Reducing the number of leads for ECG Imaging with Graph Neural Networks and meaningful latent space

Giacomo Verardo[1 [0000−0001−7367−9200]], Daniel F. Perez-Ramirez[1,2 [0000−0002−1322−4367]], Samuel Bruchfeld[3 [0000−0002−0284−5176]], Magnus Boman[3,4 [0000−0001−7949−1815]], Marco Chiesa[1 [0000−0002−9675−9729]], Sabine Koch[3 [0000−0001−7144−8740]], Gerald Q. Maguire Jr.[1 [0000−0002−6066−746X]], and Dejan Kostic[1 [0000−0002−1256−1070]]

[1] KTH Royal Institute of Technology, Stockholm, Sweden
{verardo, mchiesa, maguire, dmk}@kth.se
[2] RISE Computer Science, Stockholm, Sweden
[3] Karolinska Institutet, Stockholm, Sweden
{magnus.boman,samuel.bruchfeld,sabine.koch}@ki.se
[4] MedTechLabs, Stockholm, Sweden

Abstract. ECG Imaging (ECGI) is a technique for cardiac electrophysiology that allows reconstructing the electrical propagation through different parts of the heart using electrodes on the body surface. Although ECGI is non-invasive, it has not become clinically routine due to the large number of leads required to produce a fine-grained estimate of the cardiac activation map. Using fewer leads could make ECGI practical for clinical patient care. We propose to tackle the lead reduction problem by enhancing Neural Network (NN) models with Graph Neural Network (GNN)-enhanced gating. Our approach encodes the leads into a meaningful representation and then gates the latent space with a GNN. In our evaluation with a state-of-the-art dataset, we show that keeping only the most important leads does **not** increase the cardiac reconstruction and onset detection error. Despite dropping almost 140 leads out of 260, our model achieves the same performance as another NN baseline while reducing the number of leads. Our code is available at `github.com/giacomoverardo/ecg-imaging`.

Keywords: Deep Learning · ECG Imaging · Graph Neural Networks

1 Introduction

ECG Imaging (ECGI) is based on solving the inverse problem of cardiac electrophysiology to reconstruct the electrical propagation through different parts of the heart by using electrodes on the body surface together with cardiac and torso geometry. It offers a non-invasive method to map and analyze cardiac electrical activity with high precision, aiding in the diagnosis and treatment of arrhythmias such as atrial fibrillation and ventricular tachycardia. ECGI has typically required hundreds of leads. This paper presents an approach to reduce the number of leads needed to reconstruct the cardiac activation map from the Body Surface Potential Mapping (BSPM) measurements.

Our main contributions are:

1. We present a novel approach for redundant lead removal in ECGI, which relies on a Graph Neural Network (GNN)-gating mechanism in the Frequency Modulated Möbius (FMM) latent space. To the best of our knowledge, we are the first to propose using this combination when solving the inverse problem of ECGI.
2. We evaluate our model on a state-of-the-art dataset[5] and show that we can reduce the number of leads required for ECGI by up to a hundred out of 260, thus potentially making ECGI feasible for clinical use. Moreover, we show that our model outperforms a state-of-the-art model in terms of cardiac reconstruction loss and onset detection even when a large number of leads are removed.

2 Background

This section provides information about Artificial Intelligence (AI) and non-AI-based ECGI (Section 2.1), and the sub-networks composing our model: convolutional layers (Section 2.2), FMM encoding (Section 2.3), and the GNN architecture (Section 2.4).

2.1 ECG Imaging

ECGI is a non-invasive cardiophysiology technique that combines information from BSPM and geometry information (usually obtained with a CT scan) to reconstruct the electrical propagation inside the heart. ECGI is a broad field and includes different tasks, such as reconstruction of the cardiac activation map [2,3], detection of the earliest activation sites [13] during a heartbeat, and estimation of the cardiac Transmembrane Potential (TMP) [5]. The forward problem of ECGI refers to the computation of the BSPM (*i.e.*, the potentials measured with electrodes) from a known electrical cardiac source. This relationship can be expressed with a linear model:

$$\mathbf{x}_{bspm}(t) = \mathbf{H}\mathbf{x}_{acti}(t) + \mathbf{n}(t) \tag{1}$$

where the matrix $\mathbf{H}$ can be estimated from geometric information (e.g., torso and heart meshes) and the electrical conductivity of tissues within the torso, $\mathbf{x}_{acti}(t)$ is the source cardiac electrical activity, $\mathbf{x}_{bspm}(t)$ is the signal measured from the electrodes, and $\mathbf{n}(t)$ represents random additive noise.

Although the linearity of the equation may suggest that reconstructing the cardiac source is a straightforward task, the ECGI inverse problem is ill-posed [14]: even though the relation between the cardiac electrical source and the corresponding BSPM measurement is linear, a negligible error in the measurements may lead to considerable errors in the estimated cardiac source:

$$\hat{\mathbf{x}}_{acti}(t) = \mathbf{H}^{-1}\mathbf{x}_{bspm}(t) \tag{2}$$

Tikhonov regularization has been used to include *a priori* knowledge of x_{acti} and reduce the impact of the ill-posed problem. AI has also been proposed to learn the relation between x_{acti} and x_{bspm} or a suitable regularization [1,12].

[5] The number of employed leads in the evaluated dataset is $N = 260$. The cardiac and torso geometry is provided by Computed Tomography (CT) scans.

