

Physics-Informed Surrogate Model using Graph Neural Network for Cardiac Electrophysiology

Vu Anh Tran¹, Mengxiao Wang^{1,2}, Yilin Lyu¹, Jagmohan Chauhan^{3,4}, Vivek Muthurangu⁵, Ching-Hui Sia^{6,7}, and Lei Li^{1✉}

¹ Department of Biomedical Engineering, National University of Singapore, Singapore

lei.li@nus.edu.sg

² School of Mechanical Engineering, Shanghai Jiao Tong University, Shanghai, China

³ Department of Computer Science, University College London, London, United Kingdom

⁴ UCL AI Centre, University College London, London, United Kingdom

⁵ Institute of Cardiovascular Science, University College London, London, United Kingdom

⁶ Department of Medicine, National University of Singapore, Singapore

⁷ Department of Cardiology, National University Heart Centre Singapore, Singapore

Abstract. High-fidelity cardiac electrophysiology (EP) simulations are critical for understanding and treating heart disease, but conventional solvers are often too slow for real-time applications. Recent deep learning approaches have explored physics-informed neural networks as surrogate models to accelerate EP prediction while preserving biophysical fidelity. However, these methods typically struggle to incorporate key anatomical features, such as cardiac geometry and fiber orientation, and lack explicit integration of EP-specific physical constraints, limiting their reliability and interpretability. In this study, we present a physics-informed graph attention network that serves as an efficient surrogate model for cardiac EP. By embedding the monodomain equation and Aliev–Panfilov ionic model into the loss function and leveraging graph attention mechanisms to capture anatomical structure from cardiac meshes, our approach substantially outperforms methods that disregard fiber architecture. Experiments on unseen 3D biventricular geometries demonstrate strong generalization and accurate wavefront dynamics without fine-tuning, establishing our method as a promising surrogate for anatomically faithful, real-time EP modeling.

Keywords: Cardiac Electrophysiology · Surrogate Model · Graph Neural Network · Physics-Informed Model.

1 Introduction

Accurate characterization of cardiac electrophysiology (EP) dynamics is fundamental for investigating heart disease, but obtaining such information in practice still relies on time-consuming mathematical modeling [25]. This is largely

because comprehensive, high-resolution EP measurements are limited by several constraints in clinical data acquisition, including the need for invasive procedures [14, 7, 8]. By solving coupled systems of partial differential equations (PDEs) and ordinary differential equations (ODEs) on finite-element cardiac meshes, these simulations can accurately represent both the cellular excitation dynamics (reaction) and the intercellular spread of electrical activity (diffusion). Despite their accuracy, these methods are still far from enabling real-time cardiac digital twins due to either model oversimplification or prohibitive computational cost [12]. Specifically, fast phenomenological approaches such as the Eikonal model [11] are often too simplified because they do not explicitly model transmembrane potential dynamics. In contrast, more detailed ionic models (e.g., ten Tusscher and Panfilov, 2006 [24]) are often too computationally expensive due to the large number of state variables and parameters they must evolve. Moreover, these modeling methods are also highly sensitive to parameter choices [15], further complicating calibration and making them unable to rapidly generate sufficiently detailed personalized simulations for broad clinical use.

To tackle that problem, many deep learning solutions have been developed as surrogate models approximating the aforementioned simulation process [17]. Due to the strongly physics-governed nature of cardiac EP, many studies adopt Physics-Informed Neural Networks (PINNs) [21], which added the residual of the governing equations as a soft constraint alongside the data-fitting loss, into their solutions. Specifically, while some early attempts [19, 22] regularized basic FCNN/RNN surrogates using simplified EP formulations such as the eikonal or FitzHugh–Nagumo models, others [2, 9] leveraged CNN-style designs to learn EP dynamics directly from Cartesian-grid slab simulations. Although these studies demonstrated the benefit of enforcing physical consistency in learning cardiac EP, their formulations remain limited by grid-bound, fixed-resolution representations and overly idealized geometries.

