



HAL
open science

EP-Net 2.0: Out-of-Domain Generalisation for Deep Learning Models of Cardiac Electrophysiology

Victoriya Kashtanova, Ibrahim Ayed, Nicolas Cedilnik, Patrick Gallinari,
Maxime Sermesant

► **To cite this version:**

Victoriya Kashtanova, Ibrahim Ayed, Nicolas Cedilnik, Patrick Gallinari, Maxime Sermesant. EP-Net 2.0: Out-of-Domain Generalisation for Deep Learning Models of Cardiac Electrophysiology. FIMH 2021 - 11th International Conference on Functional Imaging and Modeling of the Heart, Jun 2021, Stanford, CA (virtual), United States. pp.482-492, 10.1007/978-3-030-78710-3_46 . hal-03369201

HAL Id: hal-03369201

<https://inria.hal.science/hal-03369201v1>

Submitted on 7 Oct 2021

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.

EP-Net 2.0: Out-of-Domain Generalisation for Deep Learning Models of Cardiac Electrophysiology

Victoriya Kashtanova^{1,2}, Ibrahim Ayed^{3,5}, Nicolas Cedilnik¹, Patrick Gallinari^{3,4}, and Maxime Sermesant^{1,2,*}

¹ Inria, Université Côte d’Azur, Nice

² 3IA Côte d’Azur, Sophia Antipolis

³ Sorbonne University, LIP6, Paris

⁴ Criteo AI Lab, Paris

⁵ Theresis lab, Thales, France

`maxime.sermesant@inria.fr`

Abstract. Cardiac electrophysiology models achieved good progress in simulating cardiac electrical activity. However, it is still challenging to leverage clinical measurements due to the discrepancy between idealised models and patient-specific conditions. In the last few years, data-driven machine learning methods have been actively used to learn dynamics and physical model parameters from data. In this paper, we propose a principled deep learning approach to learn the cardiac electrophysiology dynamics from data in the presence of scars in the cardiac tissue slab. We demonstrate that this technique is indeed able to reproduce the transmembrane potential dynamics in situations close to the training context. We then focus on evaluating the ability of the trained networks to generalize outside their training domain. We show experimentally that our model is able to generalize to new conditions including more complex scar geometries, multiple signal onsets and various conduction velocities.

Keywords: Electrophysiology · Deep learning · Simulation.

1 Introduction

Mathematical modelling of the heart has been an active research area for the last decades, and it is now more and more coupled with artificial intelligence approaches, see for instance [12]. Among the multi-physics phenomena involved in the cardiac function, cardiac electrophysiology 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 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: [14] gives a thorough introduction to the closely related field of Nonlinear System Identification while [6] gives an early example

of such endeavours. More recently, those questions have seen a renewed interest with works such as [16,15,20] proposing to use Deep Neural Network models for solving differential equations while [1,19,22] use 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: [11,10] endow neural layers with additional structure, useful for learning PDEs while [5,2] use the adjoint method to learn differential equations parametrised with neural models and learn them in fully and partially observable settings. More generally, [21] propose a broader survey of ML in physics-based modeling.

We propose a framework for learning cardiac electrophysiology dynamics from data and we experimentally evaluate its ability to forecast cardiac dynamics on new conditions, unseen during training. Our models are trained and evaluated using data simulated from an electrophysiology model [13]. This is a classical experimental setting in the domain [7] which, although offering a simplification over real cardiac data, allows us to assess our framework using controlled conditions. This work builds on initial results [2,3] that were evaluated in an idealised setting with simple boundary conditions corresponding to a healthy slab of cardiac tissue with a uniform conductivity and a single onset of transmembrane potential.



Fig. 1: Example of transmembrane potential (yellow) propagation in the cardiac tissue slab in absence (a) and presence (b) of scar tissue, through successive time steps.

Our extension here introduces more complex conditions. First, we consider 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 they are the main cause of cardiac arrhythmias. For example, in Fig. 1 we can clearly see the changes in the dynamics of the depolarisation wave in the presence of scar tissue (black area). Second, we introduce in our simulations multiple onsets and colliding fronts, as it is a classical situation in cardiac electrophysiology. The focus of the paper is on the evaluation of the ability of our model to generalise to unseen conditions. The model is then trained on simulated data corresponding to relatively simple context (one type of scar, one front and set of several discrete conduction velocities) and its generalisation ability is evaluated on more challenging contexts like more complex scars, multiple fronts and any real conduction velocity sampled from a given interval. Fig. 2 presents the general experimental setting used in the manuscript.




