Physics-based deep learning framework to learn and forecast cardiac electrophysiology dynamics

Victoriya Kashtanova [1] **Patrick Gallinari** [2] **Maxime Sermesant** [1]

Abstract

Biophysically detailed mathematical modeling of cardiac electrophysiology is often computationally demanding, for example, when solving problems for various patient pathological conditions. Furthermore, it is still difficult to reduce the discrepancy between the output of idealised mathematical models and clinical measurements, which are usually noisy. In this work, we propose a fast physics-based deep learning framework to learn complex cardiac electrophysiology 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 data, we demonstrate that this framework can reproduce the complex dynamics of transmembrane potential, even in presence of noise in the data. This combined physics-based data-driven approach may improve cardiac electrophysiology modeling by providing a robust biophysical tool for predictions.

1. Introduction

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 routinely employed to diagnose arrhythmia, as well as to help planning the best therapy on an individual basis. 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.

Despite the fact that biophysically detailed cardiac electrophysiology (EP) models (such as (Ten Tusscher et al., 2004)) can accurately reproduce electrical behaviour of cardiac cells, these models are complex and computationally expensive, and have many hidden variables which are impossible to measure, making model parameters difficult to personalise. The phenomenological models (FitzHugh, 1961; Nagumo et al., 1962; Aliev & Panfilov, 1996; Nash & Panfilov, 2004; Mitchell & Schaeffer, 2003), simplified models derived from biophysical models, have fewer parameters and are therefore especially useful for rapid computational modelling of wave propagation at the organ level. However, they are less realistic and therefore need a complementary mechanism to fit them to the measured data. Machine learning and in particular deep learning (DL) approaches could help providing such a correction mechanism. The combination of rapid phenomenological models and machine learning components could then allow the development of rapid and accurate models of transmembrane dynamics (as in (Herrero Martin et al., 2022; Fresca et al., 2021; Sahli Costabal et al., 2020)). Nevertheless, the majority of existing 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.

To address this critical limitation, here we propose a *Our framework's name* framework to learn complex cardiac electrophysiology dynamics, 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. The proposed framework closely follows the approach of Yin et al. (2021). However, in contrast to

[*]Equal contribution [1]Inria d'Université Côte d'Azur, Sophia Antipolis, France [2]Sorbonne University, Criteo AI Lab, Paris, France. Correspondence to: Victoriya Kashtanova <victoriya.kashtanova@inria.fr>, Maxime Sermesant <maxime.sermesant@inria.fr>.

Accepted after peer-review at the 1st workshop on Synergy of Scientific and Machine Learning Modeling, SynS & ML ICML, Honolulu, Hawaii, USA. July, 2023.

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.

Fig. 1 presents the general framework of our approach. 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.

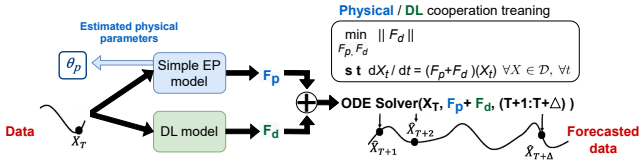

Figure 1. General *Our framework's name* framework scheme. During the training phase two-component framework learn the parameters for the physical (F_p) and the data-driven (F_d) components from data. Then via an ODE solver the framework can forecast further the learned dynamics.

2. Learning Framework

In order to learn the cardiac EP dynamics (X_t), in this work we solved an optimisation problem via our physics-based data-driven *Our framework's name* 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:

$$\min_{F_p \in \mathcal{F}_p, F_d \in \mathcal{F}_d} \|F_d\| \text{ subject to } \\ \forall X \in \mathcal{D}, \forall t, \frac{dX_t}{dt} = F(X_t) = (F_p + F_d)(X_t). \tag{1}$$

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

Specifically, our incomplete physical model is the two-variable (v, h) model (Mitchell & Schaeffer, 2003) for cardiac EP simulations, as described by equations (2). 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):

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

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 (Relan et al., 2011), 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 (2) $v(t = 0) = 0$ 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 work, we used a ResNet network (He et al., 2016), because it could accurately reproduce complex cardiac EP dynamics (Ayed et al., 2019; Kashtanova et al., 2021). However, a simpler neural network could also be used for more rapid computations.

