Meta-Learning Physics-Informed Neural Networks for Personalized Cardiac Modeling

Maryam Toloubidokhti[1 [0000−0003−2923−441X]], Ryan Missel[1 [0000−0002−9509−6775]], Shichang Lian[1 [0009−0004−6757−6212]], and Linwei Wang[1 [0000−0002−5678−2369]]

Rochester Institute of Technology (RIT), Rochester, NY, 14623, USA
{mt6129,rxm7244,sl5074,lxwast}@rit.edu

Abstract. The advancement of personalized cardiac modeling, particularly through digital cardiac twins, enables tailored treatments based on the physiology of the individual patient. Traditional physics-based methods for optimizing the parameters of these cardiac models face challenges in clinical adoption due to their computational cost. Recent shifts towards data-driven approaches offer improved efficiency, but struggle with generalization and integration of core electrophysiological principles. The emerging use of physics-informed neural networks (PINNs) has the potential to combine the advantages of these two approaches, although still requiring retraining from scratch for each subject. This paper introduces a novel framework for meta-learning PINNs to overcome these challenges, enabling rapid personalization of a PINN to new subjects' data via simple feedforward computation. We instantiate this meta-PINN framework using the Eikonal model as the governing physics, demonstrating its efficacy in significantly reducing computational demands while improving the predictive accuracy of personalized cardiac models. [1]

Keywords: PINNs · meta-learning · cardiac electrophysiology.

1 Introduction

The progress in cardiac modeling, particularly through the lens of digital twins, marks a significant leap in cardiac care. These models, when properly personalized for an individual, allow the planning of treatments [18, 1, 6] and prediction of interventional outcomes [22, 8] for individual subjects. However, rapidly personalizing the functional parameters of these models remains an open challenge.

Earlier progress in estimating cardiac model parameters is deeply rooted in physics-based methodologies, primarily involving iterative optimization of a cardiac model to fit its outputs to available measurements [22, 23, 5]. Due to the need for repeated simulations of complex physics-based equations over a large parameter space, these approaches are associated with significant time and computational cost, resulting in limited applicability for real-time or near-real-time clinical use. Furthermore, because these optimizations are independently

[1] Source code available at https://github.com/temporary-repos/MICCAI2025

conducted for each personalization task, they miss the opportunity to accumulate knowledge about *how to rapidly personalize* to different individuals' data.

To address these limitations, recent advancements have pivoted towards data-driven approaches. Techniques such as reinforcement learning [17, 16] and direct modeling of input-output relationships [10, 9] have been explored. The emergence of deep learning has further enabled neural surrogates capable of approximating complex physiological phenomena. Examples include neural surrogates for cardiovascular hemodynamics that leverage transfer learning for rapid adaptation to new subjects [25], the application of conditional variational autoencoders for predicting ventricular activation using 3D bi-ventricular meshes and electrocardiogram simulations [15], and the meta-learning of spatiotemporal graph convolutional neural networks as a patient-specific surrogate of cardiac electrophysiological (EP) models [14]. While these approaches prioritize computational efficiency, their interpretability is often obscured by neglecting the underlying physical principles of cardiac behavior. Furthermore, their training often requires time-consuming generation of a large quantity of high-fidelity simulation data, while their ability to generalize from simulation to clinical data remains unclear.

Physics-informed neural networks (PINNs) recently emerged as a potential solution to bridge the gap between physics-based and data-driven approaches [19]. PINNs integrate known physical laws – in the form of a partial differential equation (PDE) – through a PDE residual loss to encourage the neural network to approximate PDE solutions without direct supervision from these solutions [7, 19]. Recent studies have successfully utilized PINNs in personalized cardiac EP modeling [24, 13, 20]. In EP-PINN [13], a PINN is trained to generate action potentials in a 2D grid while estimating the PDE's key EP parameters. In [24], a spatial-temporal adaptive strategy is introduced to enhance the ability of PINNs to handle complex geometry, long time domain, and sharp gradients in PDE solutions. In FiberNet [20], a PINN is learned to solve for the Eikonal equation when inferring its spatially-varying fiber orientations from electroanatomical maps.

Despite this progress, a significant challenge remains for the widespread application of PINNs in personalized cardiac modeling: because the optimization of PINN is governed (in part) by a known PDE with its associated physics-based parameters, to personalize it requires the PINN to be re-trained from scratch along with estimating patient-specific parameters for the governing PDE. This results in a bottleneck similar to traditional physics-based approaches: the computational burden for clinical adoption, and a missed opportunity to *learn how to personalize* a PINN to shift this burden from deployment to offline development.

