

Characterising cardiac tissue properties with graph neural networks

Ching-En Chiu¹[0009-0006-4419-0522], Yoo Ri Kim²[0000-0003-3897-5869], Magdi Saba^{3,4}, Danilo Mandic¹[0000-0001-8432-3963], and Marta Varela^{4,5}[0000-0003-4057-7851]

¹ Department of Electrical and Electronic Engineering, Imperial College London, London, United Kingdom

² Department of Internal Medicine, Chonnam National University College of Medicine, Gwangju, Republic of Korea

³ St. George's University Hospital, London, United Kingdom

⁴ Cardiovascular & Genomics Research Institute, City St George's University of London, United Kingdom

⁵ National Heart & Lung Institute, Imperial College London, London, United Kingdom

ching-en.chiu18@imperial.ac.uk

mvarela@citystgeorges.ac.uk

Abstract. Characterising electrophysiological properties of cardiac tissue efficiently and accurately from spatially sparse intracardiac measurements is clinically important for localising ablation targets and improving arrhythmia treatment. We developed a graph neural network-based framework trained on synthetic electrogram signals on 2D flat surfaces to identify areas of interest in the context of cardiac ablation for premature ventricular complexes (PVCs). Our method achieved an average precision of 0.96, 0.97, and 0.95 for the detection of single-patch fibrosis, rapid depolarisation and high excitability, respectively. The trained model can then be applied to 2D curved surfaces with few-shot fine-tuning, demonstrating its generalisation capability. Future work will develop this framework further for clinical use in PVC ablation.

Keywords: Cardiac electrophysiology · Graph neural networks · Intracardiac electrograms.

1 Introduction

Premature ventricular complexes (PVCs) are amongst the most common arrhythmias, detected in 40–75% of apparently healthy individuals on ambulatory Holter monitoring and affecting approximately 1–4% of the general population [17]. They are characterised by early ventricular depolarisations that originate outside the normal conduction pathway. A small proportion of patients have symptoms severe enough to require treatment or a sufficiently high PVC burden to raise concerns about PVC-induced cardiomyopathies.

Cardiac ablation, a common treatment option, involves creating a focal lesion that eliminates or electrically isolates the tissue responsible for the initiation or propagation of the abnormal electrical signals. The arrhythmogenic substrate is typically a small focal region of viable myocardium with abnormal electrophysiological properties (such as abnormal calcium handling or enhanced catecholamine sensitivity). In addition, PVCs in structural heart disease may arise at scar border zones, where surviving myocyte bundles, fibrosis, slow or heterogeneous conduction, and altered cell-to-cell coupling facilitate abnormal automaticity or re-entry.

During the ablation procedure, catheters equipped with electrodes are typically inserted into the heart to measure electrical signals at the endocardium, known as contact electrograms (EGMs). Current approaches to identify ablation targets from EGMs in a clinical setting remain largely empirical and rely heavily on the operator’s expertise. A common strategy for PVCs is *activation mapping*, which identifies the earliest local activation site. The ablation outcome is often limited by infrequent intraprocedural PVCs, however. When PVCs are absent, *pace mapping* may be adopted. It relies on the assumption that pacing at the PVC origin should produce waveforms similar to the clinical PVC [11]. Nevertheless, PVCs are often not inducible even with deliberately pacing, with [5] reporting only 7% of reproducibly inducible PVCs amongst patients with infrequent PVC during procedures. It is therefore useful to be able to locate ablation targets based on tissue properties during sinus rhythm. This would increase the procedure’s success and reduce procedural time by limiting the need for extensive mapping.

Aim This study aims to characterise the underlying electrophysiological properties of the myocardium from spatially sparse EGM signals, using convolutional (CNNs) and graph neural networks (GNNs). Our framework identifies, at high spatial resolution, regions of fibrosis, abnormal repolarisation and enhanced automaticity. The proposed method may help improve ablation procedures for PVCs in the future.