Fig. 2: General setting used throughout the manuscript. Once trained, the EP-Net 2.0 model takes as input a context consisting of a few (4 here) observations plus an indication of the scar area (left), and forecasts the depolarization wave dynamics (right).

2 Problem Formulation and Learning Framework

We used the Mitchell–Schaeffer model [13] for cardiac electrophysiology simulation. This two variables model has been successfully used in patient-specific modelling [18]. The variable v in equation 1 represents normalised ($v \in [0, 1]$) dimensionless transmembrane potential while the "gating" variable h controls the repolarisation:

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

In practice, since h is a hidden variable, it is difficult to measure. Only the measurement of the potential v is available. Therefore, as in [3], we modify the system (1) by replacing variable h with an observation operator \mathcal{H} which extracts the corresponding information from the current state X_t . This allows us to rewrite this model in a vector form:

$$\begin{cases} X_0 = g_\theta(V_{-k}) \\ \frac{dX_t}{dt} = F_\theta(X_t) \\ V_t = \mathcal{H}(X_t) \end{cases} \quad (2)$$

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 the sequence of past observations of transmembrane potential.

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

$$\Omega_{scar} \subset \Omega \subset \mathbb{R}^2 : (X_t)_{\Omega_{scar}} \equiv 0 \quad (3)$$

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

In order to enforce the constraints (2) and (3), we compare the sequence of observations V_t generated by the parameterised model to data simulated from the actual equations and minimise the following loss:

$$\mathcal{L}(V, \tilde{V}) = \mathcal{L}_{obs}(V, \tilde{V}) + \lambda_{scar} \mathcal{L}_{scar}(\tilde{V}), \quad (4)$$

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

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 as:

$$\begin{aligned} & \underset{\theta}{\text{minimize}} && \mathbb{E}_{V \in \text{Dataset}} \mathcal{L}(V, \mathcal{H}(X^\theta)) \\ & \text{subject to} && \frac{dX_t}{dt} = F_\theta(X_t), \\ & && X_0 = g_\theta(V_{-k}) \end{aligned} \quad (5)$$

Learning Method Operators F, g in problem (5) are implemented via Deep Neural Networks. We chose to use a ResNets [8] (illustrated in Fig. 3) to parameterise both F and g . Optimisation is performed via stochastic gradient descent, precisely ADAM algorithm [9], according to the following algorithm:

0. Randomly initialise θ (denoting the parameters of F and g);
1. Solve the forward state equation 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. 3: The ResNet architecture used in EP-Net 2.0. It has 5 input frames (mask frame plus 4 frames of transmembrane potential) and 1 output frame of forecast.

3 Experiments

Data collection. We generated 2D data frames using the Lattice Boltzmann method to solve the EP model [17] on a Cartesian grid. The Mitchell-Schaeffer model parameters are taken as in the original paper [13]: $\tau_{in} = 0.3$, $\tau_{out} = 6$, $\tau_{open} = 120$, $\tau_{close} = 150$, $v_{gate} = 0.13$. The computational domain represents a slab of cardiac tissue of size $24 \times 24 \text{ mm}^2$ discretised with 1 mm^2 pixels. A stimulation current was applied for 10 ms to initiate the propagation in selected pixels (J_{stim}). We superposed a mask with randomly generated rectangular area (with random size and position) of excluded domain to simulate the cardiac scars. Training was performed with 5 different discrete conduction velocities, each corresponding to a given parameter of conductivity ($\sigma = 1, 2, 3, 4, 5$) in equation 1. The conductivity was applied uniformly on whole cardiac slab except scar area. One value of σ is used per 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 have a database of 30000 training and 12000 validation samples.

Training settings. Parameter λ_{scar} in loss (4) was set to 0.1 and a learning rate for ADAM optimiser was set to 10^{-3} . We use ResNet with 64 filters at the initial stage, three downsampling initial layers and three intermediary blocks and start with a reweighted orthogonal initialisation for its parameters. We also use exponential scheduled sampling [4] 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 (6 frames) to compute the losses.

3.1 Results

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

Testing environment similar to the training one: scars of rectangular shape. Figure 4 illustrate 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. The figure 4a shows the forecast over 9 time frames (9 ms) after assimilating the first 4 frames (not presented in the Fig. 4a). We observe 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 early depolarisation to full depolarisation. Figure 4b shows that EP-Net 2.0 model has a very good precision on depolarization during more than 50 ms, an equilibrium state for the model, but cannot predict a repolarisation (Fig. 4c). Quantitative results provided in the table 1 for different forecasting

horizons T (6, 12 and 24 ms) show excellent performance, while the training time horizon was only 6 ms.