To address this critical gap, we introduce a novel framework for meta-learning the personalization of PINNs, enabling the rapid adaptation of a PINN to new patient-specific data without the need for retraining. It includes three key components: 1) an *adaptive* PINN whose weight parameters are generated from a hypernetwork from inferred patient-specific embeddings, 2) a conduction velocity network that personalizes the parameters of the governing PDE in a similar fashion, and 3) a meta-learning approach that learns to extract patient-specific embeddings from few-shot context observations in a rapid feedforward fashion.

As a proof of concept, we developed the proposed meta-PINN framework for the isotropic Eikonal PDE [21] with spatially varying tissue properties. In *in-silico* and *in-vivo* experiments, we evaluated meta-PINN against physics-based and neural approaches to personalized EP modeling, as well as state-of-the-art PINN approaches where each PINN is individually optimized for each individual heart. We show that, at deployment time given sparse measurements from a patient, meta-PINN is able to deliver a personalized and physics-informed neural surrogate for cardiac EP via real-time feedforward computation, making an important stride towards fast adaptation of cardiac digital twins.

2 Background

Eikonal PDE: The Eikonal PDE is widely used to model the rapid propagation of electrical currents across the myocardium. We consider the isotropic Eikonal:

$$|\nabla T(\boldsymbol{x})|F(\boldsymbol{x}) = 1 \tag{1}$$

where $T(\boldsymbol{x})$ is a scalar field representing the arrival time of the electrical wavefront at spatial position $\boldsymbol{x} \in \mathbb{R}^3$, and $F(\boldsymbol{x})$ is a scalar field with units of speed (mm/ms) representing the local conduction velocity. While $F(\boldsymbol{x})$ denotes speed in an isotropic medium, we use the term *conduction velocity* following standard cardiac electrophysiology terminology. This equation effectively captures the isotropic propagation of electrical signals through the heterogeneous medium of the heart, accounting for varying conductivity in space.

PINNs: PINNs embed known physical laws into the learning algorithm of neural networks. In the context of the Eikonal PDE, a PINN $T_{NN_\phi}(\boldsymbol{x})$ parameterized by ϕ can be designed to output the arrival time of the electrical wavefront at any input spatial position $\boldsymbol{x}$. Instead of conventional supervised training, the key to PINN training is a PDE residual loss L_{PDE} that encapsulates the adherence of $T_{NN_\phi}(\boldsymbol{x})$ to the governing Eikonal physics as:

$$L_{PDE}(T_{NN_\phi}; F(\boldsymbol{x})) = \frac{1}{N_{PDE}} \sum\nolimits_{j=1}^{N_{PDE}} \left(|\nabla T_{NN_\phi}(\boldsymbol{x}_j)|F(\boldsymbol{x}_j) - 1\right)^2 \tag{2}$$

where N_{PDE} is the number of *residual points* on which the PDE constraint is enforced. When observations of the arrival time $Y = T_{\text{obs}}(\boldsymbol{x})$ are available at N_{data} number of locations $\boldsymbol{x}_i$, a data-fidelity loss L_{data} can also be enforced:

$$L_{data}(T_{NN_\phi}, Y) = \frac{1}{N_{data}} \sum\nolimits_{i=1}^{N_{data}} \left(T_{NN_\phi}(\boldsymbol{x}_i) - Y(\boldsymbol{x}_i)\right)^2 \tag{3}$$

If the conduction velocity $F(\boldsymbol{x})$ is unknown, it must be simultaneously estimated when $\boldsymbol{\phi}$ is optimized, resulting in the following optimization objective:

$$\{\hat{\boldsymbol{\phi}}, \hat{F}(\boldsymbol{x})\} = \arg\min_{\boldsymbol{\phi}, F(\boldsymbol{x})} \{L_{data}(T_{NN_{\boldsymbol{\phi}}}, Y) + \alpha L_{PDE}(T_{NN_{\boldsymbol{\phi}}}; F(\boldsymbol{x}))\} \tag{4}$$

where the hyperparameter α balances the two losses. In state-of-the-art PINNs, $T_{NN_{\hat{\boldsymbol{\phi}}}}(\boldsymbol{x})$ and $\hat{F}(\boldsymbol{x})$ must be optimized from scratch for each set of measurements.