CNNs excel at signal classification tasks, including electrocardiogram (ECG) analysis [15], whereas GNNs are a specialised class of deep learning technique used on data represented as graphs. GNNs leverage *message-passing* operations where each node iteratively exchanges information with its neighbours and update its own state. They have been used for cardiac electrophysiology applications, in which cardiac geometries are represented as graphs. For example, reconstructing atrial fibrillation dynamics from sparse contact mapping measurements [14] and modelling action potential propagation [16]. The latter work particularly highlights GNNs’ ability to generalise to complex geometries after only being trained on simple ellipsoid meshes, hinting at fast patient-specific cardiac applications.

70-80% of idiopathic PVCs originate in the Right Ventricular Outflow Tract (RVOT) [7]. As RV is relatively thin, it is also well suited for the development and testing of methods that encode the geometry of myocardial tissue with triangular meshes, a type of graph.

Several automated approaches have been proposed to support PVC ablation. These methods can improve the reproducibility and efficiency of activation mapping, but they remain dependent on the presence of PVCs during the procedure and primarily aim to localise the ectopic origin without necessarily characterising the underlying electrophysiological properties. This is done either coarsely, using the surface ECG [21] or, at higher resolution, from intraprocedural EGMs. Examples of EGM-based automation include automated activation mapping, in which algorithms are developed to annotate local activation times using fixed features such as bipolar EGM onset or the maximal negative slope of the unipolar EGM [2,1,13]. In contrast, our work aims to identify pathological regions associated with abnormal electrophysiological properties, providing a potential route towards substrate-informed ablation planning when intraprocedural PVCs are infrequent.

2 Methods

All code used in this study is available at https://github.com/annien094/GNN_cardiac_fibrosis.git.

Data Generation We use synthetic data to train and test the performance of our CNN and GNN-based tissue characterisation. For this, we solve the isotropic Aliev-Panfilov monodomain model [3] in 2D geometries representative of RV patches. We solve for the transmembrane potential, $V(\mathbf{x}, t)$ and a latent field $W(\mathbf{x}, t)$ that controls the recovery of the action potential:

$$\frac{\partial V}{\partial t} = \nabla \cdot (\mathbf{D} \nabla V) - kV(V - a)(V - 1) - VW \quad (1a)$$

$$\frac{\partial W}{\partial t} = \left(\epsilon + \frac{\mu_1 W}{V + \mu_2} \right) (-W - kV(V - b - 1)). \quad (1b)$$

Here, $\mathbf{D}$ is the diffusion tensor which reduces to a scalar diffusion coefficient in isotropic and homogeneous conduction we use to model healthy tissue. k controls the steepness of depolarisation, and a is the excitation threshold potential. A no-flux Neumann condition: $\frac{\partial V}{\partial \mathbf{n}} = 0$ is enforced at the boundary. The parameter values are chosen as in [12] to model the healthy tissue.

We solve the action potential in the following 2D geometries:

- A square (dimensions: $100 \times 100 \text{ mm}^2$), using a central finite differences solver ($dx = 1 \text{ mm}$) and 4-step Runge-Kutta method ($dt = 0.01 \text{ TU} = 0.129 \text{ ms}$).
- Two curved 2D surfaces corresponding to ellipsoid patches of comparable physical size (surface area ≈ 420 and 440 mm^2) but different curvature (maximum Gaussian curvature $K = 1.19 \times 10^{-2}$ and $2.23 \times 10^{-3} \text{ mm}^{-2}$ for the higher- and lower-curvature cases, respectively). Mesh size: 19,481 nodes, 38,400 triangles, triangle size: $0.17 \pm 0.05 \text{ mm}$). An implicit-explicit time-stepping scheme is implemented in FEniCS [4].

We pace every 645 ms (50 TU in the Aliev-Panfilov model) for 3 times, each at a random different location that is outside of any pathological regions. The total duration of each recorded signal is 1935 ms, with a sampling frequency of 775 Hz.

