Accelerating Cardiac Myocyte Simulations with Physics-Informed Neural Networks for Sodium Channel Gating

Jorge Sánchez[1,2], Axel Loewe[1]

[1] Institute of Biomedical Engineering, Karlsruhe Institute of Technology (KIT), Karlsruhe, Germany
[2] iBIO, Universidad Internacional de Valencia, Valencia, Spain

Abstract

The increasing availability of human electrophysiological data has enabled the development of refined myocyte models with detailed transmembrane current descriptions. This fidelity is typically achieved via stiff systems of ordinary differential equations (ODEs) that span disparate time scales and force small integration steps, especially for the fast sodium current I_{Na} conducted by Nav1.5. Explicit schemes may require sub-microsecond steps (e.g., $0.1\,\mu\mathrm{s}$) to resolve upstroke and refractory dynamics, causing substantial computational costs for single-cell, tissue, and organ simulations. We present a hybrid approach in which a physics-informed neural network (PINN) surrogates the Nav1.5 gating kinetics for the activation m and dual inactivation gates h_1, h_2 and replaces the corresponding ODE subsystem in a human atrial action potential model. The PINN is trained to match the observed current-voltage relation while penalizing violations of the known gating ODE structure. We assess fidelity and quantify speedups up to 2.5x faster and stability with enlarged time steps in paced simulations (0.18 mV RMSE). The resulting surrogate reproduces gating-dependent I_{Na} behavior while materially reducing wall-clock time, suggesting a tractable path to high-fidelity, large-scale simulations and downstream personalization workflows.

1. Introduction

Cardiac cellular electrophysiology has advanced rapidly with high-quality experimental data from patch-clamp experiments, optical mapping, and clinical recordings. Modern myocyte models leverage these data to describe transmembrane currents and ion channel gating with growing biophysical detail. However, such fidelity is typically achieved through systems of stiff ordinary differential equations (ODEs) whose dynamics span several orders of magnitude in time. Explicit integrators can demand extremely small time steps to remain stable and accurate, which makes large-scale or long-duration simulations computationally expensive. The fast gating kinetics of the voltage-gated sodium channel Nav1.5 are a prime example: rapid state transitions can force sub-microsecond steps (e.g., $0.1\,\mu\mathrm{s}$) for explicit schemes to capture upstroke dynamics and refractory behavior. Implicit or adaptive solvers alleviate some of this burden but add algorithmic complexity and remain costly when embedded in tissue- or organ-level models, or during parameter estimation and personalization.

Data-driven surrogates offer a promising route to accelerate these computations, but purely black-box neural networks risk violating physical constraints and can generalize poorly beyond the training regime. Physics-informed neural networks (PINNs) incorporate governing equations and known structure directly into the training objective, encouraging solutions that satisfy the dynamics while retaining the expressive power of neural function approximators [1]. Ion channel gating is well suited to this paradigm because variables are bounded, voltage-dependent, and governed by ODEs with known qualitative behavior.

Here, we investigate a hybrid formulation in which a PINN models the gating kinetics (m, h_1, h_2) of Nav1.5, replacing the corresponding ODE subsystem in a human atrial action potential model (following Skibsbye et al. [2]. The surrogate is trained to satisfy gating differential relations across a range of voltages, with additional constraints enforcing physiological bounds and smoothness. We evaluate the approach on the single-cell level under voltage-clamp and action potential-clamp protocols and on the tissue level where sodium kinetics strongly influence excitability, conduction, and refractory properties. We examine fidelity to the reference model, stability under larger time steps, and computational efficiency.