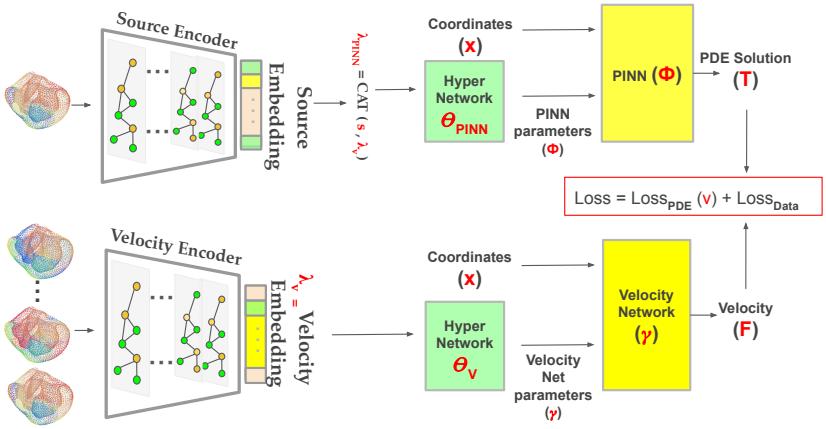

Fig. 1. Overview of meta-PINN. At inference time, sparse context AT samples on the ventricular geometry are used to extract a patient-specific embedding that adapts the PINN and velocity network to estimate personalized conduction velocity and the full volumetric AT map of a query sample, given only the first 20% of its activation.

3 Methodology

Fig. 1 outlines the proposed meta-PINN framework for personalized cardiac EP modeling. Instead of optimizing T_{NN_ϕ} and $F(\boldsymbol{x})$ from scratch for each set of observations Y using Equation (4), we aim to *learn to personalize* T_{NN_ϕ} and $F(\boldsymbol{x})$ via a meta-learning approach to automatically adapt them to patient-specific Y in real time. This is achieved with two key components:

- Learn-to-personalize meta-inference: a feedforward meta-learner that extracts patient-specific embeddings from context observations of a subject.
- Adaptive generative models: a PINN $T_{NN_\phi}(\boldsymbol{x})$ and velocity network $F_{NN_\gamma}(\boldsymbol{x})$, with $\boldsymbol{\phi}$ and $\boldsymbol{\gamma}$ adaptable to patient-specific embeddings through hypernetworks, to predict for any other samples for the same patient.

Adaptable PINN and Velocity-Net: To model the spatially-varying conduction velocity to capture local abnormality, we model $F(\boldsymbol{x})$ with a neural network $F_{NN_\gamma}(\boldsymbol{x})$ and term it the *velocity-net*, similar to the framework presented in [20, 21]. To avoid repeating the optimization of $\boldsymbol{\phi}$ and $\boldsymbol{\gamma}$ for T_{NN_ϕ} and F_{NN_γ} for every observation Y, we model $\boldsymbol{\phi}$ and $\boldsymbol{\gamma}$ with hypernetworks [12, 2] as:

$$\boldsymbol{\phi} = H_{\boldsymbol{\theta}_{\text{PINN}}}(\boldsymbol{\lambda}_{\text{PINN}}), \quad \boldsymbol{\gamma} = H_{\boldsymbol{\theta}_{\text{v}}}(\boldsymbol{\lambda}_{\text{v}}) \tag{5}$$

where $\boldsymbol{\lambda}_{\text{PINN}}$ or $\boldsymbol{\lambda}_{\text{v}}$ are learnable patient-specific embeddings that will be described in the next section, while $\boldsymbol{\theta}_{\text{PINN}}$ and $\boldsymbol{\theta}_{\text{v}}$ represent the learnable weight parameters of the hypernetworks. These adaptive generative models mitigate the need for inefficient optimization-from-scratch for each new set of observations.

Learning to Identify the PINN and Velocity-Net: To identify the patient-specific embedding $\boldsymbol{\lambda}_{\text{PINN}}$ and $\boldsymbol{\lambda}_{\text{v}}$ that can adapt $T_{NN_{\boldsymbol{\phi}}}$ and $F_{NN_{\boldsymbol{\gamma}}}$, we consider their respective generative factors. For an observed activation sequence to be modeled by the PINN $T_{NN_{\boldsymbol{\phi}}}$, there are two primary generative factors: the spatially-varying tissue property $\boldsymbol{c}$ and the sources of activation $\boldsymbol{s}$. They each correspond to the parameter and initial condition for the governing Eikonal PDE: in other words, $\boldsymbol{c}$ is also the generative factor for the conduction velocity $F_{NN_{\boldsymbol{\gamma}}}$. Therefore, we can define $\boldsymbol{\lambda}_{\text{PINN}} = (\boldsymbol{c}, \boldsymbol{s})$ and $\boldsymbol{\lambda}_{\text{v}} = \boldsymbol{c}$.