Several more complicated deep learning techniques have been explored recently to improve the generalization further. Examples include applying Physics-Informed Neural Operators (PINO) [13] into EP prediction [16], or using a Graph Neural Network to learn electrical wave propagation from mesh-based data [18]. However, these approaches often either lack problem-tailored physical constraints [18] or ignore realistic anisotropic anatomical information (e.g., fiber orientation and true 3D geometry), largely because they are validated on low-dimensional or idealized datasets [16, 10]. As a result, their generalization capacity and interpretability remain constrained when confronted with the topological variability. These limitations highlight the need for an interpretable physics-informed deep learning framework that can leverage anatomical information to predict personalized cardiac EP behaviors, and remain robust across varying 3D mesh resolutions as well as realistic geometries.

To address these challenges, we propose a novel physics-informed graph neural network (GNN) that leverages the biventricular mesh-based anatomical structure to learn and predict cardiac electrophysiology dynamics under explicit physical constraints. We first convert the biventricular meshes and corresponding

simulation results into a graph representation, where anatomical and temporal EP information are encoded as node and edge features. These spatiotemporal properties include anisotropic conduction varying with fibre orientation across locations, as well as the EP state of each cell, supporting the model to capture wavefront evolution. Second, we introduce a GNN-based architecture supported by a Long-Short Term Memory block to auto-regressively predict future electrical wave propagation in the biventricle from a few initial simulation frames. To enhance interpretability, we constrain this architecture with physics-guided loss components that enforce the reaction-diffusion activities during training.

2 Methodology

2.1 Anatomy-Aware Graph Data Representation

The graph representation $G^t = (\mathcal{V}^t, \mathcal{E}^t)$ of each action potential map at time t is constructed to preserve the original mesh connectivity. For each node $i \in \mathcal{V}^t$, since the local conductivity matrix $\mathbf{D}_i$, which is determined by the fibre orientation at each location, strongly influences electrical wave propagation [26], this anatomical information must be explicitly encoded in the graph representation. Accordingly, we define a node feature vector $\mathbf{x}_i$ composed of the action potential value V_m^t , the recovery variable z_i^t at time t , and the fibre, sheet, and normal vectors $(\mathbf{f}_i, \mathbf{s}_i, \mathbf{n}_i)$. For each edge $(i, j) \in \mathcal{E}^t$, we define edge feature vectors $\mathbf{e}_{ij}$ to encode both geometric and anisotropic conduction information. The relative displacement vector $\mathbf{d}_{ij} = \mathbf{x}_j - \mathbf{x}_i$, and its length $\|\mathbf{d}_{ij}\|$ describe the local spatial relationship between two connected nodes. To capture how the propagation direction aligns with the local tissue structure, we further include the inner products between the unit edge direction $\hat{\mathbf{d}}_{ij}$ and the averaged fibre, sheet, and normal vectors along the edge, denoted by $\mathbf{f}_{ij}$, $\mathbf{s}_{ij}$, and $\mathbf{n}_{ij}$, respectively, since there are significant differences [5] in the conduction velocity between directions within cardiomyocytes. In addition, to be more straightforward, we define an effective conductivity feature σ_{eff} as:

$$\sigma_{\text{eff}} = \sigma_f \left(\hat{\mathbf{d}}_{ij} \cdot \mathbf{f}_{ij} \right)^2 + \sigma_s \left(\hat{\mathbf{d}}_{ij} \cdot \mathbf{s}_{ij} \right)^2 + \sigma_n \left(\hat{\mathbf{d}}_{ij} \cdot \mathbf{n}_{ij} \right)^2, \quad (1)$$

where σ_f , σ_s , and σ_n are predefined conductivities along the fibre, sheet, and normal directions, respectively. By this design, each graph $\mathcal{G}^t$ is represented by a node feature tensor $\mathbf{X}^t \in \mathbb{R}^{|\mathcal{V}| \times 11}$ and an edge feature tensor $\mathbf{E}^t \in \mathbb{R}^{|\mathcal{E}| \times 8}$, where $|\mathcal{V}|$ and $|\mathcal{E}|$ denote the numbers of nodes and edges, respectively, capturing both anatomical and electrophysiological information at time t .

2.2 Spatiotemporal Physics-Informed GNN

Figure. 1 illustrates the physics-constrained GNN framework for learning the temporal evolution of the V_m distribution, which builds upon the Encoder-Processor-Decoder architecture introduced in [20] and is further enhanced with an additional LSTM block to better capture temporal dependencies.




