accurately reproduce complex cardiac EP dynamics [2, 19]. However, a simpler neural network could also be used for more rapid computations, as discussed in Section 5.

In APHYN-EP framework the physical (F_p) and the data-driven (F_d) components are trained simultaneously. This insures the finding of the best minimising pair for (3) determined by the set of parameters $\theta = (\theta_p, \theta_d)$ based on observed measurements $X_{h\Delta t}^{(i)}, i \in [1, N]$ and $h \in [1, T/\Delta t]$.

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:

$$\mathcal{L}(\theta) = \lambda * \mathcal{L}_{traj}(\theta) + \left\| F_d^{\theta_d} \right\| = \lambda * \sum_{i=1}^N \sum_{h=1}^{T/\Delta t} \|X_{h\Delta t}^{(i)} - \tilde{X}_{h\Delta t}^{(i)}(\theta)\| + \left\| F_d^{\theta_d} \right\| \quad (5)$$

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 [7, 6].

The key role of the λ coefficient is to balance the two parts of the loss. During training, we used λ in a dynamic state, as $\lambda_{j+1} = \lambda_j + \gamma \mathcal{L}_{traj}(\theta_{j+1})$ (j 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.* [42], is presented below.

Figure 2(a) 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

Algorithm 1 APHYN-EP training

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

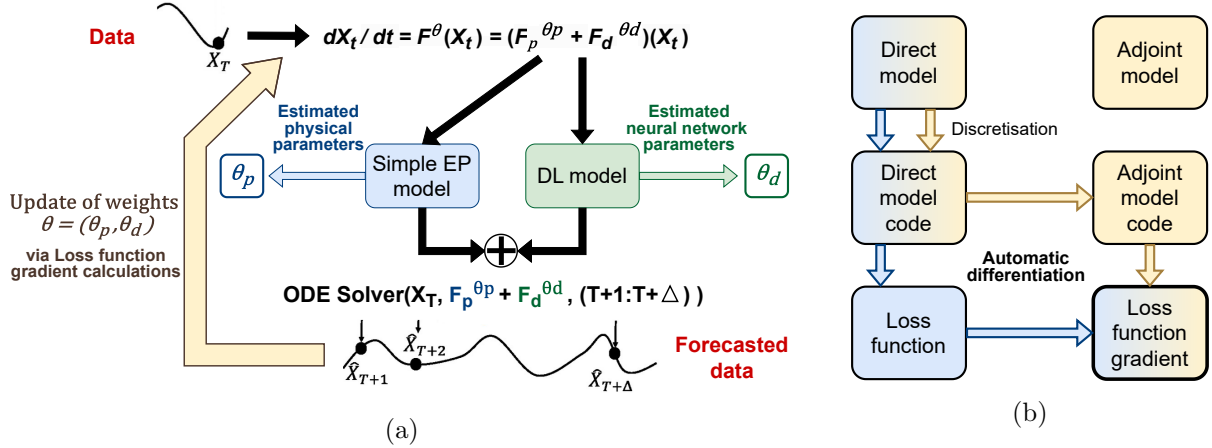


Figure 2: (a) General APHYN-EP framework training scheme. (b) Calculation of the Loss function gradient via Pytorch framework (blue part) vs. Adjoint model (yellow part).

Problems, which is based on the solution of an adjoint problem to calculate the gradient of the Loss function [16] (Fig. 2(b), yellow part)). However, thanks to the automatic differentiation tools provided by the Pytorch library [27], the gradient of (5) is calculated automatically inside of PyTorch framework (Fig. 2(b), blue part)).

Additionally, in order to train simultaneously the physical and the data-driven components of APHYN-EP, we implemented the Laplace operator in (4) 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 computing the physical component and for computing the final forecast given by the framework, respectively.

3 Experiments

In order to test the performance of our APHYN-EP framework and to further show its capability to reproduce transmembrane potential dynamics of different complexities, we chose two types of experiments. First, using synthetic *in silico* data, we tested the ability of the framework to reproduce the complex dynamics of transmembrane potential including a case where noise is present in the data. Second, using optical fluorescence imaging data of action potentials recorded *ex vivo* on explanted porcine hearts, we aimed to show that our framework: a) can identify key physical parameters for different anatomical zones having abnormal electrical function (using 0D data from a heart with an ischemic region); b) is capable to reproduce the action potential wave characteristics obtained following heart stimulation from different locations (2D mapping data from a healthy heart).

The details of data collection and training settings used for the experiments are presented in detail below.

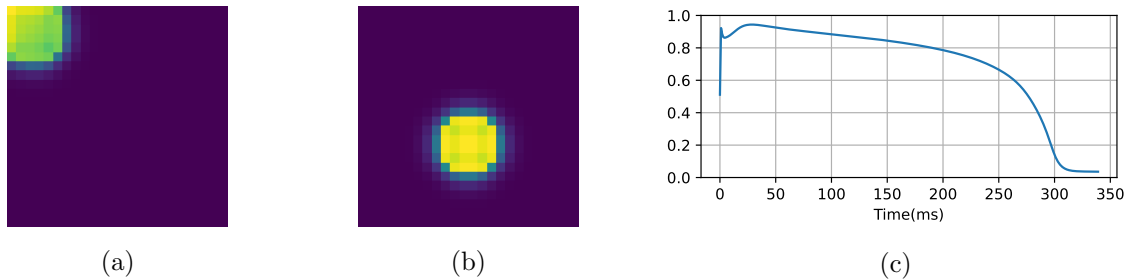


Figure 3: (a,b) Example of a selected 2D myocardial tissue slab with the transmembrane potential activation (yellow) and resting phase (dark blue), for train and validation dataset, respectively. (c) Typical temporal sequence for the simulation experiment (without noise), with the normalised amplitude of the transmembrane potential being represented on the Y-axis and the time (ms) on the X-axis.

3.1 *In silico* data

3.1.1 Data collection

To evaluate our method, we used a dataset of transmembrane potential activations simulated by employing a monodomain reaction-diffusion equation and the Ten Tusscher – Panfilov ionic model [35], which represents twelve different transmembrane ionic currents. The simulations were performed using a spatial step of 0.2 mm and a time step of 1 ms (similarly to those used in the original model [35]), with the open-source finetwave software¹. For this, our 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 400 ms of a heart beat, and was intended to achieve a full depolarisation-repolarisation cycle (note that video examples showing simulated dynamics across time are available online²). This required 20 seconds of computation time on a 8-core Intel i7-7820HQ CPU per simulation.

0D data In order to generate 0D databases, we selected the 2D 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. Next, we saved in separate files a time sequence for each pixel of the 2D cardiac slab, 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 (for non-activated at time 0 pixels), thus obtaining a time sequence of 350 ms per data sample, as shown in Figure 3(c). To test the ability of the framework to operate with noisy data, we added to each time sequence a 5 percent random noise characterised by a normal Gaussian distribution.