2. Methods

2.1. Model overview

We construct a PINN that maps transmembrane voltage V_m to the three gating variables of Nav1.5: activation m and dual inactivation gates h_1, h_2. The fast sodium current

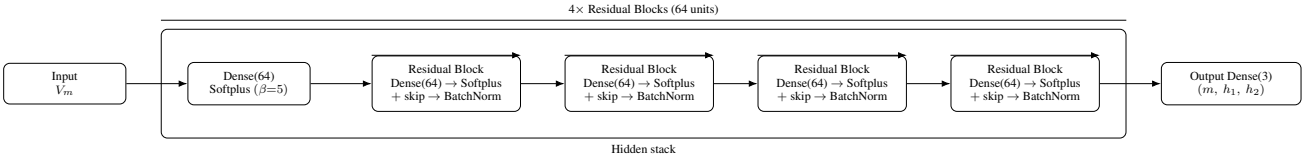

Figure 1. Architecture used to predict sodium-channel gates from transmembrane voltage. The network takes a scalar input V_m, applies a 64-unit Dense layer, then a stack of four residual blocks (each: Dense(64)→Softplus(β=5) with a skip connection and BatchNorm), and finally a Dense(3) that outputs the Nav1.5 gating triplet (m, h_1, h_2). The top bracket annotates the 4× residual-block stack; arrows indicate forward data flow.

is computed as

$$I_{\mathrm{Na}}(V_m) = \bar{G}_{\mathrm{Na}}\, m^3 h_1 h_2 \big(V_m - E_{\mathrm{Na}}\big), \tag{1}$$

where $\bar{G}_{\mathrm{Na}}$ is the maximal conductance and E_{Na} is the sodium reversal potential.

2.2. Reference relations and target current

Voltage-dependent steady states and time constants were specified for activation and inactivation:

$$m_\infty(V_{\mathrm{m}}),\ \tau_m(V_{\mathrm{m}}), \qquad h_\infty(V_{\mathrm{m}}),\ \tau_{h_1}(V_{\mathrm{m}}),\ \tau_{h_2}(V_{\mathrm{m}}),$$

including Q10/temperature adjustments at $T = 310.15\,\mathrm{K}$. The sodium reversal potential was computed from intracellular and extracellular sodium concentrations ($\mathrm{Na_i}$=7.7869, $\mathrm{Na_o}$=130) via

$$E_{\mathrm{Na}}\,[\mathrm{mV}] = 10^3\,\frac{\mathrm{RT}}{\mathrm{F}}\,\ln\!\left(\frac{\mathrm{Na_o}}{\mathrm{Na_i}}\right). \tag{2}$$

With $\bar{G}_{\mathrm{Na}} = 576$, a target current for supervised learning was generated on a uniform grid of $N = 1000$ collocation voltages spanning $V_m \in [-100, 100]\,\mathrm{mV}$ by inserting steady states into (1):

$$I_{\mathrm{Na}}^{\mathrm{target}}(V_{\mathrm{m}}) \;=\; \bar{G}_{\mathrm{Na}}\left[m_\infty(V_{\mathrm{m}})\right]^3 h_\infty(V_{\mathrm{m}})\, h_\infty(V_{\mathrm{m}})\,\big(V_{\mathrm{m}} - E_{\mathrm{Na}}\big). \tag{3}$$

2.3. PINN architecture

The PINN was fully implemented in Python using Pytorch [3]. The architecture is a fully connected network shown in Figure 1. The input of the network is the transmembrane voltage at each time step and the output are predicted values for the Nav1.5 channel gates.

2.4. Physics-informed objective

Let $(m, h_1, h_2) = \mathcal{N}_\theta(V_{\mathrm{m}})$ denote network predictions at a collocation voltage V_{m}. The composite objective

$$\mathcal{L} \;=\; \mathcal{L}_{\text{data}} \;+\; \lambda_{\text{phys}}\,\mathcal{L}_{\text{phys}} \;+\; \lambda_{\text{grad}}\,\mathcal{L}_{\text{grad}} \tag{4}$$

includes the loss components described in the following.