Fig. 1. Overview of the proposed physics-constrained spatiotemporal graph neural network-based framework for cardiac electrical propagation prediction. ODE/ PDE: ordinary/ partial differential equation. GNN: graph neural network.

Spatiotemporal GNN-based architecture: In our model, a window of T consecutive graph frames ended at $\mathcal{G}^t$, which is also referred to as the current state, is treated as one training sample. While the node feature tensors from all frames are encoded at the first stage, the edge features are encoded only once from the last frame, since they remain unchanged over time within the same trajectory. The stacked encoded node representations are then passed through an LSTM to extract the hidden temporal state $\mathbf{h}^t$, which summarizes the temporal evolution over the entire window. This hidden state is subsequently combined with the encoded representation of the current graph $\mathcal{G}^t$ to produce a temporally informed latent representation for the prediction step. Similar to [20], we use K stacked GraphNet [23] blocks as the processor to perform message passing, followed by a decoder to transform latent representations into desired values. We aim to the predicted two state increments over $\mathcal{G}^t$, namely ΔV_m^t , and Δz^t , and the next state is updated as

$$V_m^{t+1} = V_m^t + \Delta V_m^t, \quad z^{t+1} = z^t + \Delta z^t. \quad (2)$$

Predicting state changes instead of absolute values reduces the learning burden by focusing the model on local temporal transitions, while jointly estimating both V_m^{t+1} and the recovery variable z^{t+1} encourages the network to follow the coupled cardiomyocyte dynamics described by the Aliev-Panfilov model [1].

Explicit Reaction-Diffusion Physics Constraint: To explicitly regularize the prediction with the governing electrophysiological dynamics, we impose physics-informed constraints derived from the monodomain equation coupled with the Aliev-Panfilov model [1]. Given u and z denote the normalized transmembrane potential and recovery variable at time t respectively, I_{stim} denotes the external

stimulus current, the physics governing system can be written as:

$$\frac{\partial u}{\partial t} = \nabla \cdot (\mathbf{D} \nabla u) + F(u, z) + I_{\text{stim}}, \quad \frac{\partial z}{\partial t} = G(u, z), \quad (3)$$

where

$$F(u, z) = k u(1 - u)(u - a) - uz, \quad G(u, z) = \epsilon(u)(ku - z). \quad (4)$$

On the other hand, under a one-step finite-difference scheme, we can approximate the temporal derivatives by

$$\frac{\partial u}{\partial t} \approx \frac{\hat{u}^{t+1} - u^t}{\Delta t}, \quad \frac{\partial z}{\partial t} \approx \frac{\hat{z}^{t+1} - z^t}{\Delta t}. \quad (5)$$

Moreover, to ensure a continuous chain of spatial gradients in the discrete space during the PDE loss calculation, we approximate the diffusion operator using a precomputed radial basis function finite difference (RBF-FD) [27] weights matrix $\nabla \cdot (\mathbf{D} \nabla u^t) \approx \mathbf{W} u^t$ where $\mathbf{W} \in \mathbb{R}^{|\mathcal{V}| \times |\mathcal{V}|}$ is a sparse operator matrix whose each entry w_{ij} represents the discretized contribution of node j to the diffusion term evaluated at node i . The stencil size is fixed to 20 neighboring cells selected by an m -KNN search in the 3D mesh space, and the conductivity tensor $\mathbf{D}$ is assumed to remain locally constant within each neighborhood. A unique $\mathbf{W}$ is precomputed for each mesh during the Graph Converter step and reused throughout training. Based on this discretization, the PDE and ODE residuals on the node set $\mathcal{I}$ are defined as:

$$R_u^t = \frac{\hat{u}^{t+1} - u^t}{\Delta t} - \mathbf{W} u^t - F(u^t, z^t) - I_{\text{stim}}^t, \quad R_z^t = \frac{\hat{z}^{t+1} - z^t}{\Delta t} - G(u^t, z^t). \quad (6)$$