2D data In order to generate 2D databases, we applied an excitation pulse at each grid point of 2D cardiac tissue slab. Next, 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 – 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.

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

²<https://doi.org/10.6084/m9.figshare.21648752.v3>

3.1.2 Training settings

The physical model (F_p) described by Eq. 4 was implemented with a standard finite-difference scheme for the Laplace operator, using a spatial resolution of 1 mm² 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 (4), 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 models in our simulations). The other Mitchell-Schaeffer model parameters were taken from the original paper [24], 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 [13].with 4 input/output channels (assimilating the first milliseconds of dynamics). In order to demonstrate and to compare APHYN-EP framework’s performance with simpler neural network, we also tested an MLP network (for 0D experiments) and a ConvNet (for 2D experiments) as a deep learning component (F_d) of the framework.

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. More details on the training configurations can be found in the Supplementary material.

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. Furthermore, the major part of the inference time is taken by computing the integral of $\tilde{X}_T^{(i)}$ from $X_0^{(i)}$ and could be reduced by using a larger time step.

3.2 *Ex vivo* data

3.2.1 Data collection

As mentioned above, we tested the APHYN-EP framework performance using *ex vivo* experimental datasets from optical fluorescence imaging of action potential. Briefly, the optical signals were recorded *ex vivo* from the epicardial surface of hearts explanted from juvenile swine (25kg in weight). Each heart was attached to a Langendorff perfusion system and a voltage sensitive dye (i.e., di-4-ANEPPS) was injected into the perfusate solution. In order to avoid cardiac motion artifacts during the optical recordings, the heart contraction was suppressed by a bolus of saline and Cytochalasin D, an electro-mechanical uncoupler. All optical images were acquired using a high-speed CCD camera (MICAM02, BrainVision Inc. Japan), with high-temporal resolution (3.7ms) as well as a high-spatial resolution (i.e., pixel size 0.7mm x 0.7mm). The action potential was then derived at each pixel from the relative change in the intensity of fluorescence signal. The *ex vivo* studies were described in more detail in *Pop et al.* [28]. The recorded optical signals were saved and exported from the BV-Ana acquisition software. Notably, within each pixel, the relative change in optical signal intensity is directly proportional with the action potential.

0D data Our specific objective was to learn the complex dynamics of the measured action potential (AP), and then to identify its relevant physical parameters for different zones of the heart. Note that the optical data were considered as the ground truth. For this 0D experiment, we chose to use epicardial optical data recorded from a heart having an ischaemic region (artificially induced by occluding a blood vessel). For the analysis, we manually selected two rectangular regions of interest (ROIs) with different AP dynamics across time, as seen in Figure 4. For post-processing of the data, we normalised the optical signal intensity in order to obtain a [0, 1] min/max interval for the transmembrane potential, while keeping the noise in the data. Next, 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 in separate files a time sequence for each pixel in the two selected ROI, creating two databases (ROI A and ROI B), each containing approximately 10 and 5 time-sequences for training and validation, respectively.

2D data For this experiment, we chose a healthy heart stimulated via a pacing electrode placed onto two different zones within the heart (see Fig. 5), having two recordings 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 2D data. We manually selected a rectangular region of interest (see Fig. 5), where there were no changes in the tissue physical properties. Next, we divided this region of the heart into smaller 2D squares of 10 pixels. For each square, we took

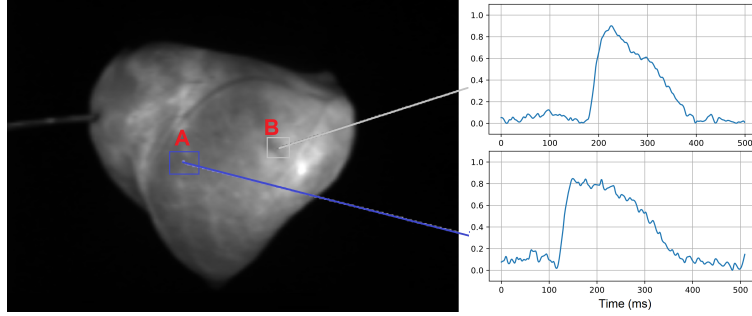


Figure 4: Example of mapping data from the optical experiment: (left) explanted porcine heart, and (right) tracings of denoised action potential waves recorded from the epicardial surface of the left ventricle, LV. Note that ROI B represents an ischaemic region characterised by a shorter action potential duration (APD) compared to the normal APD recorded from ROI A.

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 samples for training and 30 samples for validation, respectively. The optical data recorded for the case of pacing from the right ventricle, RV, (see Fig. 5(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(b)), were used for testing.

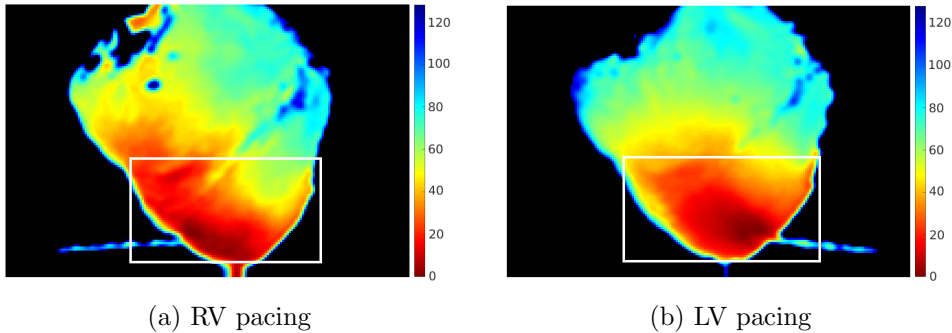


Figure 5: Example of depolarisation maps selected from the optical data recorded ex vivo in a healthy porcine heart, with the stimulating electrode (used for pacing the heart) placed onto the right ventricle (a), and left ventricle (b), respectively. Red areas correspond to early activation times (i.e., where the excitation pulse was delivered), while the late depolarisation times are depicted in green-blue, color scale is represented in ms.

3.2.2 Training settings

The training settings for this experiment were similar to the ones described in section 3.1.2, except for training horizon (300 ms).

4 Results

4.1 *In silico* data

4.1.1 0D data

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. 6), even in presence of noise in data.