While it is possible to simply infer $\boldsymbol{c}$ and $\boldsymbol{s}$ from an observed activation map Y, a challenge of *identifiability* arises: what additional learning signal is needed in order to encourage the model to learn to separate these two embeddings from the same Y? We consider an important distinction between them: the tissue property corresponding to $\boldsymbol{c}$ is specific to a subject i but can be shared by multiple activation maps $\mathcal{Y}_i = \{Y_i^{(j)}\}_{j=1}^k$ from the same subject, while the initial condition corresponding to $\boldsymbol{s}$ is specific to each activation map $Y_i^{(j)}$.

Source Encoder: The source encoder $\mathcal{E}_{\boldsymbol{\psi}_{\text{s}}}$, parameterized by $\boldsymbol{\psi}_{\text{s}}$, aims to identify the representation of the source of activation unique to each $Y_i^{(j)}$ as:

$$\boldsymbol{s}_i^{(j)} = \mathcal{E}_{\boldsymbol{\psi}_{\text{s}}}(Y_i^{(j)}) \tag{6}$$

Velocity Encoder: The velocity encoder $\mathcal{E}_{\boldsymbol{\psi}_{\text{c}}}$, parameterized by $\boldsymbol{\psi}_{\text{c}}$, aims to identify the representation for the heterogeneous tissue property by extracting the shared embedding from k activation maps $\{Y_i^{(j)}\}_{j=1}^k$ from the same subject i. To provide an explicit learning signal for a shared $\boldsymbol{c}_i$ for subject i, we obtain an embedding from each $Y_i^{(j)}$ and then aggregate them via a simple averaging:

$$c_i = \frac{1}{K}\sum_{k=1}^{K} \mathcal{E}_{\psi_{\text{c}}}(Y_i^{(k)}) \tag{7}$$

The Learn-to-Identify Meta-Objectives: Given a dataset of observed activation maps for N subjects $\mathcal{Y} = \{\mathcal{Y}_i\}_{i=1}^N$, where M activation time maps with varying sources of activation are available for each subject $\mathcal{Y}_i = \{Y_i^{(j)}\}_{j=1}^M$, we now formulate a meta-learning objective to learn to personalize $T_{NN_{\boldsymbol{\phi}}}$ and $F_{NN_{\boldsymbol{\gamma}}}$. More specifically, for each subject i, we consider k activation maps as context samples $\mathcal{Y}_i^c = \{Y_i^{c,(j)}\}_{j=1}^k$ while the remaining as query samples $\mathcal{Y}_i^q = \{Y_i^{q,(j)}\}_{j=1}^{M-k}$.

For each query sample Y_i^q, the model extracts a subject-specific embedding $\boldsymbol{c}$ from the context set $\mathcal{Y}_i^c$ and and activation-specific embedding $\boldsymbol{s}$ from only the first 20% activation of Y_i^q. These embeddings are then used to adapt the PINN and velocity-net for predicting the full Y_i^q while minimizing the PDE residual loss as:

$$\hat{\boldsymbol{\Theta}} = \arg\min_{\boldsymbol{\Theta}} \sum\nolimits_{i=1}^{N} \sum\nolimits_{Y_i^q \in \mathcal{Y}_i^q} \left\{ L_{data}(T_{NN_{\boldsymbol{\phi}}}, Y_i^q) + \alpha L_{PDE}(T_{NN_{\boldsymbol{\phi}}}; F_{NN_{\boldsymbol{\gamma}}}) \right\} \tag{8}$$

where $\boldsymbol{\Theta} = \{\boldsymbol{\psi}_s, \boldsymbol{\psi}_c, \boldsymbol{\theta}_{\text{PINN}}, \boldsymbol{\theta}_v\}$. This optimization is carried out in an episodic training scheme, where the division of context and query samples is changed at each epoch for each subject. The outcome is a meta-model that, given k (which can be variable) observed activation maps for a subject, can obtain a subject-specific PINN along with an estimate of the spatially-varying conduction velocity for that subject — all via simple and rapid feedforward computation.