Therefore, the training objective that combines data fidelity with explicit reaction-diffusion constraints is:

$$\mathcal{L} = \lambda_{\text{data}} \frac{1}{|\mathcal{V}|} \|\hat{u}^{t+1} - u_{\text{gt}}^{t+1}\|_2^2 + \lambda_{\text{PDE}} \frac{1}{|\mathcal{I}|} \sum_{i \in \mathcal{I}} (R_{u,i}^t)^2 + \lambda_{\text{ODE}} \frac{1}{|\mathcal{I}|} \sum_{i \in \mathcal{I}} (R_{z,i}^t)^2. \quad (7)$$

This formulation encourages the model to match the next-step target while remaining consistent with the underlying reaction-diffusion dynamics.

2.3 Multi-Step Rollout Training

Multi-Step Rollout Training is adopted to improve the model's ability to perform autoregressive prediction over longer horizons in practical rollout settings. Given an initial window of consecutive frames $\{\mathcal{G}^{t-T+1}, \dots, \mathcal{G}^t\}$, the model is recursively unrolled for R future steps; at rollout step r , it estimates $(\hat{u}^{t+r}, \hat{z}^{t+r})$, computes the corresponding loss against the target at the same step, and then feeds the updated frame back into the next input window. The total training objective is accumulated across the rollout horizon,

$$\mathcal{L}_{\text{rollout}} = \sum_{r=1}^R \mathcal{L}^{(t+r)}, \quad (8)$$

so that gradients are not only driven by the immediate one-step target but also propagated through later rollout steps, encouraging the network to better capture long-term temporal dynamics. However, to reduce early error accumulation, teacher forcing is applied at the beginning of training by appending either the ground-truth or predicted frame, and is then gradually decreased to enable fully autoregressive rollout.

3 Experiments

3.1 Dataset

We conduct experiments on a simulation dataset generated from 130 volumetric biventricular meshes reconstructed from MRI [6]. For each mesh, a patient-specific Purkinje tree and the corresponding Purkinje–Myocyte Junction (PMJ) points are generated following [4]. EP simulations are then performed using MonoAlg3D [3] for 600 ms with $dt = 0.02$ ms, on coarse meshes ($dx = dy = dz = 1200 \mu\text{m}$, 40,000–60,000 hexahedral elements/ mesh) for training/validation, and on finer meshes ($dx = dy = dz = 600 \mu\text{m}$, 300,000–550,000 elements/ mesh) for testing. The mono-domain system is solved with anisotropic conductivities $\sigma_f = 1.76 \times 10^{-4}$ and $\sigma_s = \sigma_n = 1.334 \times 10^{-4}$ ($\mu\text{m}^2/\text{ms}$), using the Aliev–Panfilov ionic model with default parameters $k = 8$, $a = 0.15$, $\epsilon_0 = 0.002$, $\mu_1 = 0.2$, and $\mu_2 = 0.3$ [1]. The dataset is split by trajectory with a train:val:test ratio of 8:1:1, ensuring that no mesh geometry is shared across the three subsets.

3.2 Implementation Details

Training Setup: The model is trained using Adam ($\text{lr} = 10^{-4}$, weight decay $= 10^{-5}$) with gradient clipping at 5.0. We use a latent dimension of 64, a 2-layer LSTM (dropout = 0.1), and $K = 10$ stacked GraphNet blocks. Training is performed for 100 epochs with batch size 2, using $T = 4$ history frames and a 4-step rollout ($\sim 0.5\text{h}/\text{epoch}$). After warm-up, the loss weights are set to $(\lambda_{\text{data}}, \lambda_{\text{PDE}}, \lambda_{\text{ODE}}) = (1, 0.3, 0.2)$; only the data term is used in the first 15 epochs, and the physics weights are linearly increased until epoch 30. Teacher forcing is reduced from epoch 20 and reaches zero at epoch 50, after which training is fully autoregressive. All experiments were performed on a single NVIDIA A40 GPU.

Autoregressive Testing: At test time, the model is evaluated in a fully autoregressive manner. Given $T_{\text{init}} = 20$ ground-truth frames, the last four frames are used to initialize the temporal hidden state, after which the model recursively predicts $(\hat{V}_m^{t+1}, \hat{z}^{t+1})$ and feeds each prediction back as the next input. The predicted voltage is clipped to $[-90, 35]$ mV for stability. The test set includes different geometries with different mesh resolutions to evaluate the generalization of the model.