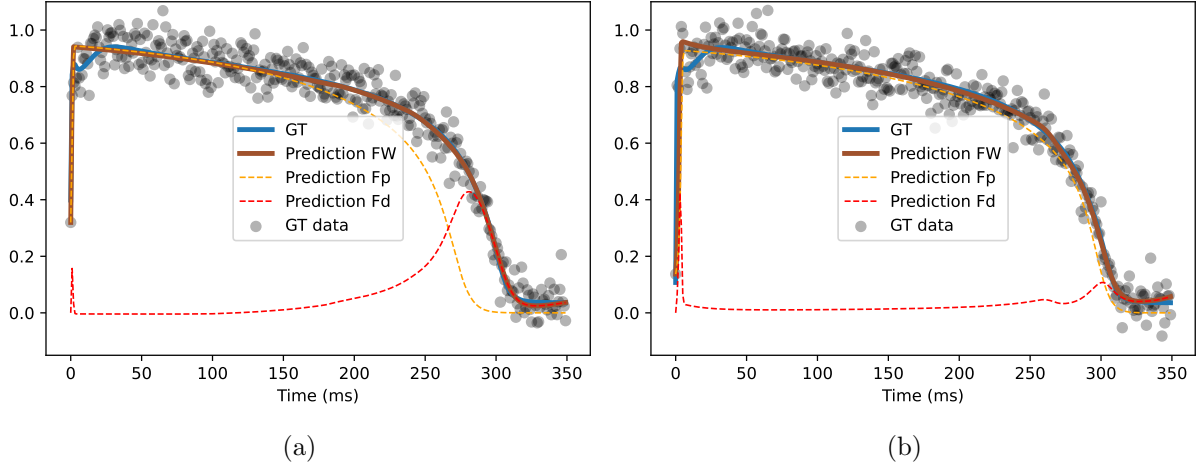


Figure 6: Validation results of the trained framework with learning of: (a) 2 (τ_{in} and τ_{out}) and (b) 3 (τ_{in} , τ_{out} and τ_{close}) physical parameters. Legend: ground truth (GT), prediction of the framework (Prediction FW), decomposition of prediction on physical (F_p) and DL (F_d) components.

Figure 6(a) shows that in 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 [24]. In the presence of learning of τ_{close} the error of dynamics corrected by data-driven component is minimal (Fig. 6(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 6.

4.1.2 2D data

We include here our qualitative results obtained for the forecast over 8 ms, after assimilating only one first frame of dynamics (see Fig. 7). These first 8 ms (i.e., the AP upstroke) represent an important part of the cardiac dynamics, ranging from the earliest depolarisation phase to the full depolarisation phase. Importantly, one can observe a very good agreement between the ground truth and the forecast transmembrane potentials generated by APHYN-EP, as illustrated in Figure 8. The effect of the correction term introduced by F_d is clearly visible. Moreover, in Figure 8(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.

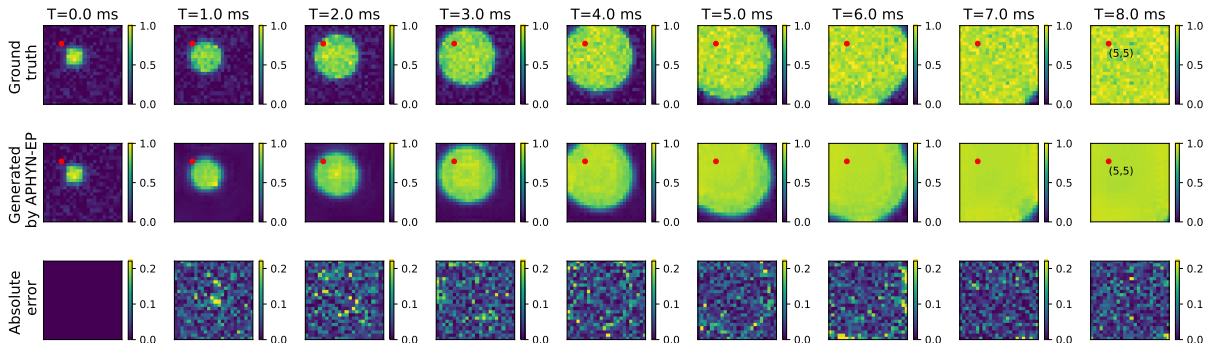


Figure 7: APHYN-EP predicted dynamics for the transmembrane potential diffusion. Figure shows a period of 8 ms of the forecast. Red point is the reference point for Figure 8.

Figures 8 and 9 present the performance of different components of APHYN-EP and associated contribution to the final result. We can observe which part of the generated transmembrane potential

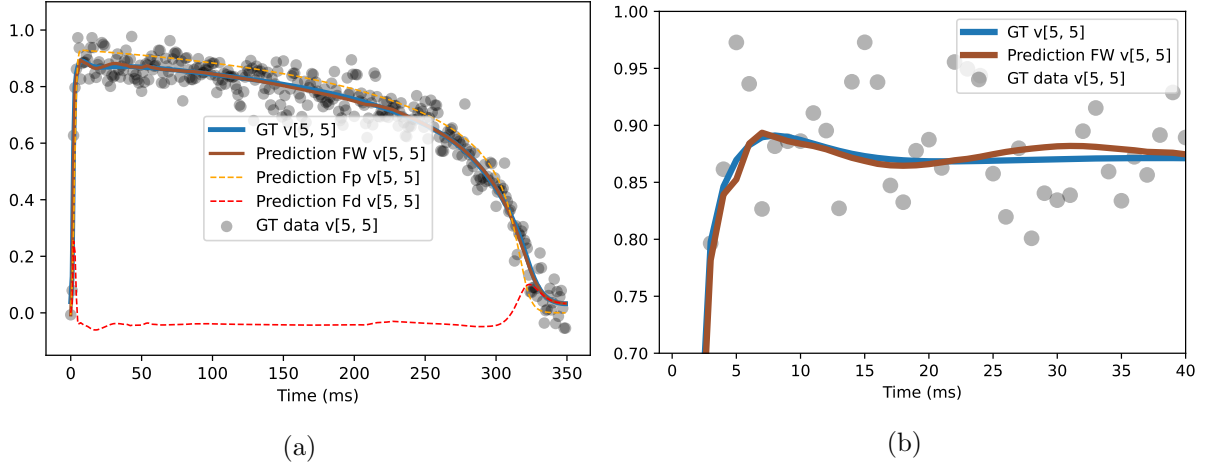


Figure 8: Transmembrane potential at point (5,5) in the cardiac slab (red point, see Fig. 7): (a) Original, (b) Zoom-in of first 40 ms. Legend: ground truth (GT), APHYN-EP, physical (F_p) and data-driven (F_d) component of APHYN-EP.

was created by the physical component of the framework (as seen in the second row of Fig. 9). The data-driven component was used only to correct the difference between the ground-truth dynamics and the physical part (Fig. 9 (third and fourth rows)).