4 Experiments and Results

In all experiments, meta-PINN employed encoders with six SplineConv layers (kernel size = 5, dim = 3) followed by adaptive max pooling, effectively capturing spatial relations in the ventricular geometry. The PINN hypernetwork consisted of four fully connected layers, generating parameters for the six-layer PINN with 20 neurons per layer and Tanh activation. Similarly, the velocity hypernetwork had five hidden layers, producing parameters for the eight-layer velocity network with ReLU activation. Optimization was carried out using the AdamW optimizer, starting with a learning rate of 2×10^{-5} that we adjusted using a cosine annealing scheduler. We gradually increased the PDE loss hyperparameter, α, from 0 to 1 throughout the training epochs. Each training episode, executed on an NVIDIA RTX 3090 with 24GB of memory, took ~ 3 seconds.

4.1 Synthetic Data Experiments

Data: We created activation-time (AT) maps using the Eikonal model in Equation (1), utilizing the fast marching method from the *fim-python* library [11]. We obtained four unique human ventricular geometries from MRI images, aligned via non-rigid registration. Two tissue classes were modeled: scar (0.1 m/s) and healthy (0.6 m/s), with scars placed at five distinct locations per geometry. For each scar-geometry combination, we generated 20 AT maps with random placements of the site of origin, resulting in 400 AT maps, 200 for training and 200 for test. Train and test samples were split by activation origin per task; all tasks appear in both splits. We set $k = 5$ for the context set. To reflect practice, we use endocardial AT measurements (200-400 nodes), to predict AT and conduction velocity throughout the ventricular myocardium (1600-2400 nodes).

Baselines & Metrics: We evaluated meta-PINN against three baselines: 1) direct optimization of $F(\boldsymbol{x})$ for the Eikonal PDE using the gradient-free BOBYQA algorithm [4] (*physics-based*); 2) data-driven meta-learning of T_{NN_ϕ} but without PDE-residual loss, similar to [14] – this can also be seen as an ablation of the presented meta-PINN (*meta-neural*); and 3) standalone PINN models where T_{NN_ϕ} is optimized for each observed activation time and F_{NN_γ} is shared across PINNs of the same subject, similar to that described in [20] (*multi-PINN*). Because *meta-neural* and *meta-PINN* have predictive abilities once personalized from context samples, they were tested for that purpose. *Physics-based* optimization and *multi-PINN* were optimized to each test sample instead.

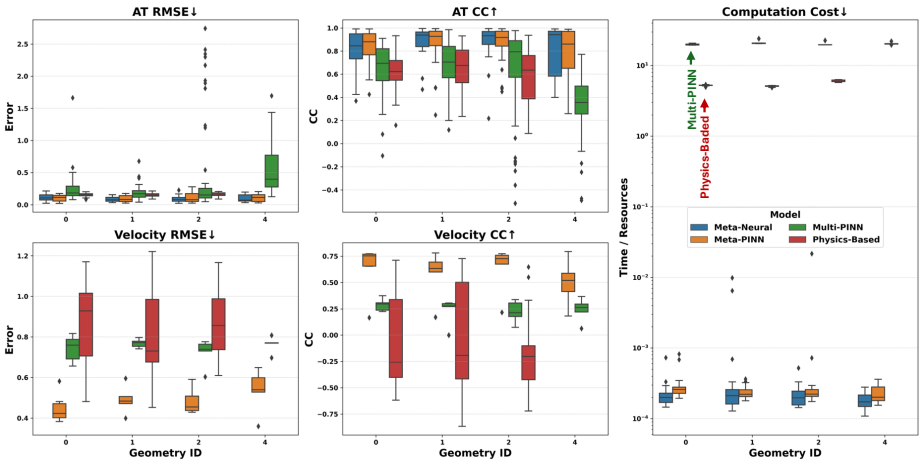

Fig. 2. Quantitative results on synthetic data. Meta-PINN gained significant improvements in AT and velocity accuracy compared to physics-based optimization and multi-PINN at a fraction of computation. Its performance and computation were comparable to meta-neural, with added benefit of interpretable personalized velocity estimation.

For quantitative metrics, we considered 1) relative mean squared error (RMSE) and 2) spatial correlation coefficient (SCC) between predicted and true AT maps and velocity fields, as well as 3) computational cost to achieve personalization.