Electrogram (EGM) Computation From the action potential field, $V(\mathbf{x}, t)$, we can compute the (unscaled) unipolar extracellular potentials, $\phi_e(\mathbf{x}_e, t)$, measured by electrodes placed on the endocardium at locations $\mathbf{x}_e$ through an unbounded Poisson equation solver [18]:

$$\phi_e(\mathbf{x}_e, t) \propto \iint \frac{\nabla \cdot (\mathbf{D} \nabla V(\mathbf{x}, t))}{|\mathbf{x} - \mathbf{x}_e|} dS. \quad (2)$$

In the square geometry, $\phi_e(\mathbf{x}_e, t)$ are calculated at 100 locations placed uniformly across the square. For the curved tissue, $\phi_e(\mathbf{x}_e, t)$ are calculated at 192 locations uniformly distributed across the surface, offset by a small distance (0.01 mm) from the tissue surface along the outward normal.

Modelling of Focal Pathology We introduce heterogeneities in model parameters to model localised pathological changes that could be related to the sources or pathways of PVCs, including:

- Fibrosis: assigned a lower diffusion coefficient D [19] sampled uniformly from $[0.02, 0.04] \text{ mm}^2 \text{ TU}^{-1}$, 60-80% lower than the baseline $D = 0.1 \text{ mm}^2 \text{ TU}^{-1}$.
- Fast depolarisation: assigned a higher k sampled uniformly from $[9.6, 16.0]$, 20-100% higher than the baseline $k = 8.0$.
- High excitability: assigned a negative excitation threshold $a = -0.025$. These regions thus have the ability to initiate action potentials, although they are not sources of action potentials in our model – the frequency of spontaneous depolarisation at the chosen a is lower than the external pacing frequency of 7.75 Hz.

For each type of abnormality, we generate 6000 simulations. Each simulation is independently generated with heterogeneities of different irregular polygonal shapes, location, dimensions, and, for fibrotic patches, a different number of heterogeneities. The shapes are generated by sampling $n_v \in [6, 12]$ vertices at random radii around a circle, applying optional elongation, and enforcing separation.

Deep Learning Model for Focal Pathology Detection To detect, from the EGM signals, the location of focal pathology introduced in the Aliev-Panfilov model, we use a deep learning model composed of two modules (Fig. 1):

1. A temporal encoder module based on a CNN to learn low-dimensional latent representations of the EGM waveforms, and
2. A spatial module based on a GNN to propagate the information in the sparsely distributed latent EGMs waveforms to a high-resolution representation of the myocardium.




Fig. 1. A schematic of our proposed framework. The temporal encoder module based on a CNN learns latent representations of the EGM waveforms. Then, the GNN based spatial module propagates the latent information from sparsely distributed electrodes to high-resolution tissue nodes. A small classification head then takes the final embeddings and performs binary classification.

CNN Module The temporal module is composed of two blocks, each with two 1D-CNN layers (kernel sizes 64 and 32, respectively), followed by max pooling, group norm, and dropouts, producing a 64-dimensional latent embedding for each EGM trace. Designed as an encoder to be coupled to the spatial module, it is initially pretrained on its own on a per-electrode binary EGM classification task: from a single electrode’s ϕ_e time series, it predicts whether that electrode overlies an abnormal region (characterised by either lower D , higher k , or negative a).

GNN Graph Construction We represent the spatial relationships between each tissue element and the electrodes using a graph. We describe two types of nodes, though the model does not explicitly distinguish between them: 1) i_ϕ , the electrodes, at which we have ϕ_e measurements, 2) i_V , the tissue nodes. Each i_ϕ node is connected to its 20 nearest i_V neighbours. Apart from the connections to i_ϕ , each i_V is connected to all the other i_V nodes within a radius of 1.0 mm.

For the 2D square, we have 3600 i_V and 100 i_ϕ nodes (i_V are $6\times$ finer than the electrodes along each axis). For the 2D curved surfaces, we have 1,728 i_V 192 i_ϕ nodes (i_V are around $3\times$ finer than the electrodes along each direction). The same graph resolution is used for both the higher- and lower-curvature geometries.