2.4.1. Data loss ($\mathbf{I_{Na}}$ current residual)

$$\mathcal{L}_{\mathrm{data}} = \Big\| \bar{G}_{\mathrm{Na}}\, m^3 h_1 h_2\,(V_{\mathrm{m}} - E_{\mathrm{Na}}) - I_{\mathrm{Na}}^{\mathrm{target}}(V_{\mathrm{m}}) \Big\|_2^2. \tag{5}$$

2.4.2. Physics loss (gating ODE residuals)

We regularize toward first-order kinetics

$$\begin{aligned} \frac{dm}{dt} &= \frac{m_\infty(V_{\mathrm{m}}) - m}{\tau_{\mathrm{m}}(V_{\mathrm{m}})}, \\ \frac{dh_1}{dt} &= \frac{h_\infty(V_{\mathrm{m}}) - h_1}{\tau_{\mathrm{h_1}}(V_{\mathrm{m}})}, \\ \frac{dh_2}{dt} &= \frac{h_\infty(V_{\mathrm{m}}) - h_2}{\tau_{\mathrm{h_2}}(V_{\mathrm{m}})}. \end{aligned} \tag{6}$$

Using automatic differentiation with respect to V_m at training points,

$$\mathcal{R}_m = \frac{dm}{dV_{\mathrm{m}}} - \frac{m_\infty(V_{\mathrm{m}}) - m}{\tau_m(V_{\mathrm{m}})}, \tag{7}$$

$$\mathcal{R}_{h_1} = \frac{dh_1}{dV_{\mathrm{m}}} - \frac{h_\infty(V_{\mathrm{m}}) - h_1}{\tau_{h_1}(V_{\mathrm{m}})}, \tag{8}$$

$$\mathcal{R}_{h_2} = \frac{dh_2}{dV_{\mathrm{m}}} - \frac{h_\infty(V_{\mathrm{m}}) - h_2}{\tau_{h_2}(V_m\mathrm{m}}, \tag{9}$$

and

$$\mathcal{L}_{\text{phys}} \;=\; \|\mathcal{R}_{\mathrm{m}}\|_{\in}^{\in} + \|\mathcal{R}_{\mathrm{h_1}}\|_{\in}^{\in} + \|\mathcal{R}_{\mathrm{h_2}}\|_{\in}^{\in}. \tag{10}$$

Remark: Under prescribed voltage protocols, time and voltage derivatives relate via the chain rule $d(\cdot)/dt = (d(\cdot)/dV_m)\,(dV_m/dt)$.

2.4.3. Smoothness penalty

A small penalty discourages oscillatory solutions:

$$\mathcal{L}_{\text{grad}} \;=\; \left\| \frac{dm}{dV_{\mathrm{m}}} \right\|_2 + \left\| \frac{dh_1}{dV_{\mathrm{m}}} \right\|_2 + \left\| \frac{dh_2}{dV_{\mathrm{m}}} \right\|_2. \tag{11}$$

We set $\lambda_{\text{grad}} = 0.01$ and use a dynamic physics weight

$$\lambda_{\text{phys}} = \min\Big(0.1,\ \frac{\mathcal{L}_{\text{data}}}{100}\Big), \tag{12}$$

so the physics term gains relative weight as current mismatch decreases.

2.5. Training

We used $N = 1000$ collocation voltages in full-batch mode and the AdamW optimizer. A cosine-annealing-with-restarts scheduler was applied initially, followed by a plateau scheduler to reduce the learning rate when the loss stalled. Early stopping was used so that training continued until the composite loss satisfied $\mathcal{L} < 10^{-4}$.

3. Results

3.1. Single-cell fidelity

Across the test grid, the PINN-based current closely matched the supervised I–V relation, preserving the sharp rise of I_{Na} near activation thresholds and the attenuation associated with inactivation. Predicted gates were smooth and monotonic in physiologically expected regions, and no spurious oscillations were observed. Figure 2 shows the predictions made by the PINN model for the I_{Na} gates. Figure 2A shows the prediction for the m gate where there is no difference compared to the ODE dynamics. Figure 2B+C show the prediction for the inactivation gates. The predictions by the PINN differ from the ODE especially for the h_2 gate.

