

AdaINR: Locally Adaptive Implicit Neural Representation of Diffusion Currents for Mechanistic Electrophysiology Simulations

Sumeet Vadhavkar¹, Stephanie Appel², Cristian Barrios Espinosa², Axel Loewe², Casey Meisenzahl¹
Linwei Wang¹

¹ Rochester Institute of Technology (RIT), Rochester, New York, USA

² Institute of Biomedical Engineering, Karlsruhe Institute of Technology (KIT), Karlsruhe, Germany

Abstract

Mechanistic electrophysiology simulations provide detailed insights into arrhythmia mechanisms but are hindered by high computational cost. Recent emergence of reaction-eikonal models rely on efficient surrogates for the diffusion current to maintain scalability, although existing approximations struggle to generalize across heterogeneous conduction properties and complex arrhythmic dynamics. This paper presents adaINR – a novel implicit neural representation (INR) for the diffusion current that is adaptive to local conduction characteristics. Trained and tested on data from a simulated reentry scenario in cardiac tissue, we demonstrated the ability of adaINR to efficiently capture the key features of diverse diffusion current morphologies - including planar, curved, and colliding wavefronts - with an average mean squared error of 33.8 ($\mu\text{A}/\text{cm}^2$)². It has the potential to facilitate faster and more scalable simulations of complex arrhythmias.

1. Introduction

Cardiac electrophysiology (EP) simulations have become valuable tools for understanding arrhythmia mechanisms and supporting patient-specific therapies [1, 2]. These simulations are typically based on the monodomain model, which resolves the spatiotemporal dynamics of the transmembrane voltage V_m across cardiac tissue. A major computational bottleneck is the diffusion term due to the need for fine spatiotemporal grids to resolve rapid, localized changes in diffusion currents, which are essential for accurately capturing wavefront propagations.

Reaction-eikonal models have emerged to approximate wave propagation by solving the eikonal equation for activation times [3]. These activation times serve as a basis for approximating the diffusion term, allowing simulations on coarser meshes with substantially reduced runtime. One representative is the Diffusion-Reaction Eikonal Alternant Model (DREAM) [4]. DREAM replaces the dif-

fusion term with a time-shifted triple-Gaussian waveform aligned to the activation time. This surrogate captures the typical biphasic shape of diffusion currents (DCs) during planar wavefront propagation. However, amplitude and time course vary in complex arrhythmias and tissue structures. For example, for curved wavefronts, the shape becomes broader and more variable with usually smaller amplitudes; for colliding wavefronts, the course becomes almost monophasic with a dominant positive peak and little to no negative component. While DREAM effectively models planar wavefronts, its fixed Gaussian template is less suited for highly fragmented or reentrant wavefronts.

In this work, we propose a neural surrogate model for the DC to replace template-based approximations in hybrid EP models while retaining biophysical ionic dynamics. Specifically, we introduce adaINR with two innovations: (i) we represent DC as a continuous function over time t (I_{diff}) using an implicit neural representation (INR), enabling flexible approximation without rigid spatial or temporal discretization; and (ii) we adaptively generate the weight parameters of the INR with a hypernetwork based on local conduction information, including local activation times (LATs) and diastolic intervals (DIs). As a proof of concept, we trained and tested adaINR on simulated data of reentrant wave propagation in cardiac tissue, demonstrating its ability to capture diverse and physiologically relevant DC morphologies.

2. Methods

Fig 1 outlines the key components in adaINR, which includes a compact INR representation of I_{diff} adapted by a hypernet based on local features.

2.1. Graph convolution

To generalize across different geometries, the tissue was represented as a graph, where each node i is associated with a feature vector $\mathbf{x}_i \in \mathbb{R}^d$, containing spatial coordinates, LAT, DI, and a binary indicator for boundary nodes.

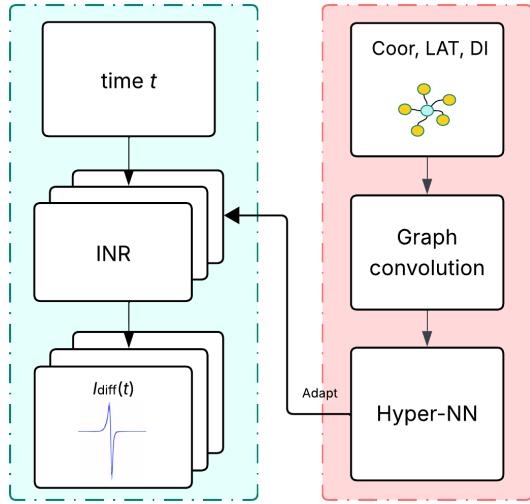


Figure 1: Schematic of adaINR: graph convolution processes local features of a node and its neighbors (coordinates, LAT, DI) to generate an embedding, which a hypernetwork maps to the weights of an implicit neural representation (INR) to predict I_{diff} at any time t .