Evaluation Metrics: For AP map prediction, we report the node-wise Root mean square error (RMSE) of V_m and the Pearson spatial correlation coefficient (PCC) between predicted and ground-truth voltage maps at each rollout

Table 1. Summary of performance of predicted activation time map at different future horizons. Time is reported with a frame interval of 2 ms compared to the start time.

Time (ms)	RMSE (V_m)	MAE (V_m)	PCC (V_m)	MSE (z)	MAE (z)	PCC (z)
2	0.48 (0.12)	0.28 (0.06)	1.00 (0.00)	0.01 (0.00)	0.01 (0.00)	0.85 (0.03)
6	0.75 (0.04)	0.48 (0.01)	1.00 (0.00)	0.02 (0.00)	0.01 (0.00)	0.75 (0.04)
10	1.59 (0.24)	0.97 (0.13)	1.00 (0.00)	0.03 (0.00)	0.02 (0.01)	0.66 (0.05)
20	8.47 (2.24)	4.97 (1.01)	0.96 (0.04)	0.05 (0.01)	0.04 (0.01)	0.54 (0.07)
40	21.03 (9.44)	13.42 (5.38)	0.75 (0.22)	0.11 (0.02)	0.08 (0.02)	0.42 (0.07)
100	13.80 (2.97)	12.08 (2.85)	0.33 (0.20)	0.27 (0.04)	0.21 (0.06)	0.34 (0.09)
200	20.74 (1.65)	17.32 (0.77)	0.45 (0.06)	0.84 (0.13)	0.62 (0.13)	0.35 (0.07)

step. For activation-time analysis, the activation time map is defined as the first threshold-crossing time of V_m at -60 mV, and evaluated using Mean Absolute Error (MAE) and RMSE over all valid activated nodes.

4 Results

Table 1 presents the performance of the proposed model in predicting activation time maps across various future time points. The model demonstrates strong predictive capability, particularly over short-to-moderate horizons, though performance predictably degrades for longer-term predictions as error accumulates. For near-term predictions (2–20 ms), the model exhibits high fidelity, with a membrane potential (V_m) PCC remaining at or near 1.00, indicating excellent alignment with the ground-truth wavefront morphology. The next-step potential MSE is 0.01 ± 0.00 for the recovery variable (z), underscoring the model’s ability to generalize to unseen 3D geometries without fine-tuning over a moderate rollout horizon (2–120 ms). As the prediction horizon extends, a gradual decline in performance is observed. Despite this, the spatial activation maps maintain moderate correlation (0.35 ± 0.07) even at 200 ms, suggesting the model retains some structural information over extended rollouts. Overall, the close alignment between predicted and in-silico ground-truth wavefronts (see Fig. 2), particularly during phases of rapid shape change, confirms that the model effectively captures the spatiotemporal electrophysiological dynamics of the biventricles.

As the rollout horizon increases, the error gradually accumulates, as expected in autoregressive prediction. However, we noticed a slight improvement around $t + 100$ to $t + 140$ ms likely occurs when most of the tissue has repolarized, reducing wavefront complexity and making the dynamics easier to predict. This trend suggests that model performance remains closely related to the underlying EP pattern complexity over time. To verify the role of the physics-informed constraints, we compare the proposed model with a purely data-driven variant in the Activation Time map predictions task. As Fig. 4 suggests, the Temporal Physics-Informed GNN shows a better performance, confirming the benefit of physical supervision. It also outperforms the result produced by AGATA reported in [18], achieving an average AT-MAE of 4.20 ± 3.79 ms. The role of




Fig. 2. Autoregressive rollout prediction on an unseen higher-resolution ($d = 600 \mu\text{m}$) biventricular mesh over time given the first 40 ms of simulation trajectory.