GNN Graph Attention Layers Each i_ϕ node is initialised with the 64-dimensional latent representation from the CNN Module. i_V node initialisation is set by taking the bilinear interpolation of neighbouring i_ϕ initialisations. For information exchange across the graph nodes, we adopt two graph attention (GATv2) layers [6], which allows each node to learn to pay different level of attention to its neighbours during two rounds of message passing. The attention is learnt and initialised based on edge features, where the edge features are the inverse of distances between nodes, similar to the formulation in Eq. 2.

Experimental Setup 80% of the simulations were used for training, leaving 20% for testing. The loss function is a combined cross entropy and Dice score loss, equally weighted, calculated on the i_V nodes. The CNN is trained using an Adam optimiser for 100 epochs, whereas the GNN is trained for 50, both with early stopping.

To assess the performance of the GNN module, we compare the performance of the full CNN+GNN (‘High Resolution (HR) Model’ below) approach against a CNN-only classification (‘Low Resolution (LR) Model’ below). For this comparison, after pretraining the temporal CNN module, the Low Resolution Model assigns each i_V the label of its nearest i_ϕ . We also evaluate both methods’ robustness against noisy signals and electrode sparsity.

Finally, to investigate the GNN models’ generalisation capability, we apply the model trained on 2D flat surfaces directly on simulations on curved surfaces to detect focal pathology, without any further training (‘zero-shot’). Then, we fine-tune the model over 10 additional simulations on the lower-curvature surface, each with a different pathology location (‘few-shot’), and apply it to detect focal pathology again on held-out graphs.

3 Results

2D Square The performance of focal pathology detection across a range of metrics is summarised in Table 1, with example datasets displayed in Fig. 2.

Table 1. Performance evaluated by average precision, macro F1, precision/recall (of the abnormal class at a decision threshold of 0.5), and ROC AUC for low resolution (LR) and GNN-based high resolution (HR) model.

	Average precision	Macro F1	Precision	Recall	ROC AUC
Low Resolution (LR) Model (CNN only)					
Fibrosis ($D \downarrow$), single patch	0.76	0.80	0.45	0.96	0.990
Fibrosis ($D \downarrow$), multiple patches	0.78	0.82	0.57	0.90	0.960
Rapid depolarisation ($k \uparrow$)	0.74	0.86	0.61	0.92	0.991
High excitability ($a < 0$)	0.75	0.85	0.59	0.93	0.990
High Resolution (HR) Model (GNN+CNN)					
Fibrosis ($D \downarrow$), single patch	0.96	0.93	0.86	0.89	0.998
Fibrosis ($D \downarrow$), multiple patches	0.93	0.91	0.84	0.84	0.986
Rapid depolarisation ($k \uparrow$)	0.97	0.94	0.90	0.89	0.999
High excitability ($a < 0$)	0.95	0.93	0.89	0.84	0.998

The CNN-only LR model achieves good recall, indicating that it is able to pick up the pathological EGMs. However, the precision is consistently poor due to the limited spatial resolution of electrodes. In contrast, the GNN-based HR method localises the areas with abnormal tissue properties in detail beyond the electrode resolution, achieving higher F1 scores, demonstrating a good balance

between precision and recall. It also shows better performance in threshold-independent metrics such as average precision⁶ and ROC AUC across all pathology types. When detecting multiple fibrosis patches, both models perform slightly worse in the prevalence-independent metric, ROC AUC, compared to the single-patch cases, reflecting the greater difficulty of ranking several dispersed regions.




Fig. 2. Visualisation of the low resolution model and our proposed GNN-based high resolution model for detecting single & multi-patch lower D , higher k , and negative a . The performance with sparse electrodes and noisy signals are also shown. Examples of healthy signals and signals from the pathological regions (in the locations indicated by a matching colour star) with and without noise are shown in the rightmost column.

The average precision of the HR model with respect to increasing electrode sparsity and signal noise is shown in Fig. 3, with visual examples in Fig. 2. The

⁶ Here and throughout, average precision means precision averaged over ranked predictions as the threshold varies, weighted by changes in recall.