Let \mathcal{E} be the set of edges between neighboring nodes i and j , with edge attributes e_{ij} representing the difference in the LAT between the nodes. We applied a spline-based graph convolution to incorporate local neighborhood information. For each node i , an embedding \mathbf{c}_i was computed by aggregating features across its neighbors:

$$\mathbf{c}_i = \rho \left(\sum_{j \in N(i)} B(e_{ij}) W \mathbf{x}_j \right), \quad (1)$$

where $N(i)$ are the neighbors of node i , W is a learnable weight matrix, B is a B-spline basis function that transforms the edge attribute e_{ij} into a modulating factor, and $\rho(\cdot)$ is a non-linear activation function. This embedding captures local conduction patterns and enables the model to account for spatial dependencies, while naturally adapting to complex tissue geometry.

2.2. INR representation of I_{diff}

To estimate I_{diff} continuously in time, we used an INR, which models $I_{\text{diff}}(t)$ as a function that can be queried at arbitrary time points. This enabled flexible modeling of rapid waveform changes without the need for time-stepping or large memory storage. Specifically,

$$I_{\text{diff}}(t) = f_{\theta}(t), \quad (2)$$

where f_{θ} is a 3-layer MLP with weights θ . The first layer maps the scalar input t to 64 dimensions with Leaky ReLU activation. The second layer reduces it to 32 dimensions (also with Leaky ReLU). The final layer projects to a scalar bounded by tanh activation.

2.3. Hypernetwork for Adaptive I_{diff}

To account for spatial variability in I_{diff} morphology, the INR weights θ were generated per node using a hypernetwork conditioned on the local embedding \mathbf{c}_i :

$$\theta = h_{\phi}(\mathbf{c}_i), \quad (3)$$

where h_{ϕ} is a 5-layer MLP with Leaky ReLU activations and weights ϕ . This setup enabled the model to adapt DC predictions to local conduction properties, improving accuracy in heterogeneous regions.

2.4. Training Objective

To train the model to capture the full range of I_{diff} morphologies, we need to ensure that representative I_{diff} morphologies receive adequate attention from the model during training. To this end, we first apply principal component analysis and K-means clustering to the I_{diff} data available for training, to group waveforms into distinct clusters of morphologies. As seen in Fig. 2, clusters 5 and 6 are located in regions far from the center of the reentrant circuit, where wavefronts propagate nearly planar and produce sharp, biphasic DC profiles. Clusters 1–4 lie closer to the core, where wavefronts become increasingly curved, resulting in broader, lower-amplitude, and more variable signals. Cluster 7 is found along tissue boundaries and reflects wavefront collisions, while cluster 0 corresponds to subthreshold activity with minimal temporal variation.

Capturing the locally varying morphologies in clusters 1–4 and 7 is key to advancing beyond template-based surrogates such as the one used in DREAM. These clusters, which we defined as regions of interest (ROIs), are critical for accurately simulating reentrant dynamics but were underrepresented in the present dataset. To address this imbalance, we adopted two training strategies: (i) construct mini-batches that include samples from all clusters, and (ii) use a weighted mean squared error ($wMSE$) loss that scales each sample inversely to its cluster size:

$$wMSE = \frac{1}{N} \sum_{i=1}^N w_{c_i} \left[\frac{1}{T} \sum_{t=1}^T (\hat{I}_{\text{diff}}^{(i)}(t) - I_{\text{diff}}^{(i)}(t))^2 \right]. \quad (4)$$

where T is the number of sampled time points, N is the number of training samples, and $\hat{I}_{\text{diff}}(t)$ the model prediction. The weight w_{c_i} is inversely proportional to the size of the cluster node i belongs to, emphasizing spatially rare but physiologically important morphologies.

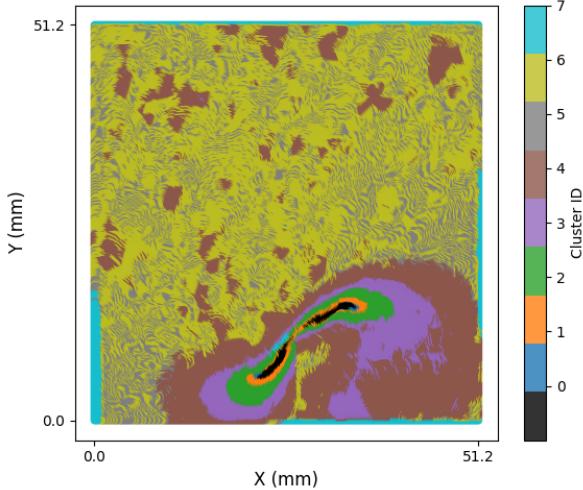


Figure 2: Spatial cluster map for a representative activation cycle. Clusters reflect distinct I_{diff} morphologies. Nodes without activation in the current cycle are shown in black.