Fig. 3. Comparison between the physics-informed model (blue) and the purely data-driven model (orange) ($\lambda_{PDE} = \lambda_{ODE} = 0$).

the explicit reaction-diffusion loss component is further demonstrated in Fig. 3, where the physics-informed (PI) model is compared with an equally structured data-only baseline ($\lambda_{PDE} = \lambda_{ODE} = 0$), showing consistently better performance across all metrics. Additionally, our method just takes 0.128s on average to autoregressively produce an AP map given the only biventricle mesh in the training resolution, which is much faster compared to the FEM solvers, e.g., MonoAlg3D [3], which requires ~ 2.15 s per time step.

5 Conclusion

We proposed a novel physics-informed GNN framework to estimate future activation maps from an initial cardiac state. By integrating an explicit physics-constrained loss function derived from the monodomain model, our approach ensures that predictions remain grounded in realistic cardiac electrophysiology, enhancing both interpretability and generalization across different mesh geometries and resolutions. Experimental results demonstrate that the model can accurately predict electrophysiological dynamics on unseen meshes over moderate time horizons, offering a computationally efficient alternative to conventional simulation methods. Although long-horizon forecasting remains challenging and the model has not yet been validated on clinical electrophysiology data, this




Fig. 4. Illustration of the ground truth and predicted activation time map. The average MAE and RMSE of the proposed method were 4.20 ± 3.79 ms and 7.76 ± 6.04 ms.

work represents a meaningful step toward enabling real-time digital twins for cardiac electrophysiology.

Acknowledgments. This work was supported by Singapore National Medical Research Council Open Fund - Young Individual Research Grant (25-1321-A0001), Singapore Ministry of Education Tier 1 grant (25-1097-P0001), and NUS-UCL Research and Innovation Collaboration Fund (25-2338-A0001).

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

References

1. Aliev, R.R., Panfilov, A.V.: A simple two-variable model of cardiac excitation. *Chaos, Solitons & Fractals* **7**(3), 293–301 (1996)
2. Ayed, I., Cedilnik, N., Gallinari, P., Sermesant, M.: Ep-net: Learning cardiac electrophysiology models for physiology-based constraints in data-driven predictions. In: *International Conference on Functional Imaging and Modeling of the Heart*. pp. 55–63. Springer (2019)
3. Berg, L.A., Oliveira, R.S., Camps, J., Wang, Z.J., Doste, R., Bueno-Orovio, A., dos Santos, R.W., Rodriguez, B.: Monoalg3d: Enabling cardiac electrophysiology digital twins with an efficient open source scalable solver on gpu clusters. *bioRxiv* pp. 2025–04 (2025)
4. Berg, L.A., Rocha, B.M., Oliveira, R.S., Sebastian, R., Rodriguez, B., de Queiroz, R.A.B., Cherry, E.M., Dos Santos, R.W.: Enhanced optimization-based method for the generation of patient-specific models of purkinje networks. *Scientific Reports* **13**(1), 11788 (2023)
5. Clerc, L.: Directional differences of impulse spread in trabecular muscle from mammalian heart. *The Journal of physiology* **255**(2), 335–346 (1976)
6. Dillon, J.R., Mauger, C., Zhao, D., Deng, Y., Petersen, S.E., McCulloch, A.D., Young, A.A., Nash, M.P.: An open-source end-to-end pipeline for generating 3d+t biventricular meshes from cardiac magnetic resonance imaging. In: *Functional Imaging and Modeling of the Heart*. Lecture Notes in Computer Science, vol. 15673. Springer (2025)
7. Dominguez-Gomez, P., Zingaro, A., Baldo-Canut, L., Balzotti, C., Darpo, B., Morton, C., Vázquez, M., Aguado-Sierra, J.: Fast and accurate prediction of drug induced proarrhythmic risk with sex specific cardiac emulators. *npj Digital Medicine* **7**(1), 380 (2024)