HR model is generally robust against electrode sparsity across all three heterogeneity types, with average precision above 70% up to half of the electrodes masked. Electrode sparsity reduces the model’s ability to identify the gap between closely placed fibrotic patches in the multi-patch case. For noisy signals, the performance varies across pathologies, with fibrosis detection being particularly sensitive to noise for both LR and HR models. For high excitability and fast depolarisation detection, the LR model can be seen in Fig. 2 to be less robust to noise, typically incorrectly confusing noise with small pathological regions. Quantitatively, LR model achieves a precision of 0.44 and 0.25, compared to the HR precision of 0.91 and 0.89 for higher k and negative a , respectively, at 5% noise.





Fig. 3. Average precision of the HR model with increasing electrode sparsity and signal noise. Performance for both random and regular downsampling of electrodes are shown. Error bars are computed over 3 different rounds of random sampling.

2D Curved Surfaces Table 2 summarises the zero- (trained on the flat surfaces only) and few-shot performance of the HR model on curved surfaces for focal pathology detection. In the zero-shot setting, fibrosis (low D) and fast depolarisation (high k) detection both achieve average precisions of above 0.850, correctly identifying the approximate location although the shape is not accurate. Detection of heightened excitability (negative a) does not generalise well when directly applied to the curved surfaced in a zero-shot setting, at only 0.11 of average precision.

After fine-tuning just over 10 graphs, however, the model’s performance improves and is exceptional in all three cases, achieving average precisions of above 0.97 across all three types of pathology and the two types of curved surfaces. The model has a higher performance in the lower curvature setting than in the higher curvature one. The few-shot performance visualised on two held-out graphs is shown in Fig. 4, demonstrating the fine-tuned models’ great performance.

Table 2. Zero-shot and few-shot performance on curved surfaces.

Metric	Lower curvature ($\max K=2.23 \times 10^{-3} \text{ mm}^{-2}$)		Higher curvature ($\max K=1.19 \times 10^{-2} \text{ mm}^{-2}$)	
	Zero-shot	Few-shot	Zero-shot	Few-shot
Fibrosis ($D \downarrow$)				
Average Precision	0.857	0.986	0.632	0.958
ROC-AUC	0.966	0.998	0.944	0.994
Max-F1 score	0.822	0.937	0.704	0.880
Precision at Max-F1	0.785	0.966	0.693	0.971
Recall at Max-F1	0.862	0.911	0.715	0.805
PVC foci ($k \uparrow$)				
Average Precision	0.853	0.998	0.930	0.997
ROC-AUC	0.966	1.000	0.989	1.000
Max-F1 score	0.792	0.988	0.830	0.976
Precision at Max-F1	0.976	0.976	0.762	0.976
Recall at Max-F1	0.667	1.000	0.911	0.976
PVC foci ($a < 0$)				
Average Precision	0.107	0.977	0.126	0.988
ROC-AUC	0.508	0.996	0.598	0.998
Max-F1 score	0.274	0.918	0.269	0.963
Precision at Max-F1	0.159	0.973	0.156	0.975
Recall at Max-F1	0.984	0.870	1.000	0.951

**Fig. 4.** Visualisation of focal pathology detection results after few-shot (10 graphs) fine-tuning on curved surfaces.

4 Discussion & Conclusions

We present a CNN and GNN-based method to characterise cardiac tissue properties from electrode measurements, specifically detecting regions of fibrosis, abnormal repolarisation, or heightened excitability, relevant for PVC ablation. The proposed method shows a good performance on simulated data, and demonstrates the potential to adapt to irregular distributions of sparse electrodes, and from flat to curved geometries with minimal fine-tuning, which is crucial for clinical settings and personalisation.

For flat surfaces, the LR 1D-CNN-only method does not see beyond the signals of one electrode at a time, whilst the HR CNN+GNN approach learns to process signals from all neighbouring electrodes, showing robustness against electrode sparsity. The GNN also naturally incorporates irregular layouts through the graph structure, which makes it more suitable for cardiac geometries compared to a 2D-CNN, which could use neighbouring electrode information but is better suited to regular grids. The HR method is more robust to noise compared to the LR one, as the latter often falsely picks up noisy signals as abnormalities. However, fibrosis detection is brittle against noise for both methods. It may be because the D heterogeneities modelling fibrosis affect local conduction rather than single cell properties, making them harder to identify.