The hybrid PINN-ODE model reproduced the action potential morphology with high fidelity (RMSE $0.18\,\text{mV}$ over the full trace). A summary of the action potential properties are shown in Table 1. The upstroke timing and overshoot had a small deviation (4 V/s) from the reference, and both models returned to the same resting potential. The largest visible discrepancies occurred around the plateau and early repolarization, where the PINN trace was slightly more negative, consistent with a mild acceleration of repolarization (Figure 3). Late repolarization aligns again closely with an APD_{90} difference of $6.6\,\text{ms}$. These small, localized errors are compatible with residual inaccuracies in the learned inactivation kinetics and could in future be further reduced by incorporating action potential-clamp sequences or emphasizing loss terms near plateau voltages during training.

3.2. Computational performance

The ODE model was integrated using a Rush-Larsen scheme for the gating ODEs and Euler for the rest of the ODEs. In paced single-cell simulations, the hybrid model reduced wall-clock time by a factor of $2.1\times$ at a baseline step of $10\,\mu\text{s}$. Moreover, stable simulations were obtained with a tenfold larger time step (from $10\,\mu\text{s}$ to $100\,\mu\text{s}$), while maintaining a root-mean-square error below $0.25\,\text{mV}$ during repolarization over 100 stimuli at $1\,\text{Hz}$.

Biomarker	Ground truth ODE	PINN-ODE
APD_{90} (ms)	283.8	284.4
RMP (mV)	−74.8	−74.8
dV/dt_{max} (V/s)	217.5	213.4
APA (mV)	103.0	103.0

Table 1. Action potential properties at $1\,\text{Hz}$. Biomarker values are shown for the mechanistic ground truth ODE model and the hybrid PINN-ODE model.

3.3. Action potential restitution

Because refractoriness is strongly governed by sodium inactivation, we examined restitution (Figure 4). The hybrid approach preserved key action-potential features, with a maximum deviation of $25\,\text{ms}$ in the restitution curve for pacing rates below $2.5\,\text{Hz}$. Deviations were primarily attributable to residual errors in the predicted inactivation kinetics.

4. Discussion

The results indicate that a physics-informed surrogate for Nav1.5 gating can largely reproduce the reference current while enabling materially larger time steps. Embedding the surrogate in the membrane model eliminates the stiffest subsystem without discarding mechanistic structure elsewhere, thus preserving interpretability and compatibility with existing pipelines. Compared with blackbox alternatives, the physics terms constrain learning to adhere to known kinetics, improving robustness when extrapolating across voltage protocols.

5. Limitations and future work

First, supervised training used voltage-collocation points with currents synthesized from steady-state relations; incorporating dynamic clamp data and time-resolved protocols should further improve fidelity. Second, outputs were unconstrained; sigmoidal squashing or barrier terms could enforce $m, h_i \in [0, 1]$. Finally, multi-current surrogates (e.g., I_{Kr}, I_{CaL}) or intracellular ion concentration (e.g. Ca^{2+}) could yield additional speedups for tissue and organ simulations.