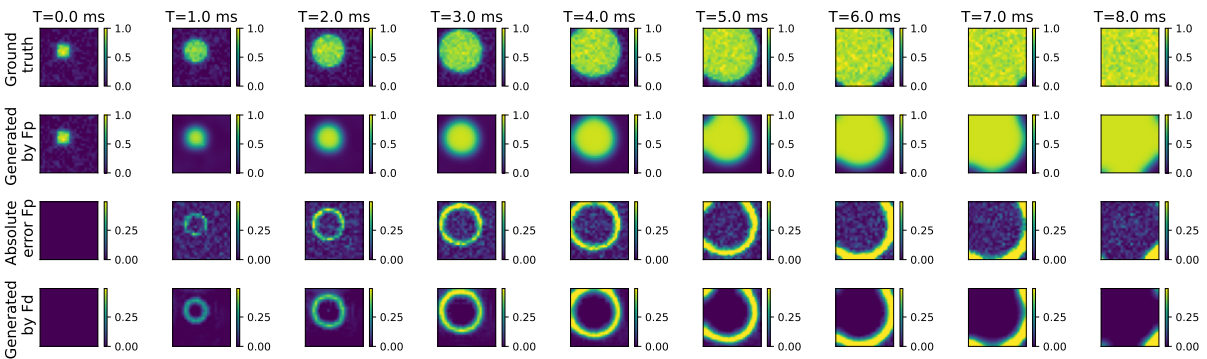


Figure 9: Exemplary results illustrating: the dynamics of the transmembrane potential diffusion predicted by the APHYN-EP physical component (second row); the error with ground-truth diffusion for this physical component of APHYN-EP (third row); and, the trained APHYN-EP data-driven component contribution (bottom row).

Table 1 presents the mean squared error (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 information. Furthermore, for comparison, we also added two baseline models: the “incomplete” physical model (F_p from APHYN-EP framework, trained alone) and a fully data-driven model (EP-Net 2.0 [19]) trained on the same dataset as APHYN-EP described in 3.1(2D). We clearly noticed that APHYN-EP captured the observed dynamics with good precision for a large time horizon (400 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.

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. 10). One can observe that APHYN-EP (trained on data from 3.1.1) has successfully generated the forecast of the new transmembrane potential wave dynamics.

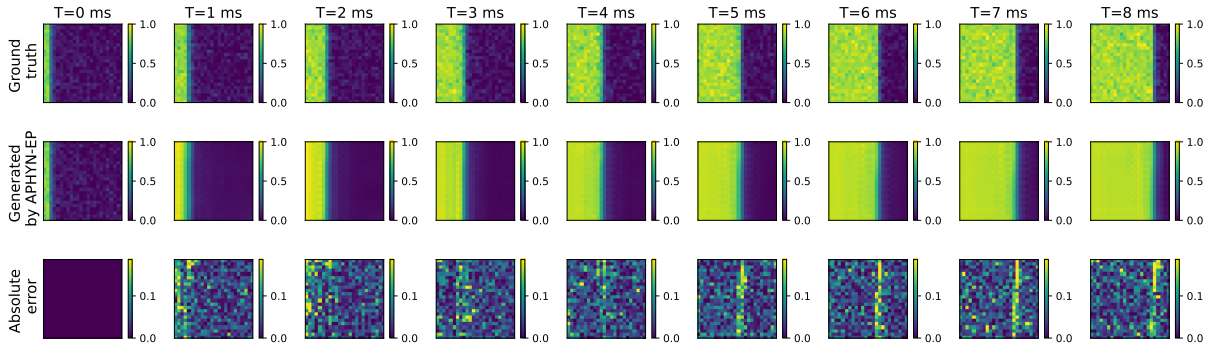


Figure 10: APHYN-EP predicted dynamics for the transmembrane potential diffusion of planar wave. The frames show a period of 8 ms of forecast obtained without re-training the APHYN-EP framework.

4.2 *Ex vivo* data

4.2.1 0D data

Using optical imaging mapping data, our APHYN-EP framework was able to reproduce the observed action potential dynamics for different ROIs within the heart, identifying the 3 major physical dynamics parameters (τ_{in} , τ_{out} and τ_{close}). Figure 11 demonstrates that the framework correctly estimated the difference in value for the parameter τ_{close} , which either increased APD or shortened it, respectively.

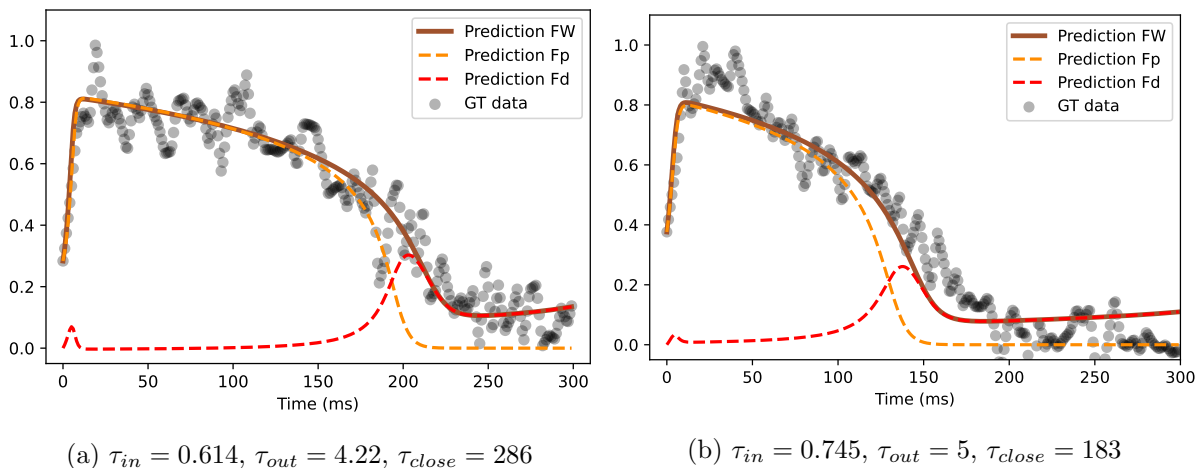


Figure 11: Validation results of the framework trained on: (a) ROI A data and (b) ROI B data, respectively. Ground truth (GT) data, prediction of the framework (Prediction FW), decomposition of prediction on physical (F_p) and DL (F_d) components.

Table 1 summarises the quantitative results for our framework forecasting on train and validation data samples, in comparison to baseline methods trained on the same data. The obtained MSE is relatively small for both ROI, 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 a good results on ROI B, the pure data-driven model encountered difficulties to learn the dynamics from ROI A data.