We also show that the model trained on flat surfaces can be used on unseen curved surface directly or, for an improved performance, after minimal fine-tuning to learn artefacts from curvature. Generalising parameters learnt in planar tissue to curved geometries is non-trivial, because wave propagation speed and properties depend on tissue curvature [10]. In the zero-shot setting, the detection of heightened excitability ($a < 0$) in curved surfaces is more difficult than the identification of the other two abnormal properties. This may be related to a dependency of excitability on the tissue curvature [10,9].

Our GNN approach leverages the attention-based message passing layers, allowing each node to place learnt levels of importance on its neighbours as a function of the inverse of the distance between them, without explicitly fixing the attention weights. We hypothesise that this helps with the generalisability observed. These data suggest that the model learns features that transfer across geometries, relevant when we want to adapt to patient-specific geometries using limited data. Performance may be further increased if a variety of tissue curvatures are included in the training data.

Other techniques, beyond the GNN setup we used, could be employed to achieve higher-resolution classification performance, although they may not generalise as well to unseen geometries. These include implicit neural representations, which could parametrise the electrogram-derived features as a continuous function enabling data query at arbitrarily high spatial resolutions [20], classical spatial reconstruction methods such as radial basis function (RBF) interpolation, or physics-informed machine learning methods [8].

We note that, so far, both the planar and curved datasets are generated using the same isotropic Aliev–Panfilov model and a similar pathology-generation procedure. The current results therefore demonstrate transfer across synthetic

geometries within one simulator. Further work will address how well the approach generalises to other electrophysiological models, electrophysiological heterogeneities, anisotropic conduction and geometries with non-zero thickness (as in the left ventricle). Ultimately, we expect to show that our model performs well in realistic patient anatomies, and in the presence of clinical EGM noise and artefacts. In the long run, it could play a role in the clinic to help identify PVC ablation targets using sinus rhythm data from PVC patients, improving the success rates and decreasing the duration of PVC ablation procedures.

Acknowledgments. We acknowledge computational resources and support provided by the Imperial College Research Computing Service (<http://doi.org/10.14469/hpc/2232>). We thank St George’s Hospital Charity for research funding.

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

References

1. Acosta, J., Soto-Iglesias, D., Fernández-Armenta, J., Frutos-López, M., Jáuregui, B., Arana-Rueda, E., Fernández, M., Penela, D., Alcaine, A., Cano, L., Pedrote, A., Berruezo, A.: Clinical validation of automatic local activation time annotation during focal premature ventricular complex ablation procedures. *EP Europace* **20**(F12), f171–f178 (2018). <https://doi.org/10.1093/europace/eux306>
2. Alcaine, A., Soto-Iglesias, D., Acosta, J., Korshunov, V., Penela, D., Martínez, M., Linhart, M., Andreu, D., Fernández-Armenta, J., Laguna, P., Martínez, J.P., Camara, O., Berruezo, A.: Automatic activation mapping and origin identification of idiopathic outflow tract ventricular arrhythmias. *Journal of Electrocardiology* **51**(2), 239–246 (2018). <https://doi.org/10.1016/j.jelectrocard.2017.10.015>
3. Aliiev, R.R., Panfilov, A.V.: A simple two-variable model of cardiac excitation. *Chaos, Solitons & Fractals* **7**(3), 293–301 (1996)
4. Baratta, I.A., Dean, J.P., Dokken, J.S., Habera, M., Hale, J.S., Richardson, C.N., Rognes, M.E., Scroggs, M.W., Sime, N., Wells, G.N.: Dolfinx: The next generation fenics problem solving environment (Dec 2023). <https://doi.org/10.5281/zenodo.10447666>
5. Baser, K., Bas, H.D., Yokokawa, M., Latchamsetty, R., Morady, F., Bogun, F.: Infrequent intraprocedural premature ventricular complexes: implications for ablation outcome. *Journal of Cardiovascular Electrophysiology* **25**(10), 1088–1092 (2014)
6. Brody, S., Alon, U., Yahav, E.: How attentive are graph attention networks? arXiv preprint arXiv:2105.14491 (2021)
7. Calvo, N., Jongbloed, M., Zeppenfeld, K.: Radiofrequency catheter ablation of idiopathic right ventricular outflow tract arrhythmias. *Indian pacing and electrophysiology journal* **13**(1), 14–33 (2013)
8. Chiu, C.E., Roy, A., Cechnicka, S., Gupta, A., Pinto, A.L., Galazis, C., Christensen, K., Mandic, D., Varela, M.: Physics-informed neural networks can accurately model cardiac electrophysiology in 3d geometries and fibrillatory conditions. In: *International Workshop on Statistical Atlases and Computational Models of the Heart*. pp. 98–109. Springer (2024)