Results: Fig. 2 summarizes the quantitative results while visual examples are shown in Fig. 3. Compared to *physics*-based optimization (red) and *multi-PINN* (green), meta-PINN (orange) obtained significantly improved accuracy in both AT map prediction and velocity estimation, *e.g.*, 22.3±4.0% and 68.9±10.0% improvements in SCC respectively for AT and velocity maps, compared to physics; and $19.7 \pm 4.7\%$ and $36.2 \pm 4.1\%$ improvements compared to multi-PINN. These improvements were obtained with a fraction of computation time to personalize ($0.009 \pm 0.007\%$ of physics-based, and $0.003 \pm 0.002\%$ of multi-PINN). Compared to meta-neural (blue), meta-PINN achieved comparable accuracy in predicting unseen AT maps with slightly higher computation time. It however had the benefit of delivering subject-specific conduction velocity maps as shown in Fig. 3, leading to interpretable clinical utility.

Ablation and Analysis: We conducted three ablation-style analyses. First, we evaluated the model's sensitivity to the number of context observations k. Meta-PINN showed stable performance with fewer context maps: for k=1, AT/velocity SCC $= 0.86 \pm 0.17/0.56 \pm 0.26$; for k=3, $0.86 \pm 0.16/0.59 \pm 0.20$; and for k=5, $0.87 \pm 0.15/0.60 \pm 0.19$ suggesting strong robustness with sparse supervision. Second, to assess whether accuracy differs between observed and unobserved regions, we conducted a preliminary stratified analysis: meta-PINN achieved AT/velocity

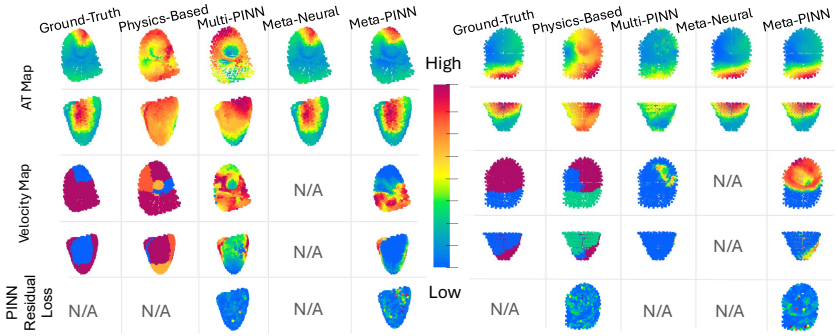

Fig. 3. Visual examples on synthetic data. Meta-PINN showed significantly improved AT and velocity compared to multi-PINN and physics-based optimization. Meta-neural showed similar AT maps without an ability to deliver personalized conduction velocity.

SCC of 0.87/0.57 on the endocardium (where measurements are) and 0.87/0.60 on the remaining myocardium; given the minimal difference, full regional analysis is left to future work. Finally, we evaluated meta-PINN's adherence to the Eikonal PDE. As shown in the last row of Fig. 3, residual heatmaps indicate smoothly minimized PDE loss. Without any test-time retraining, meta-PINN achieved a per-node residual of 0.31 ± 0.18, comparable to individually trained PINNs (0.29 ± 0.04), confirming its ability to satisfy PDE constraints.

4.2 Real Data Experiments

Data: We considered *in-vivo* data from an animal model experiment [3], which included cardiac activation sequences driven by bipolar stimulation via intramural plunge needles at six different left-ventricular locations. Epicardial potentials were recorded using a 247-electrode sock, and geometric surfaces were then generated from the electrode positions acquired during the experiment.

Baselines & Metrics: We directly tested meta-PINN and meta-neural trained from synthetic data, each time considering five epicardial maps as the context to predict the left-out activation map. Note the large generalization gap due to both the sim-to-real gap and the change of endocardial to epicardial measurements. Physics-based and multi-PINN were optimized from scratch.

Results: Fig. 4 provides several visuals of predicted AT maps. Note the relatively stronger performance of physics-based optimization and multi-PINN as they were optimized from scratch to all six available maps. Meta-neural, in comparison, completely failed in this generalization task by producing similar outputs across all test cases. Meta-PINN, in comparison, was able to capture the main activation patterns with a performance similar to multi-PINN, demonstrating the benefit of physics inductive bias compared to its fully neural counterpart.