Generalisation ability of APHYN-EP: Fast personalisation Having noted that the obtained correction of the DL component is similar for both ROIs (Fig. 11), 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 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 the F_d component trained on ROI B data ($F_{d_B}(B)$) (and vice versa), and achieved similar results in forecasting (see Fig. 12-15).

Figures 13 and 15 visually confirm our initial hypothesis, suggesting that in order to obtain an appropriate forecast of 0D 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.

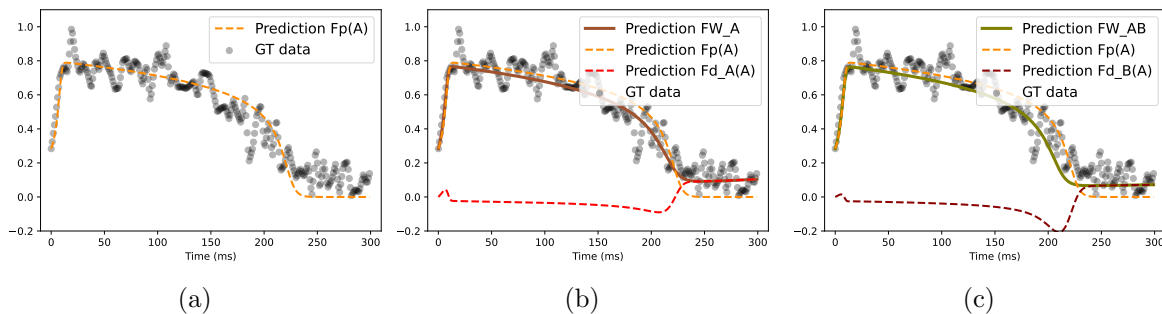


Figure 12: (a) Validation results of the F_p trained alone on ROI A data to obtain $\theta_p(A)$ parameters. (b) Validation results for APHYN-EP framework (with fixed $\theta_p(A)$ parameters) trained on ROI A data. (c) Validation results of the APHYN-EP framework (with fixed $\theta_p(A)$ parameters) 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.

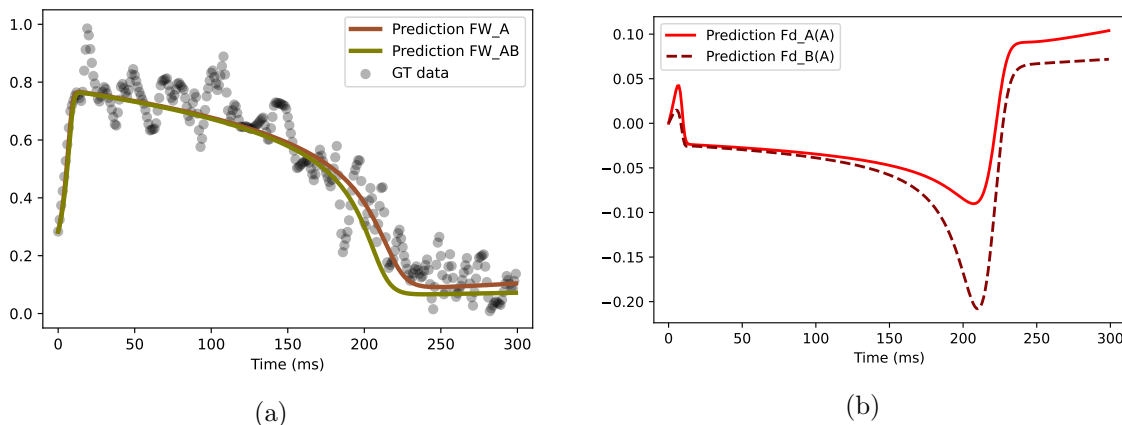


Figure 13: (a) Validation results of the framework for ROI A, with the following legend: Ground truth (GT) data, prediction of the framework (Prediction FW_A) trained on only ROI A data, and prediction of the framework composed of (F_p) component trained on ROI A data and DL (F_d) component trained on ROI B data (Prediction FW_AB). (b) Comparison of the 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.

4.2.2 2D data

Finally, we observed that APHYN-EP framework was able to reproduce the features of action potential wave from 2D optical mapping data (see Fig. 16 and Fig. 17). Table 1 shows the quantitative results for our framework forecasting on train, validation and out-of-domain test data samples in comparison

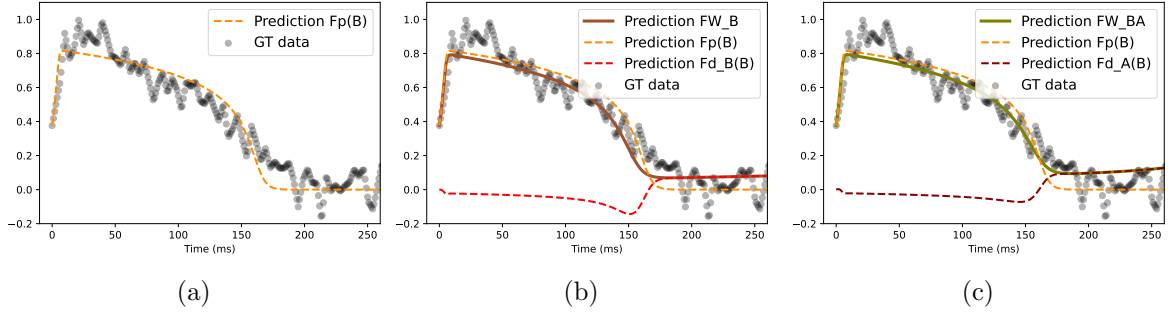


Figure 14: (a) Validation results of the F_p trained alone on ROI B data to obtain $\theta_p(B)$ parameters. (b) Validation results of the APHYN-EP framework (with fixed $\theta_p(B)$ parameters) trained on ROI B data. (c) Validation results of the APHYN-EP framework (with fixed $\theta_p(B)$ parameters) 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.