9. Connolly, A., Kelly, A., Campos, F.O., Myles, R., Smith, G., Bishop, M.J.: Ventricular endocardial tissue geometry affects stimulus threshold and effective refractory period. *Biophysical Journal* **115**(12), 2486–2498 (2018)
10. Dierckx, H., Bernus, O., Verschelde, H.: Accurate eikonal-curvature relation for wave fronts in locally anisotropic reaction-diffusion systems. *Physical Review Letters* **107**(10), 108101 (2011)
11. Enriquez, A., Muser, D., Markman, T.M., Garcia, F.: Mapping and ablation of premature ventricular complexes: state of the art. *Clinical Electrophysiology* **10**(6), 1206–1222 (2024)
12. Göktepe, S., Kuhl, E.: Electromechanics of the heart: a unified approach to the strongly coupled excitation–contraction problem. *Computational mechanics* **45**, 227–243 (2010)
13. Jáuregui, B., Fernández-Armenta, J., Acosta, J., Penela, D., Terés, C., Ordóñez, A., Soto-Iglesias, D., Silva, E., Chauca, A., Carreño, J.M., Scherer, C., Pedrote, A., Berrueto, A.: MANual vs. automatIC local activation time annotation for guiding Premature Ventricular Complex ablation procedures (MANIaC-PVC study). *EP Europace* **23**(8), 1285–1294 (2021). <https://doi.org/10.1093/europace/euab080>
14. Jenkins, A., Cini, A., Barker, J., Mandic, D., Alippi, C., Ng, F.S.: Mapping the dynamics of atrial fibrillation with spatiotemporal graph neural networks. In: *NeurIPS 2025 Workshop on Learning from Time Series for Health* (2025)
15. Li, D., Zhang, J., Zhang, Q., Wei, X.: Classification of ecg signals based on 1d convolution neural network. In: *2017 IEEE 19th international conference on e-health networking, applications and services (Healthcom)*. pp. 1–6. IEEE (2017)
16. 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)
17. Ng, G.A.: Treating patients with ventricular ectopic beats. *Heart* **92**(11), 1707–1712 (2006)
18. Plonsey, R., Barr, R.C., Bioelectricity, A.: Quantitative approach. Cham, Switzerland: Springer (2007)
19. Roy, A., Varela, M., Aslanidi, O.: Image-based computational evaluation of the effects of atrial wall thickness and fibrosis on re-entrant drivers for atrial fibrillation. *Frontiers in physiology* **9**, 1352 (2018)
20. Sitzmann, V., Martel, J., Bergman, A., Lindell, D., Wetzstein, G.: Implicit neural representations with periodic activation functions. *Advances in neural information processing systems* **33**, 7462–7473 (2020)
21. Wang, C., Pei, Z., Cai, X.: Advancements in artificial intelligence for the localization of premature ventricular contraction origins. *Frontiers in Cardiovascular Medicine* **13**, 1864886 (2026). <https://doi.org/10.3389/fcvm.2026.1864886>