8. Gutiérrez-Fernández, M., López-Linares, K., Fambuena-Santos, C., Guillem, M.S., Climent, A.M., Barquero-Pérez, Ó.: Deep learning for atrial electrogram estimation: toward non-invasive arrhythmia mapping using variational autoencoders. *Frontiers in Physiology* **16**, 1720244 (2025)
9. Kashtanova, V., Ayed, I., Cedilnik, N., Gallinari, P., Sermesant, M.: Ep-net 2.0: Out-of-domain generalisation for deep learning models of cardiac electrophysiology. In: *International Conference on Functional Imaging and Modeling of the Heart*. pp. 482–492. Springer (2021)
10. Kashtanova, V., Pop, M., Ayed, I., Gallinari, P., Sermesant, M.: Aphyn-ep: Physics-based deep learning framework to learn and forecast cardiac electrophysiology dynamics. In: *International Workshop on Statistical Atlases and Computational Models of the Heart*. pp. 190–199. Springer (2022)
11. Keener, J.P.: An eikonal-curvature equation for action potential propagation in myocardium. *Journal of mathematical biology* **29**(7), 629–651 (1991)
12. Li, L., Camps, J., Rodriguez, B., Grau, V.: Solving the inverse problem of electrocardiography for cardiac digital twins: A survey. *IEEE Reviews in Biomedical Engineering* **18**, 316–336 (2024)
13. Li, Z., Zheng, H., Kovachki, N., Jin, D., Chen, H., Liu, B., Azizzadenesheli, K., Anandkumar, A.: Physics-informed neural operator for learning partial differential equations. *ACM/IMS Journal of Data Science* **1**(3), 1–27 (2024)
14. Lim, B., Kim, J., Hwang, M., Song, J.S., Lee, J.K., Yu, H.T., Kim, T.H., Uhm, J.S., Joung, B., Lee, M.H., et al.: In situ procedure for high-efficiency computational modeling of atrial fibrillation reflecting personal anatomy, fiber orientation, fibrosis, and electrophysiology. *Scientific reports* **10**(1), 2417 (2020)
15. Lydon, H., Kazemi, M., Bishop, M., Paoletti, N.: Physics-informed neural operators for cardiac electrophysiology. *arXiv preprint arXiv:2511.08418* (2025)
16. Lydon, H., Kazemi, M., Bishop, M., Paoletti, N.: Physics-informed neural operators for cardiac electrophysiology. *arXiv preprint arXiv:2511.08418* (2025)
17. Martinez, E., Moscoloni, B., Salvador, M., Kong, F., Peirlinck, M., Lesley Marsden, A.: Full-field surrogate modeling of cardiac function encoding geometric variability (2025)
18. Morier, M., Rodríguez Padilla, J., Gallinari, P., Sermesant, M.: Learning cardiac electrophysiology with graph neural networks for fast data-driven personalised predictions. In: *International Conference on Functional Imaging and Modeling of the Heart*. pp. 138–150. Springer (2025)
19. Nazarov, I., Olakorede, I., Qureshi, A., Ogbomo-Harmitt, S., Aslanidi, O.: Physics-informed fully connected and recurrent neural networks for cardiac electrophysiology modelling. In: *2022 Computing in Cardiology (CinC)*. vol. 498, pp. 1–4. IEEE (2022)
20. Pfaff, T., Fortunato, M., Sanchez-Gonzalez, A., Battaglia, P.: Learning mesh-based simulation with graph networks. In: *International conference on learning representations* (2020)
21. 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 (2019)
22. 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)

23. Sanchez-Gonzalez, A., Heess, N., Springenberg, J.T., Merel, J., Riedmiller, M., Hadsell, R., Battaglia, P.: Graph networks as learnable physics engines for inference and control. In: International conference on machine learning. pp. 4470–4479. PMLR (2018)
24. Ten Tusscher, K.H., Panfilov, A.V.: Alternans and spiral breakup in a human ventricular tissue model. *American Journal of Physiology-Heart and Circulatory Physiology* **291**(3), H1088–H1100 (2006)
25. Trayanova, N.A., Lyon, A., Shade, J., Heijman, J.: Computational modeling of cardiac electrophysiology and arrhythmogenesis: toward clinical translation. *Physiological reviews* (2024)
26. Valderrábano, M.: Influence of anisotropic conduction properties in the propagation of the cardiac action potential. *Progress in biophysics and molecular biology* **94**(1-2), 144–168 (2007)
27. Xiang, Z., Peng, W., Yao, W., Liu, X., Zhang, X.: Solving spatiotemporal partial differential equations with physics-informed graph neural network. *Applied Soft Computing* **155**, 111437 (2024)