Fig. 4: (a) Results of trained EP-Net 2.0 model (9 ms of forecast, cardiac slab conductivity $\sigma = 2$). (b,c) Transmembrane potential graph at the leftmost upper point (0,0) and the rightmost bottom point (23,23) in the slab with different forecasting horizons.

Generalisation ability of EP-Net 2.0: scars of various shapes and multiple onsets. Our objective is to train models able to generalise to conditions outside the training environment. This is important since for example different patients will have different characteristics. In order to evaluate the capability of our model to generalise we performed two types of tests, one with scars with different shapes when training was performed only with rectangular shapes, and one with multiple onset when training considered only one onset.

As for the generalisation to different scar shapes, we evaluated our model with triangular, circular and complex scars (see Fig. 5). Table 1 shows that the model performs well on the different shapes. The errors are slightly larger than for the rectangular scars used for training, but remain low. They however increase for long term prediction (24 ms here). Figure 5 illustrates the behaviour of the model for typical test sequences.

We also performed tests with multiple onsets. The model shows good results for forecasting of multiple depolarisation waves on one cardiac slab tissue

(Fig. 6), which is essential to capture correctly for ventricular tachycardia simulation. As one can see from the table 2, relative mean-squared error is larger for multiple onsets than for one onset (as it was used for training) but still acceptable. However, this error does not increase proportionately to time of forecast (like in the table 1), because the virtual slab reaches faster the full depolarisation with multiple onsets, an equilibrium state for EP-Net 2.0 model.

Table 1: Relative mean-squared error (MSE) of transmembrane potential forecasting in presence of scars of various forms for different forecasting horizons (cardiac slab conductivity $\sigma = 2$).

	MSE (6 ms)	MSE (12 ms)	MSE (24 ms)
Rectangular shape	1.8×10^{-4}	4.45×10^{-4}	$6, 8 \times 10^{-4}$
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}




Fig. 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 $\sigma = 2$).

Generalisation ability of EP-Net 2.0: various conduction velocities To estimate the ability of EP-Net 2.0 model to learn the conduction velocity of

Table 2: Relative mean-squared error (MSE) of potential forecasting in presence of multiple onsets and scar of rectangular form 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	$4.7 * 10^{-4}$	$5.8 * 10^{-4}$	$6.9 * 10^{-4}$




Fig. 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. 7: Results of EP-Net 2.0 model on scar with circular shape and cardiac slab conductivity $\sigma = 3.8$ (top three rows), and on scar with triangular shape, two onsets and cardiac slab conductivity $\sigma = 1.5$ (bottom three rows).

Table 3: Relative mean-squared error (MSE) of potential forecasting in presence of various conduction velocities of cardiac slab and scar of rectangular form for different forecasting horizons.

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

cardiac tissue we performed tests with with various cardiac slab conductivities (σ). The tests 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 the figure 7, EP-Net 2.0 model keeps the capability to generalise to unseen conditions, such as scars of various shapes and multiple onsets. Quantitative results provided in the table 3 show that model achieves a good precision in forecasting depolarisation waves in cardiac tissue slabs for any conductivity (real number) from the interval $[0.7, 6]$.

3.2 Limitations and Discussion

The section 3.1 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 limitations. For example, as shown in the figure 8, EP-Net 2.0 model does not work properly on thin scars (thickness less than 2 pixels) and produces additional transmembrane potential diffusion through the scar from generated noise. The current model has been trained only to model depolarization of the cardiac slab tissue and cannot predict its repolarisation (see Fig. 4c). This is left for future work.




Fig. 8: Results of trained EP-Net 2.0 model on thin scar with circular shape.

4 Conclusion and Future Work

In this paper, we proposed the EP-Net 2.0 Deep Learning model to learn the cardiac electrophysiology dynamics in presence of complex initial boundary conditions (like scar area, multiple onsets and various conduction velocities). The obtained results show a great generalisation ability of this model to unseen conditions. Despite training on data with scars of only rectangular shape and one onset activation, EP-Net 2.0 model works on data with scars of any possible geometric shape and multiple onsets, even when normalised transmembrane potential stimulations were applied at different moments⁶.

We believe that in future our approach can help upgrade and personalise mathematical model via additional data. However in clinical practice tissue properties are not binary and real data is always noised and sparse, future work includes looking into more complex formulations considering model parameter identification for continuously varying cardiac tissue properties and into possible strategies to complete real data via simulation data.