In *Our framework's name* framework the physical (F_p) and the data-driven (F_d) components are trained simultaneously, using automatic differentiation tools provided by the Pytorch library (Paszke et al., 2019). This insures the finding of the best minimising pair for (1) determined by the 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:

$$\mathcal{L}(\theta) = \lambda * \mathcal{L}_{traj}(\theta) + \left\| F_d^{\theta_d} \right\| \\ = \lambda * \sum_{i=1}^{N} \sum_{h=1}^{T/\Delta t} ||X_{h\Delta t}^{(i)} - \tilde{X}_{h\Delta t}^{(i)}(\theta)|| + \left\| F_d^{\theta_d} \right\| \tag{3}$$

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 et al., 2018; 2021). The *Our framework's name* training uses algorithm similar to Yin et al. (2021).

Additionally, in order to train simultaneously the physical and the data-driven components of *Our framework's name*, we implemented the Laplace operator in (2) with a simple finite-difference scheme.

3. Experiments and Results

In order to test the performance of our *Our framework's name* framework and to further show its capability to reproduce complex transmembrane potential dynamics simulated via a biophysically detailed cardiac EP model, we chose two types of experiments. First, we tested the ability of the framework to learn the complex dynamics of transmembrane potential including a case where noise is present in the data. Second, using test data samples, we showed that our framework is able to generalise to new conditions, outside of its training environment.

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

3.1. 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 (Ten Tusscher & Panfilov, 2006), 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 (Ten Tusscher & Panfilov, 2006)), with the open-source finetwave software[1]. 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.

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 *Our framework's name* 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

3.2. Results

We include here our qualitative results obtained for the forecast over 8 ms, after assimilating only one first frame of dynamics (see Fig. 2). 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. Importantly, one can observe a very good agreement between the ground truth and the forecast transmembrane potentials generated by *Our framework's name*, as illustrated in Figure 3. The effect of the correction term introduced by F_d is clearly visible. Moreover, in Figure 3(b), one can also observe that *Our framework's name* framework achieves a good precision in transmembrane potential forecasting even when noise is present in the data.

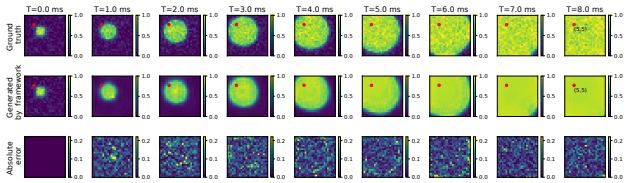

Figure 2. Our framework's name predicted dynamics for the transmembrane potential diffusion. Figure shows a period of 8 ms of the forecast. Red point is the reference point for Figure 3.

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 *Our framework's name* framework, trained alone) and a fully data-driven model (EP-Net 2.0 (Kashtanova et al., 2021)) trained on the same dataset as *Our framework's name* described in 3.1. We clearly noticed that *Our framework's name* 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 *Our framework's name*: 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. 4). One can observe that *Our framework's name* (trained on data from 3.1) has successfully generated the forecast of the new transmembrane potential wave dynamics.