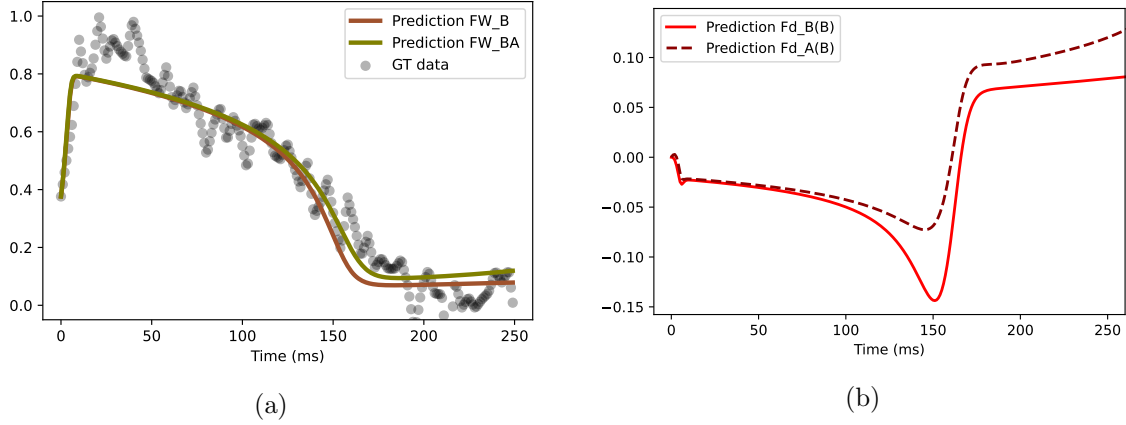


Figure 15: (a) Validation results of the framework for ROI B, with the following legend: Ground truth (GT) data, prediction of the framework (Prediction FW_B) trained on only ROI B data, and prediction of the framework composed of (F_p) component trained on ROI B data and DL (F_d) component trained on 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.

to baseline methods trained on the same data. One can notice that the Physical model and a simple Data-driven model outplay the APHYN-EP framework on training and validation datasets, but they have worse performance on a test dataset, which indicates either their possible overfitting or disability to generalise to new conditions. It is also important to note that APHYN-EP has the best results of forecasting for first 150 ms (see Table 1) and retains persistent performance on out-of-domain test data.

Generalisation ability of APHYN-EP: LV pacing As shown in Figure 17, the framework keeps the capability to generalise to unseen conditions (e.g. LV pacing). The absolute error was slightly larger than that on images obtained with RV pacing (used for training); however, this error was still acceptable. The quantitative results are provided in Table 1.

5 Discussion

In this work, we developed and tested a robust learning framework that is able to assimilate cardiac electrophysiology dynamics. Overall, our results suggest that automated learning of cardiac EP dynamics is feasible and has great potential in predicting the features of action potential wave. In particular, we

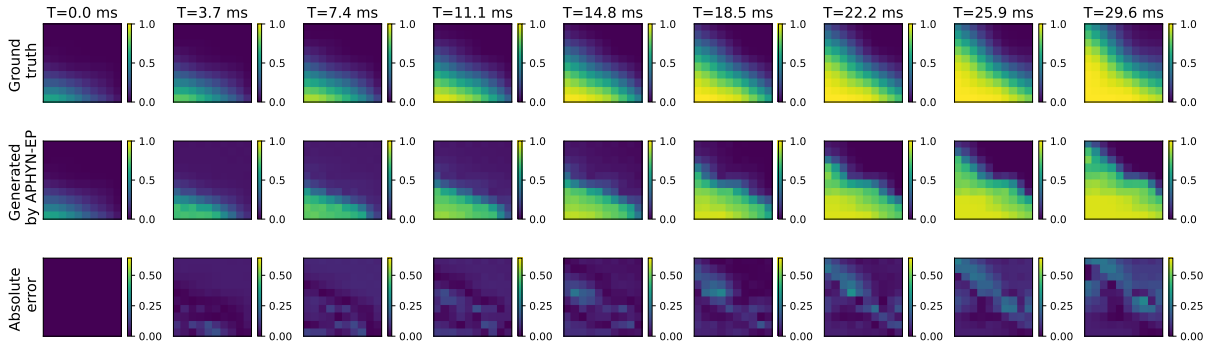


Figure 16: APHYN-EP predicted dynamics for the transmembrane potential diffusion, RV pacing. The frames illustrate time sequences from a 30 ms period of forecast.

demonstrated that the APHYN-EP framework can master the EP data of different origin (i.e, data simulated via the Ten Tusscher – Panfilov ionic model, as well as real optical mapping data), even in the presence of noise in the data. This framework may be useful for applications concerning fast parameterisation of computational models that predict electrical wave propagation through the heart.

In the section 4.2.1, we demonstrated that APHYN-EP framework can identify the major physical dynamics parameters for different heart regions. Alongside, we also demonstrated an example of fast personalisation for the framework.

Moreover, in sections 4.1.2 and 4.2.2 we illustrated the generalisation abilities of our framework, by applying the framework on new data without the need of further re-training or a change in its parameters.

However, we acknowledge that the performance of the framework depends on many hyper-parameters, such as a choice of DL component of the framework, or of the training strategy, etc. For instance, due to the data-driven architecture, our framework training could lead to a local minimum for the physical component parameters and additional involvement of the DL component. Several solutions for this problem may be given by the rigid boundaries on physical model parameters or a more strict training protocol such as separate sequential training of physical and DL components.

Furthermore, we observed that the usage of simple neural networks for the DL component of our APHYN-EP framework, is able to improve the results (see Table 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 work we have focused mostly on results generated by the framework with DL component 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 [2, 19].

6 Conclusion

In this article we successfully demonstrated the ability of a novel APHYN-EP framework to learn cardiac EP dynamics from data of different complexity. The main advantage of our proposed framework is its coupled architecture, which allowed us to use a simplified low-fidelity EP model as a physical component of the framework, along with a neural network as a data-driven correction mechanism for the EP model. Our original framework opens up several possibilities in order to introduce prior knowledge in deep learning approaches through explicit equations, as well as to correct the physical model errors from assimilated data.

Acknowledgments

The authors are grateful to the OPAL infrastructure from Université Côte d’Azur for providing resources and support and to Dr. A. Panfilov and his PhD student T. Nezlobinsky for showing us how to use their software FiniteWave to easily simulate the EP data.