3. Experiments and Results

3.1. Data & Implementation Details

The approach was evaluated on a 2D cardiac tissue dataset generated using the openCARP simulator [5] with an S1S2 pacing protocol. The $51.2 \text{ mm} \times 51.2 \text{ mm}$ tissue slab was discretized into 72,449 nodes at $200 \mu\text{m}$ resolution. The dataset had up to ten activation cycles per node, including both, planar and reentrant wavefronts. Each node was annotated with spatial coordinates (x, y) , neighboring nodes, I_{diff} during the depolarization phase of V_{m} , LATs defined as the time when V_{m} crosses -35 mV with positive slope, and DI measured as the interval between V_{m} drops below -40 mV with negative slope and the LAT of the next activation. Each node-activation pair was treated as an independent training instance.

For validation and testing, five disconnected regions of 1,000 nodes each were randomly sampled. The rest were used for training. To facilitate numerical stability and robust gradient flows during training, we normalized all input features including spatial coordinates, LAT, and DI, as well as I_{diff} amplitudes. We used the Adam optimizer with a learning rate of 10^{-3} , and the best model was selected based on validation loss. adaINR was evaluated using node-level mean squared error (MSE) and the correlation coefficient (CC) between $\hat{I}_{\text{diff}}(t)$ and $I_{\text{diff}}(t)$.

3.2. Results

An important assumption for our locally adaptive INR is that the temporal morphology of a I_{diff} can be inferred

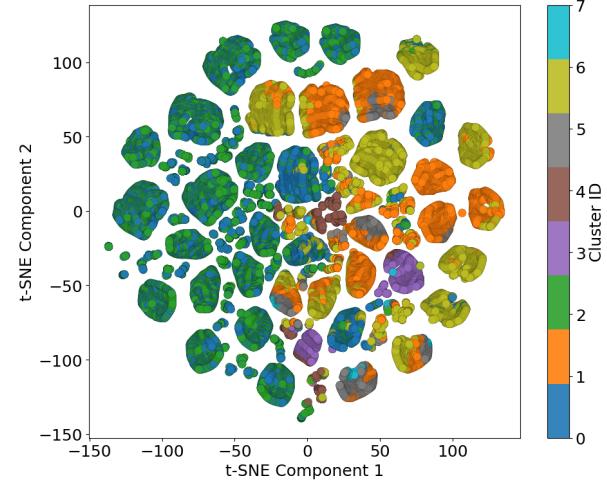


Figure 3: t-SNE visualization of node embeddings learned via local graph convolution, with points labeled according to I_{diff} morphology clusters.

Table 1: Mean squared error (MSE) and correlation coefficient (CC) of adaINR predictions across I_{diff} morphology clusters, reported as mean \pm std. Best (blue) and worst (red) performing clusters are highlighted.

Cluster	MSE $(\mu\text{A}/\text{cm}^2)^2$	CC
0	98.36 ± 107.35	0.17 ± 0.23
1	114.52 ± 62.49	0.67 ± 0.19
2	29.51 ± 23.98	0.95 ± 0.04
3	34.79 ± 23.95	0.95 ± 0.03
4	37.71 ± 31.86	0.97 ± 0.05
5	52.56 ± 27.18	0.97 ± 0.02
6	47.59 ± 29.22	0.98 ± 0.02
7	64.79 ± 58.64	0.89 ± 0.09

from local conduction features and the spatial relationship of a node and its neighbors. To verify this, we examined the embedding \mathbf{c}_i extracted via graph convolution across all nodes (see Fig 3). Clear separation was observed between the I_{diff} morphology clusters identified in the data. ROIs were well distinguished from planar wave clusters 5 and 6, with some overlap between clusters 1 and 5 near the rotor. This suggests that the embeddings effectively capture local morphological differences and are well-suited for generating INR weights to model diverse diffusion current profiles.

Table 1 summarizes the quantitative results of adaINR in I_{diff} prediction, and Fig. 4 provides visual examples. As shown, adaINR accurately reconstructed sharp, biphasic I_{diff} signals of planar wavefronts (cluster 6), capturing relatively high-amplitude and well-defined waveform profiles with low MSE and the highest CC value. It captured lower-

amplitude, curved morphologies (clusters 1-4) mostly with similarly strong performance, represented as high CC and low MSE values (lowest MSE in cluster 2, second closest to the rotor). However, cluster 1 posed the greatest challenge, with occasional underestimation of peaks or broader waveform transitions, leading to higher MSE and lower CC values. Prediction in the boundary regions with wave-front collisions (cluster 7) also proved difficult, the predominantly monophasic morphology remained captured. adaINR’s forward pass requires only 6.1 s to generate predictions for a test set of 5,000 nodes across 10 activation cycles, measured on a single commodity GPU.