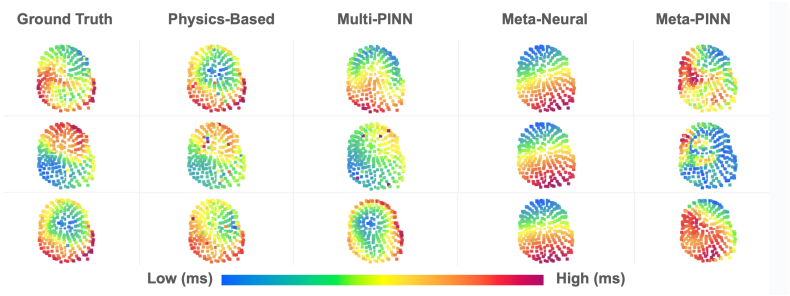

Fig. 4. Visual examples on real-data AT maps.

4.3 Conclusion & Discussion

We presented an innovative meta-PINN for learning to rapidly adapt a personalized and physics-informed surrogate to sparse surface AT measurements from a subject. Results showed improved accuracy and computational efficiency over existing approaches. As a proof-of-concept, this study can be strengthened by a systematic generalization study considering unseen geometry or scar settings in synthetic experiments, as well as evaluation on a larger number of real-data experiments. Future works will extend meta-PINN to anisotropic Eikonal PDE, higher-fidelity cardiac EP models, as well as beyond the domain of cardiac EP.

Acknowledgments. This study was supported by NIH National Heart, Lung, And Blood Institute (NHLBI) grant R01HL145590, NIH National Institute of Nursing Research (NINR) grant R01NR0183 and NSF OAC-2212548.

Disclosure of Interests. The authors have no competing interests to declare that are relevant to the content of this article.

References

1. Arevalo, H.J., Vadakkumpadan, F., Guallar, E., Jebb, A., Malamas, P., Wu, K.C., Trayanova, N.A.: Arrhythmia risk stratification of patients after myocardial infarction using personalized heart models. Nature communications **7**(1), 11437 (2016)
2. de Avila Belbute-Peres, F., Chen, Y.f., Sha, F.: Hyperpinn: Learning parameterized differential equations with physics-informed hypernetworks. The symbiosis of deep learning and differential equations **690** (2021)
3. Bergquist, J.A., Good, W.W., Zenger, B., Tate, J.D., Rupp, L.C., MacLeod, R.S.: The electrocardiographic forward problem: A benchmark study. Computers in biology and medicine **134**, 104476 (2021)
4. Cartis, C., Fiala, J., Marteau, B., Roberts, L.: Improving the flexibility and robustness of model-based derivative-free optimization solvers. ACM Transactions on Mathematical Software (TOMS) **45**(3), 1–41 (2019)
5. Chinchapatnam, P., Rhode, K.S., Ginks, M., Mansi, T., Peyrat, J.M., Lambiase, P., Rinaldi, C.A., Razavi, R., Arridge, S., Sermesant, M.: Estimation of volumetric