5 Acknowledgments

This work has been supported by the French government, through the 3IA Côte d’Azur Investments in the Future project managed by the National Research Agency (ANR) with the reference number ANR-19-P3IA-000 and through the ”Research and Teaching chairs in artificial intelligence (AI Chairs)” funding for DL4Clim project. The research leading to these results has also received European funding from the ERC starting grant ECSTATIC (715093). The authors are grateful to the OPAL infrastructure from Université Côte d’Azur for providing resources and support.

References

1. Alvarez, M.A., Luengo, D., Lawrence, N.D.: Linear latent force models using gaussian processes. *IEEE PAMI* **35**(11), 2693–2705 (2013)
2. Ayed, I., de Bézenac, E., Pajot, A., Brajard, J., Gallinari, P.: Learning dynamical systems from partial observations. *arXiv preprint:1902.11136* (2019)
3. Ayed, I., Cedilnik, N., Gallinari, P., 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)
4. Bengio, S., Vinyals, O., Jaitly, N., Shazeer, N.: Scheduled sampling for sequence prediction with recurrent neural networks. *arXiv preprint:1506.03099* (2015)
5. Chen, R.T.Q., Rubanova, Y., Bettencourt, J., Duvenaud, D.: Neural ordinary differential equations. *Adv Neural Inform Process Syst* (2018)
6. Crutchfield, J.P., McNamara, B.: Equations of motion from a data series. *Complex systems* **1**(417-452), 121 (1987)
7. Fresca, S., Manzoni, A., Dedè, L., Quarteroni, A.: Deep learning-based reduced order models in cardiac electrophysiology. *PLOS ONE* **15**(10), e0239416 (2020)

⁶ Visit our github page with trained EP-Net-2.0, for more detailed examples.

8. He, K., Zhang, X., Ren, S., Sun, J.: Deep residual learning for image recognition. In: IEEE conf. CVPR). pp. 770–778 (2016)
9. Kingma, D.P., Ba, J.: Adam: A method for stochastic optimization. arXiv preprint:1412.6980 (2014)
10. Long, Z., Lu, Y., Dong, B.: Pde-net 2.0: Learning PDEs from data with a numeric-symbolic hybrid deep network. *J. Comput. Phys.* **399**, 108925 (2019)
11. Long, Z., Lu, Y., Ma, X., Dong, B.: Pde-net: Learning pdes from data. In: Int. Conf. ICML. pp. 3208–3216. PMLR (2018)
12. Mansi, T., Passerini, T., Comaniciu, D.: *Artificial Intelligence for Computational Modeling of the Heart*. Elsevier (2020)
13. Mitchell, C.C., Schaeffer, D.G.: A two-current model for the dynamics of cardiac membrane. *Bull. Math. Biol.* **65**(5), 767–793 (2003)
14. Nelles, O.: *Nonlinear System Identification*. Springer Berlin Heidelberg (2001)
15. Raissi, M.: Deep hidden physics models: Deep learning of nonlinear partial differential equations. *J Mach Learn Res* **19**(1), 932–955 (2018)
16. Raissi, M., Perdikaris, P., Karniadakis, G.E.: Machine learning of linear differential equations using gaussian processes. *J. Comput. Phys.* **348**, 683–693 (2017)
17. Rapaka, S., Mansi, T., Georgescu, B., Pop, M., Wright, G.A., Kamen, A., Comaniciu, D.: LBM-EP: Lattice-boltzmann method for fast cardiac electrophysiology simulation from 3d images. In: Int. Conf. MICCAI. pp. 33–40. Springer (2012)
18. Relan, J., Chinchapatnam, P., Sermesant, M., Rhode, K., Ginks, M., Delingette, H., Rinaldi, C.A., Razavi, R., Ayache, N.: Coupled personalization of cardiac electrophysiology models for prediction of ischaemic ventricular tachycardia. *Interface focus* **1**(3), 396–407 (2011)
19. Rudy, S.H., Brunton, S.L., Proctor, J.L., Kutz, J.N.: Data-driven discovery of partial differential equations. *Science Advances* **3**(4), e1602614 (2017)
20. Sirignano, J., Spiliopoulos, K.: DGM: A deep learning algorithm for solving partial differential equations. *J. Comput. Phys.* **375**, 1339–1364 (2018)
21. Willard, J.D., Jia, X., Xu, S., Steinbach, M., Kumar, V.: Integrating physics-based modeling with machine learning: A survey. arXiv preprint:2003.04919 (2020)
22. Zhang, S., Lin, G.: Robust data-driven discovery of governing physical laws with error bars. *Proc. Math. Phys. Eng. Sci.* **474**(2217), 20180305 (2018)