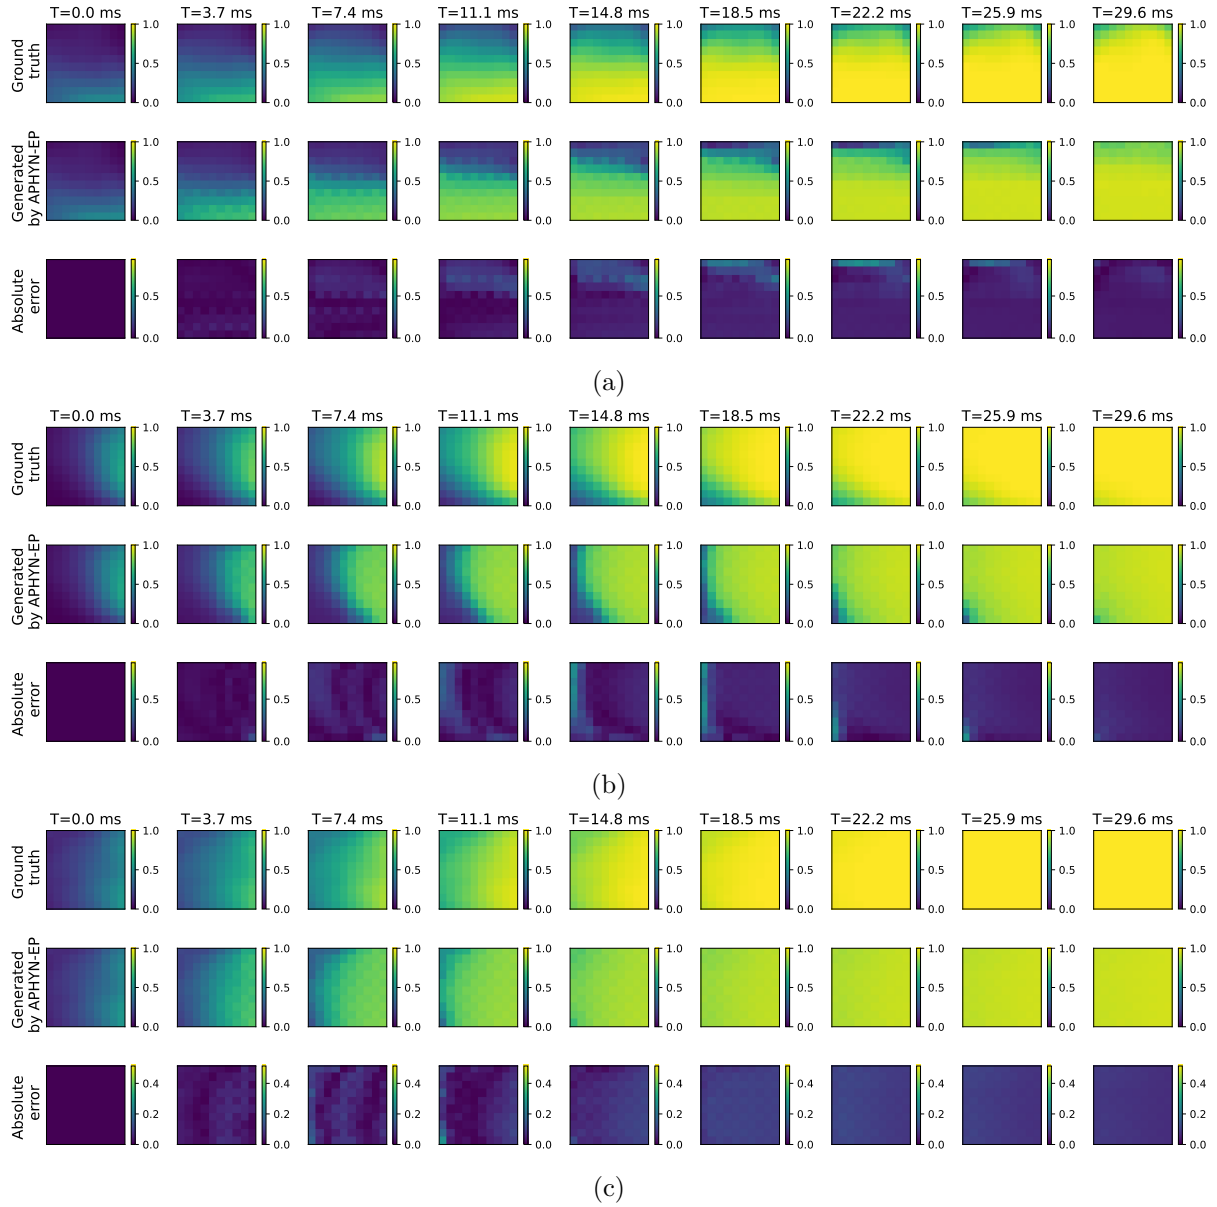


Figure 17: APHYN-EP predicted dynamics for the transmembrane potential diffusion, LV pacing. The frames show a period of 30 ms of forecast obtained without re-training the APHYN-EP framework.

References

- [1] Rubin R. Aliev and Alexander V. Panfilov. A simple two-variable model of cardiac excitation. *Chaos, Solitons & Fractals*, 7(3):293–301, 1996.
- [2] 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*, pages 55–63. Springer, 2019.
- [3] Ibrahim Ayed, Emmanuel de Bézenac, Arthur Pajot, Julien Brajard, and Patrick Gallinari. Learning dynamical systems from partial observations. *arXiv preprint:1902.11136*, 2019.
- [4] Jieyun Bai, Jichao Zhao, Haibo Ni, and Dechun Yin. Editorial: Diagnosis, monitoring, and treatment of heart rhythm: new insights and novel computational methods. *Front Physiol*, 14:1272377, 2023.
- [5] O. Camara, M. Sermesant, P. Lamata, L. Wang, M. Pop, J. Relan, M. De Craene, H. Delingette,

Table 1: Mean-squared error, MSE ($\times 10^{-3}$) of the normalised transmembrane potential (adimensional) forecasting (forecasting horizon of 400 ms for *In silico* experiments and of 300 ms for *Ex vivo* experiments). Baseline models: the Physical model (4) and a fully data-driven model (EP-Net 2.0 [19]) trained on the same dataset as APHYN-EP. Out-of-domain tests: Planar wave for *In silico* 2D dataset, LV pacing for *Ex vivo* 2D dataset.

Dataset		Method	Training data ($\ F_d\ ^2$)	Validation data ($\ F_d\ ^2$)	Out-of-domain test ($\ F_d\ ^2$)		
<i>In silico</i>	2D	APHYN-EP ResNet	2.54 (0.47)	2.54 (0.472)	4.2 (0.4)		
		APHYN-EP ConvNet	2.5 (0.81)	2.5 (0.8)	4.4 (1)		
		Physical model	5.7	5.6	4.3		
		Data-driven model	10	10	100		
<i>Ex vivo</i>	0D	ROI A	APHYN-EP ResNet	9.12 (0.16)	5.72 (0.08)		
			APHYN-EP MLP	9 (0.0785)	5.37 (0.077)		
		ROI B	Physical model	14	10		
			Data-driven model	20	9.78		
	2D	RV	APHYN-EP ResNet	54.5 (92.6)	53.26 (94.2)		49.9 (134.6)
			APHYN-EP ConvNet	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 150	APHYN-EP ResNet	8.9 (110.85)	7.66 (109.9)		8.93 (135.8)
			APHYN-EP ConvNet	8.26 (101.16)	6.93 (107.5)		7.47 (124.4)
			Physical model	20	18.42		32.5
			Data-driven model	16.84	15.23		14.9