4. Discussion

This work presented adaINR as a novel surrogate for diverse DC morphologies. With local graph convolution, adaINR can be directly applied to a node and its neighbors

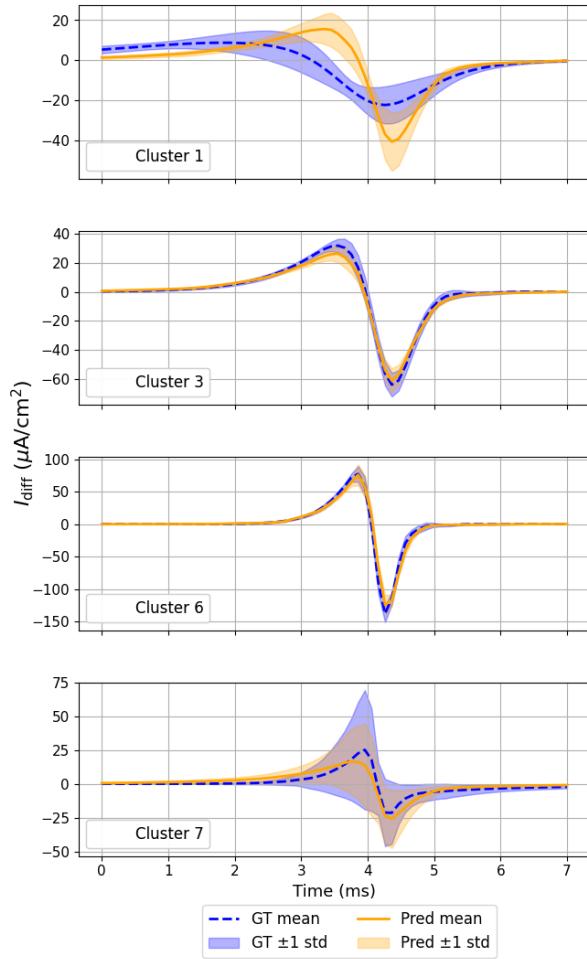


Figure 4: Mean \pm standard deviation of the ground truth (GT) and predicted (Pred) I_{diff} for representative clusters.

on different geometries. With feedforward hypernetwork and the compact INR, adaINR is computationally efficient. It thus provides an important building block for hybrid neural-mechanistic EP models to support high-resolution mechanistic investigation of complex arrhythmias.

As a proof of concept, this work considered simulation data on a 2D grid of cardiac tissue, and the training of adaINR assumed access to true I_{diff} signals for supervision. Future works will consider complex 3D cardiac geometries, and incorporate adaINR into a mechanistic EP model where direct supervision of I_{diff} is not available; instead adaINR needs to work with Eikonal calculation of activation times and mechanistic formulation of ionic currents to resolve transmembrane voltages.

Acknowledgments

This research was funded in part by the National Institutes of Health award R01HL145590, National Science Foundation award OAC-2212548, and the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) – Project-ID 258734477 – SFB 1173.

References

- [1] Heijman J, Sutanto H, Crijns HJ, Nattel S, Trayanova NA. Computational models of atrial fibrillation: Achievements, challenges, and perspectives for improving clinical care. *Cardiovascular Research* 2021;117(7):1682–1699.
- [2] Loewe A, Hunter PJ, Kohl P. Computational modelling of biological systems now and then: revisiting tools and visions from the beginning of the century. *arXiv preprint arXiv:250113142* 2025;.
- [3] Neic A, Campos FO, Prassl AJ, Niederer SA, Bishop MJ, Vigmond EJ, Plank G. Efficient computation of electrograms and ecgs in human whole heart simulations using a reaction-eikonal model. *Journal of computational physics* 2017;346:191–211.
- [4] Barrios Espinosa C, Sánchez J, Appel S, Becker S, Krauß J, Martínez Díaz P, Unger L, Houillon M, Loewe A. A cyclical fast iterative method for simulating reentries in cardiac electrophysiology using an eikonal-based model. *Engineering with Computers* 2025;1–24.
- [5] Plank* G, Loewe* A, Neic* A, Augustin C, Huang YLC, Gsell M, Karabelas E, Nothstein M, Sánchez J, Prassl A, Seemann* G, Vigmond* E. The openCARP simulation environment for cardiac electrophysiology. *Computer Methods and Programs in Biomedicine* 2021;208:106223.

Address for correspondence:

Sumeet Vadhwkar
718 Park Point dr, Rochester, NY, 14623
sv6234@rit.edu