myocardial apparent conductivity from endocardial electro-anatomical mapping. In: 2009 Annual International Conference of the IEEE Engineering in Medicine and Biology Society. pp. 2907–2910. IEEE (2009)
6. Corral-Acero, J., Margara, F., Marciniak, M., Rodero, C., Loncaric, F., Feng, Y., Gilbert, A., Fernandes, J.F., Bukhari, H.A., Wajdan, A., et al.: The 'digital twin'to enable the vision of precision cardiology. European heart journal **41**(48), 4556–4564 (2020)
7. Cuomo, S., Di Cola, V.S., Giampaolo, F., Rozza, G., Raissi, M., Piccialli, F.: Scientific machine learning through physics–informed neural networks: Where we are and what's next. Journal of Scientific Computing **92**(3), 88 (2022)
8. Gerach, T., Schuler, S., Fröhlich, J., Lindner, L., Kovacheva, E., Moss, R., Wülfers, E.M., Seemann, G., Wieners, C., Loewe, A.: Electro-mechanical whole-heart digital twins: a fully coupled multi-physics approach. Mathematics **9**(11), 1247 (2021)
9. Giffard-Roisin, S., Delingette, H., Jackson, T., Fovargue, L., Lee, J., Rinaldi, A., Ayache, N., Razavi, R., Sermesant, M.: Sparse bayesian non-linear regression for multiple onsets estimation in non-invasive cardiac electrophysiology. In: Functional Imaging and Modelling of the Heart: 9th International Conference, FIMH 2017, Toronto, ON, Canada, June 11-13, 2017, Proceedings 9. pp. 230–238. Springer (2017)
10. Giffard-Roisin, S., Jackson, T., Fovargue, L., Lee, J., Delingette, H., Razavi, R., Ayache, N., Sermesant, M.: Noninvasive personalization of a cardiac electrophysiology model from body surface potential mapping. IEEE Transactions on Biomedical Engineering **64**(9), 2206–2218 (2016)
11. Grandits, T.: A fast iterative method python package. Journal of Open Source Software **6**(66), 3641 (2021). https://doi.org/10.21105/joss.03641, https://doi.org/10.21105/joss.03641
12. Ha, D., Dai, A., Le, Q.V.: Hypernetworks. arXiv preprint arXiv:1609.09106 (2016)
13. Herrero Martin, C., Oved, A., Chowdhury, R.A., Ullmann, E., Peters, N.S., Bharath, A.A., Varela, M.: Ep-pinns: Cardiac electrophysiology characterisation using physics-informed neural networks. Frontiers in Cardiovascular Medicine **8**, 768419 (2022)
14. Jiang, X., Li, Z., Missel, R., Zaman, M.S., Zenger, B., Good, W.W., MacLeod, R.S., Sapp, J.L., Wang, L.: Few-shot generation of personalized neural surrogates for cardiac simulation via bayesian meta-learning. In: International Conference on Medical Image Computing and Computer-Assisted Intervention. pp. 46–56. Springer (2022)
15. Li, L., Camps, J., Banerjee, A., Beetz, M., Rodriguez, B., Grau, V.: Deep computational model for the inference of ventricular activation properties. In: International Workshop on Statistical Atlases and Computational Models of the Heart. pp. 369–380. Springer (2022)
16. Neumann, D., Mansi, T.: Machine learning methods for robust parameter estimation. In: Artificial Intelligence for Computational Modeling of the Heart, pp. 161–181. Elsevier (2020)
17. Neumann, D., Mansi, T., Itu, L., Georgescu, B., Kayvanpour, E., Sedaghat-Hamedani, F., Amr, A., Haas, J., Katus, H., Meder, B., et al.: A self-taught artificial agent for multi-physics computational model personalization. Medical image analysis **34**, 52–64 (2016)
18. Prakosa, A., Arevalo, H.J., Deng, D., Boyle, P.M., Nikolov, P.P., Ashikaga, H., Blauer, J.J., Ghafoori, E., Park, C.J., Blake III, R.C., et al.: Personalized virtual-heart technology for guiding the ablation of infarct-related ventricular tachycardia. Nature biomedical engineering **2**(10), 732–740 (2018)

19. Raissi, M., Perdikaris, P., Karniadakis, G.E.: 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 (2 2019). https://doi.org/10.1016/j.jcp.2018.10.045
20. Ruiz Herrera, C., Grandits, T., Plank, G., Perdikaris, P., Sahli Costabal, F., Pezzuto, S.: Physics-informed neural networks to learn cardiac fiber orientation from multiple electroanatomical maps. Engineering with Computers **38**(5), 3957–3973 (2022)
21. Sahli Costabal, F., Yang, Y., Perdikaris, P., Hurtado, D.E., Kuhl, E.: Physics-informed neural networks for cardiac activation mapping. Frontiers in Physics **8**, 42 (2020)
22. Sermesant, M., Chabiniok, R., Chinchapatnam, P., Mansi, T., Billet, F., Moireau, P., Peyrat, J.M., Wong, K., Relan, J., Rhode, K., et al.: Patient-specific electromechanical models of the heart for the prediction of pacing acute effects in crt: a preliminary clinical validation. Medical image analysis **16**(1), 201–215 (2012)
23. Wong, K.C., Sermesant, M., Rhode, K., Ginks, M., Rinaldi, C.A., Razavi, R., Delingette, H., Ayache, N.: Velocity-based cardiac contractility personalization from images using derivative-free optimization. Journal of the mechanical behavior of biomedical materials **43**, 35–52 (2015)
24. Ye, Y., Liu, H., Jiang, X., Toloubidokhti, M., Wang, L.: A spatial-temporally adaptive pinn framework for 3d bi-ventricular electrophysiological simulations and parameter inference. In: International Conference on Medical Image Computing and Computer-Assisted Intervention. pp. 163–172. Springer (2023)
25. Zhou, Y., He, Y., Wu, J., Cui, C., Chen, M., Sun, B.: A method of parameter estimation for cardiovascular hemodynamics based on deep learning and its application to personalize a reduced-order model. International Journal for Numerical Methods in Biomedical Engineering **38**(1), e3533 (2022)