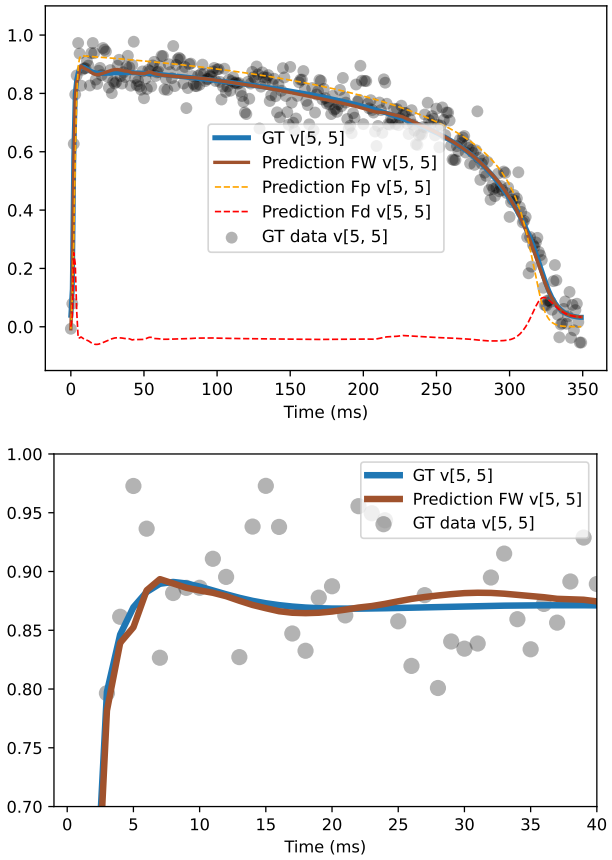

Figure 3. Transmembrane potential at point (5,5) in the cardiac slab (red point, see Fig. 2): (top) Original, (bottom) Zoom-in of first 40 ms. Legend: ground truth (GT), *Our framework's name*, physical (F_p) and data-driven (F_d) component of *Our framework's name*.

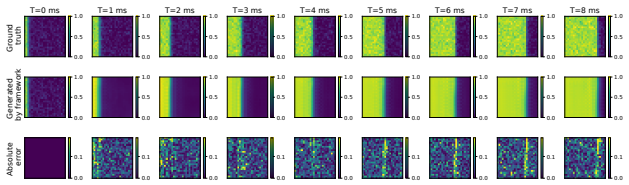

Figure 4. Our framework's name predicted dynamics for the transmembrane potential diffusion of plannar wave. The frames show a period of 8 ms of forecast obtained without re-training the *Our framework's name* framework.

4. Conclusion

We have presented the *Our framework's name* framework for modeling complex cardiac electrophysiology dynamics via a surrogate model combining simplified physics and deep neural network. We demonstrated that this framework is able to reproduce with good precision the dynamics simulated by the Ten Tusscher – Noble – Noble – Panfilov ionic model, even using a simplified electrophysiology model as a physical component of the framework. Such framework opens up possibilities in order to introduce prior knowledge in deep learning approaches through explicit equations and to correct model errors from data.

Our current work is the evaluations of this framework on more challenging settings (such as, presence of anisotropic depolarisation wave propagation and various conduction velocities in the cardiac tissue slab, as well as, the real data applications), which were left out of scope in this paper.

Table 1. Mean-squared error, MSE (x 10^{-3}) of the normalised transmembrane potential (adimensional) forecasting (forecasting horizon of 400 ms). Baseline models: the Physical model (2) and a fully data-driven model (EP-Net 2.0 (Kashtanova et al., 2021)) trained on the same dataset as *Our framework's name*. Out-of-domain tests: Plannar wave.

Method	**Training data**	**Validation data**	**Out-of-domain test**
Our framework's name framework with ResNet ($\|\|F_d\|\|^2$)	2.54 (0.47)	2.54 (0.472)	4.2 (0.4)
Our framework's name framework with ConvNet ($\|\|F_d\|\|^2$)	2.5 (0.81)	2.5 (0.8)	4.4 (1)
Physical model	5.7	5.6	4.6
Data-driven model	10	10	100

Broader impact

Modelling complex systems like the human heart has made great progress over the last decades. However, it is difficult to improve the realism of such models even with detailed measurements, as it requires to modify complex sets of equations in order to change their dynamics as well as to find sets of parameters that enable relevant simulations. In this work, we presented a fast physics-based deep learning framework to learn complex cardiac electrophysiology dynamics from data. Its coupled architecture allows the training on data of

different complexity and origin. Therefore, this 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. That may improve cardiac electrophysiology modeling by providing a robust biophysical tool for predictions.