- 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. *Progress in Biophysics and Molecular Biology*, 107(1):122–133, 2011. Experimental and Computational Model Interactions in Bio-Research: State of the Art.
- [6] Ricky T. Q. Chen, Brandon Amos, and Maximilian Nickel. Learning neural event functions for ordinary differential equations. *ICRL*, 2021.
- [7] Ricky T. Q. Chen, Yulia Rubanova, Jesse Bettencourt, and David Duvenaud. Neural ordinary differential equations. *Adv Neural Inform Process Syst*, 2018.
- [8] Richard H Clayton, Yasser Aboelkassam, Chris D Cantwell, Cesare Corrado, Tammo Delhaas, Wouter Huberts, Chon Lok Lei, Haibo Ni, Alexander V Panfilov, Caroline Roney, and Rodrigo Weber Dos Santos. An audit of uncertainty in multi-scale cardiac electrophysiology models. *Philos Trans A Math Phys Eng Sci*, 378(2173):20190335, 2020.
- [9] Sebastien Court and Karl Kunisch. Design of the monodomain model by artificial neural networks. *arXiv preprint:2107.03136*, 2021.
- [10] Anouk G W de Lepper, Carlijn M A Buck, Marcel van 't Veer, Wouter Huberts, Frans N van de Vosse, and Lukas R C Dekker. From evidence-based medicine to digital twin technology for predicting ventricular tachycardia in ischaemic cardiomyopathy. *J R Soc Interface*, 19(194):20220317, 2022.
- [11] Richard FitzHugh. Impulses and physiological states in theoretical models of nerve membrane. *Biophysical Journal*, 1(6):445–466, 1961.
- [12] 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. *Front. Physiol.*, 12, 2021.

- [13] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *IEEE conf. CVPR*, pages 770–778, 2016.
- [14] 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. *Front. Cardiovasc. Med.*, 8, 2022.
- [15] 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.
- [16] Sergey I. Kabanikhin. *Inverse and Ill-posed Problems: Theory and Applications*. De Gruyter, Berlin, Boston, 2011.
- [17] George Em Karniadakis, Ioannis G. Kevrekidis, Lu Lu, Paris Perdikaris, Sifan Wang, and Liu Yang. Physics-informed machine learning. *Nature Reviews Physics*, 3:422–440, 2021.
- [18] Victoriya Kashtanova, Ibrahim Ayed, Andony Arrieula, Mark Potse, Patrick Gallinari, and Maxime Sermesant. Deep learning for model correction in cardiac electrophysiological imaging. In *Medical Imaging with Deep Learning*, 2022.
- [19] 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*, volume 12738 of *Lecture Notes in Computer Science*, pages 482–492. Springer International Publishing, 2021.
- [20] 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 Modeling of the Heart workshop*, 2022.
- [21] Zichao Long, Yiping Lu, and Bin Dong. PDE-net 2.0: Learning PDEs from data with a numeric-symbolic hybrid deep network. *J. Comput. Phys.*, 399:108925, 2019.
- [22] Zichao Long, Yiping Lu, Xianzhong Ma, and Bin Dong. PDE-net: Learning PDEs from data. In *Int. Conf. ICML*, pages 3208–3216. PMLR, 2018.
- [23] Alejandro Lopez-Perez, Rafael Sebastian, and Jose M Ferrero. Three-dimensional cardiac computational modelling: methods, features and applications. *Biomed Eng Online*, 14:35, 2015.
- [24] Colleen C Mitchell and David G Schaeffer. A two-current model for the dynamics of cardiac membrane. *Bull. Math. Biol.*, 65(5):767–793, 2003.
- [25] J. Nagumo, S. Arimoto, and S. Yoshizawa. An active pulse transmission line simulating nerve axon. *Proceedings of the IRE*, 50(10):2061–2070, 1962.
- [26] Martyn P. Nash and Alexander V. Panfilov. Electromechanical model of excitable tissue to study reentrant cardiac arrhythmias. *Prog. Biophys. Mol. Biol.*, 85(2):501–522, 2004.
- [27] 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. *Adv. Neural Inf. Process. Syst.*, 32, 2019.
- [28] 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. *Medical Image Analysis*, 13(2):370–380, 2009. Includes Special Section on Functional Imaging and Modelling of the Heart.
- [29] 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. *Journal of Computational Physics*, 378:686–707, 2019.

- [30] Maziar Raissi and George Em Karniadakis. Hidden physics models: Machine learning of nonlinear partial differential equations. *Journal of Computational Physics*, 357:125–141, 2018.
- [31] Jatin Relan, Phani Chinchapatnam, Maxime Sermesant, Kawal Rhode, Matt Ginks, Hervé Delingette, C Aldo Rinaldi, Reza Razavi, and Nicholas Ayache. Coupled personalization of cardiac electrophysiology models for prediction of ischaemic ventricular tachycardia. *Interface Focus*, 1(3):396–407, 2011.
- [32] 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. *Engineering with Computers*, 38(5):3957–3973, 2022.
- [33] Francisco Sahli Costabal, Yibo Yang, Paris Perdikaris, Daniel E. Hurtado, and Ellen Kuhl. Physics-informed neural networks for cardiac activation mapping. *Front. Phys.*, 8:42, 2020.
- [34] K. H. W. J. Ten Tusscher, D. Noble, P. J. Noble, and A. V. Panfilov. A model for human ventricular tissue. *Am. J. Physiol. - Heart Circ. Physiol.*, 286:H1573–H1589, 2004.
- [35] K. H. W. J. Ten Tusscher and A. V. Panfilov. Alternans and spiral breakup in a human ventricular tissue model. *Am. J. Physiol. - Heart Circ. Physiol.*, 291(3):H1088–H1100, 2006.
- [36] Rui Wang, Robin Walters, and Rose Yu. Incorporating symmetry into deep dynamics models for improved generalization. In *International Conference on Learning Representations*, 2021.
- [37] Rui Wang, Robin Walters, and Rose Yu. Approximately equivariant networks for imperfectly symmetric dynamics. In *International Conference on Machine Learning*, pages 23078–23091. PMLR, 2022.
- [38] Rui Wang and Rose Yu. Physics-Guided Deep Learning for Dynamical Systems: A Survey. 2022.
- [39] Jared Willard, Xiaowei Jia, Shaoming Xu, Michael Steinbach, and Vipin Kumar. Integrating Scientific Knowledge with Machine Learning for Engineering and Environmental Systems. *ACM Computing Surveys*, 1(1):1–35, 2022.
- [40] Jared D. Willard, Xiaowei Jia, Shaoming Xu, Michael Steinbach, and Vipin Kumar. Integrating physics-based modeling with machine learning: A survey. *arXiv preprint:2003.04919*, 2020.
- [41] Liu Yang, Xuhui Meng, and George Em Karniadakis. B-pinns: Bayesian physics-informed neural networks for forward and inverse pde problems with noisy data. *Journal of Computational Physics*, 425:109913, 2021.
- [42] 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.