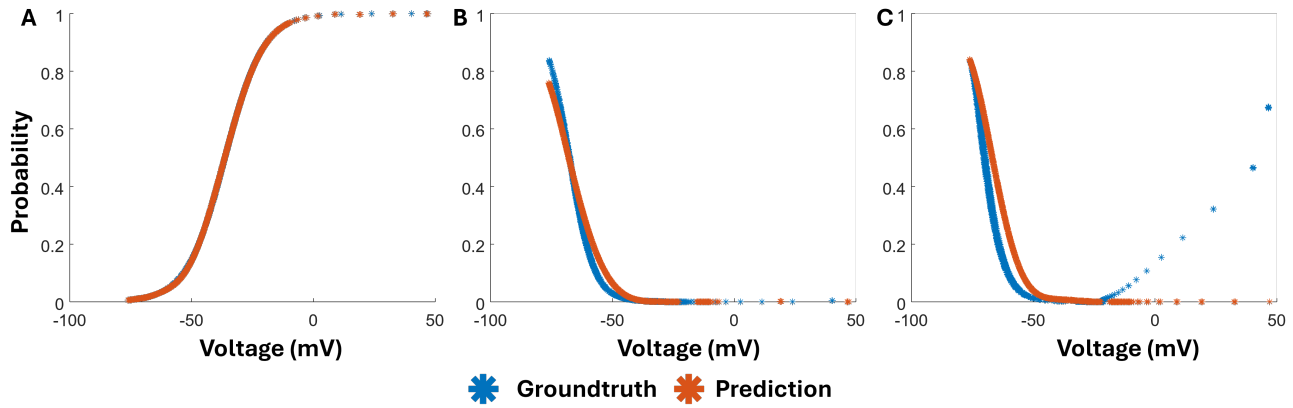

Figure 2. Gate open probabilities for the ODE (ground truth) and the prediction of the PINN model. A) Traces for the m gate. B) Traces for the h_1 gate. C) Traces for the h_2 gate

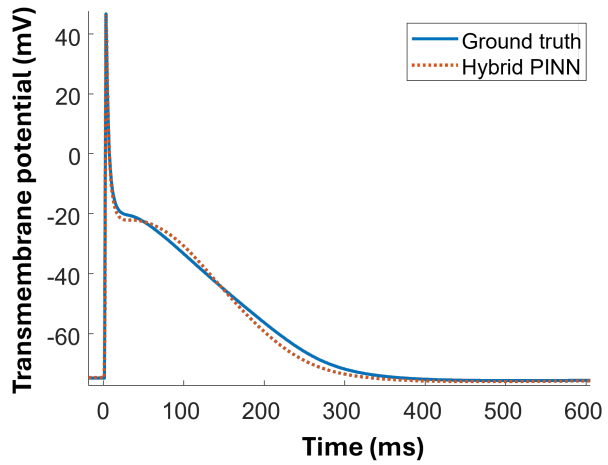

Figure 3. Transmembrane potential of the hybrid PINN-ODE model at 1 Hz. The traces correspond to the last action potential out of 100.

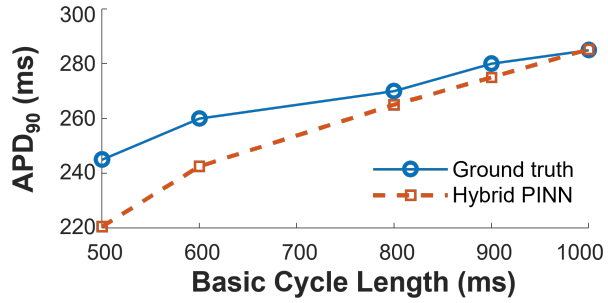

Figure 4. Restitution curve of APD at 90% repolarization.

6. Conclusion

A PINN surrogate for Nav1.5 gating reproduces key I_{Na} features while substantially reducing computational cost and permitting larger time steps in paced simulations. This modular replacement of a known stiffness bottleneck makes high-fidelity electrophysiology models more tractable for large-scale in silico experimentation and personalization.

Acknowledgments

Jorge Sánchez is partially supported by the Karlsruhe Institute of Technology Excellence Strategy via the Young Investigator Group Preparation Program.

References

[1] Raissi M, Perdikaris P, Karniadakis G. Physics-informed neural networks: A deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations. Journal of Computational Physics 2019; 378:686–707.

[2] Skibsbye L, Jespersen T, Christ T, Maleckar MM, van den Brink J, Tavi P, Koivumäki JT. Refractoriness in human atria: Time and voltage dependence of sodium channel availability. Journal of Molecular and Cellular Cardiology 2016;101:26–34. ISSN 00222828.

[3] Paszke A, Gross S, Chintala S, Chanan G, Yang E, DeVito Z, Lin Z, Desmaison A, Antiga L, Lerer A. Automatic differentiation in PyTorch. In Conference on Neural Information Processing Systems, volume 31. 2017; .

Address for correspondence:

Jorge Sánchez
Institute of Biomedical Engineering, KIT
Kaiserstr. 12, 76131 Karlsruhe, Germany
jorsana4@upv.edu.es