References

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

Ayed, I., Cedilnik, N., Gallinari, P., and Sermesant, M. EP-Net: Learning cardiac electrophysiology models for physiology-based constraints in data-driven predictions. In *Int. Conf. FIMH*, pp. 55–63. Springer, 2019.

Chen, R. T. Q., Rubanova, Y., Bettencourt, J., and Duvenaud, D. Neural ordinary differential equations. *Adv Neural Inform Process Syst*, 2018.

Chen, R. T. Q., Amos, B., and Nickel, M. Learning neural event functions for ordinary differential equations. *ICRL*, 2021.

FitzHugh, R. Impulses and physiological states in theoretical models of nerve membrane. *Biophysical Journal*, 1 (6):445–466, 1961.

Fresca, S., Manzoni, A., Dedè, L., and Quarteroni, A. POD-enhanced deep learning-based reduced order models for the real-time simulation of cardiac electrophysiology in the left atrium. *Front. Physiol.*, 12, 2021.

He, K., Zhang, X., Ren, S., and Sun, J. Deep residual learning for image recognition. In *IEEE conf. CVPR*, pp. 770–778, 2016.

Herrero Martin, C., Oved, A., Chowdhury, R. A., Ullmann, E., Peters, N. S., Bharath, A. A., and Varela, M. EP-PINNs: Cardiac electrophysiology characterisation using physics-informed neural networks. *Front. Cardiovasc. Med.*, 8, 2022.

Kashtanova, V., Ayed, I., Cedilnik, N., Gallinari, P., and 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*, pp. 482–492. Springer International Publishing, 2021.

Mitchell, C. C. and Schaeffer, D. G. A two-current model for the dynamics of cardiac membrane. *Bull. Math. Biol*, 65(5):767–793, 2003.

Nagumo, J., Arimoto, S., and Yoshizawa, S. An active pulse transmission line simulating nerve axon. *Proceedings of the IRE*, 50(10):2061–2070, 1962.

Nash, M. P. and Panfilov, A. V. Electromechanical model of excitable tissue to study reentrant cardiac arrhythmias. *Prog. Biophys. Mol. Biol.*, 85(2):501–522, 2004.

Paszke, A., Gross, S., Massa, F., Lerer, A., Bradbury, J., Chanan, G., Killeen, T., Lin, Z., Gimelshein, N., Antiga, L., et al. Pytorch: An imperative style, high-performance deep learning library. *Adv. Neural Inf. Process. Syst.*, 32, 2019.

Relan, J., Chinchapatnam, P., Sermesant, M., Rhode, K., Ginks, M., Delingette, H., Rinaldi, C. A., Razavi, R., and Ayache, N. Coupled personalization of cardiac electrophysiology models for prediction of ischaemic ventricular tachycardia. *Interface Focus*, 1(3):396–407, 2011.

Sahli Costabal, F., Yang, Y., Perdikaris, P., Hurtado, D. E., and Kuhl, E. Physics-informed neural networks for cardiac activation mapping. *Front. Phys.*, 8:42, 2020.

Ten Tusscher, K. H. W. J. and Panfilov, A. V. Alternans and spiral breakup in a human ventricular tissue model. *Am. J. Physiol. - Heart Circ. Physiol.*, 291(3):H1088–H1100, 2006.

Ten Tusscher, K. H. W. J., Noble, D., Noble, P. J., and Panfilov, A. V. A model for human ventricular tissue. *Am. J. Physiol. - Heart Circ. Physiol.*, 286:H1573–H1589, 2004.

Yin, Y., Le Guen, V., Dona, J., Ayed, I., de Bezenac, E., Thome, N., and Gallinari, P. Augmenting physical models with deep networks for complex dynamics forecasting. In *Int. Conf. ICRL*, 2021.