ECGI with Tikhonov regularization Solving the inverse problem for ECGI requires two steps: (i) defining the forward H matrix in Equation (1) from geometrical and conductivity properties of the torso (for instance, by means of the Boundary Element Method (BEM)) and (ii) estimating the cardiac electrical source while stabilizing the solution. Tikhonov regularization is a standard technique to penalize predicted $\hat{x}_{acti}$ with high norms. The general formulation is:

$$J(\mathbf{x}) = \|\mathbf{H}\mathbf{x}_{acti} - \mathbf{x}_{bspm}\|^2 + \lambda^2 \|\mathbf{R}\mathbf{x}_{acti}\|^2 \tag{3}$$

where λ is the regularization parameter and $\mathbf{R}$ the regularization matrix, which is the identity matrix for zero-order Tikhonov regularization. Tikhonov regularization thus involves solving an optimization problem that balances the fit of the forward model to the measured data with a regularization term to stabilize the solution, minimizing $J(x)$.

ECGI via AI-based models Neural Network (NN)-based solutions have been applied to ECGI. They employ Long-Short Term Memory (LSTM) layers, fully connected models, convolutional networks [4,2,5], or GNNs [8] to include the geometry input (e.g.,obtained from the CT scan or Magnetic Resonance Imaging (MRI) [21]) and the BSPM measurements to estimate the cardiac electrical source. This allows learning from available data instead of solving an optimization problem for each heartbeat. Once the NN is trained, inference can be performed to retrieve an estimate in much less time than the optimization time of standard solvers for inverse problems.

Although ECGI is a viable alternative to invasive approaches, it is clinically impractical since it requires hundreds of electrodes to be placed on the patient's body by means of either single electrodes or vests [11]. Multiple non-AI-based methods have been proposed to detect the earliest activation site from a reduced number of electrodes [13] or even from standard 12-lead Electrocardiograms (ECGs) [11] with increasing error as the number of leads is reduced. Instead, we argue that AI-based techniques can be beneficial for ECGI when a limited set of leads is available.

2.2 Convolutional neural networks

Convolutional Neural Networks (CNNs) [10] are state-of-the art NNs that detect increasingly complex patterns in the input as their depth increases. Each layer in a CNN performs a convolution operation between the inputs and a set of trainable parameters (*i.e.,* the kernel), thus producing an output which is higher where the input is closer to the kernel's representation. CNNs have been extensively used for 1D signals, such as ECGs [7,2], since they can recognize specific patterns in the ECG input, such as the shapes of the PQRST waves, and thus CNNs are employed to solve numerous tasks. However, CNNs have seldom been employed to encode ECGs into explainable, lead-specific latent space [20].

2.3 FMM Formulation

The FMM model [15] is a time-series oriented mathematical formulation that can encode simple waves as a set of interpretable parameters. In particular, they can be

used to model the depolarization and repolarization waves in each ECG lead [16], thus compressing them into a lower-dimensional representation that retains the original information. These parameters can be extracted with the following formulas:

$$W(t_i) = \mu(t_i) + e(t_i) = M + A\cos(\varphi(t_i)) + e(t_i) \tag{4}$$

$$\varphi(t) = \beta + 2\arctan(\omega\tan(t-\alpha)) \tag{5}$$

$$\mathbf{e} = (e(t_1), \ldots, e(t_n))^T \sim \mathcal{N}(\mathbf{0}, \sigma^2 \boldsymbol{I}) \tag{6}$$

where each parameter in the tuple $\theta = (A, \alpha, \beta, \omega, M)$ can be either shared or not between leads and the parameters are: $A \in \mathbb{R}^+$ is the amplitude of the wave; $\alpha \in [0, 2\pi)$ is the angular position of the peak (lead-shared); $\beta \in [0, 2\pi)$ is the peak direction; $\omega \in [0, 1]$ indicates the width of the lobe of the peak (lead-shared); $M \in \mathbb{R}$ is the constant offset of the lead; and $\mathbf{e} \in \mathbb{R}$ is the error between the real heartbeat and one reconstructed from FMM parameters.

Different NN-based models have been proposed to automatically extract the FMM parameters from ECG data [20,23]. Using a model trained with these FMM parameters results in an inference time that is three orders of magnitude less than the time to solve the optimization problem [16]. However, these studies are limited to 12-leads ECG and have never been applied for ECGI.

2.4 Graph Neural networks

GNNs have become a versatile tool for addressing various inference tasks on graph-structured data [17,22]. Conceptually, a GNN is composed of message-passing layers that recursively update each node's feature vector by considering information from its neighboring nodes [6]. Stacking K GNN layers involves creating node embedding vectors that consider their K-hop neighborhood by utilizing the graph's structure and the relationships between nodes [6,9]. Mathematically, for a graph $G = \langle V, E \rangle$ with nodes $v \in V$ and edges $(v, u) \in E$, at GNN layer i, the hidden feature vector h_v for node v is updated as:

$$h_v^i = g_1\left(\text{AGG}_{u \in \mathcal{N}(v)}\left[g_2\left(h_u^{i-1}\right)\right], h_v^{i-1}\right), \tag{7}$$

where g_1 and g_2 are non-linear transformations, AGG is an order-invariant aggregation function, and $\mathcal{N}(v)$ represents the set of neighboring node feature vectors of node v [6,9,18]. These embedding vectors are generally processed further with linear layers to produce the final output based on the specific task. For instance, node regression can be achieved by feeding each h_v $\forall v \in V$, through a final regression layer. Although GNNs have already been employed for ECGI to map torso to heart geometries [8], we claim that they are also beneficial to detect the most crucial leads to reconstruct the cardiac activation map.

3 Methodology

Our aim is to reconstruct the cardiac activation map from the BSPM measurements and the CT scans. We do so by integrating two novel components: (i) an FMM-based

encoded latent space that supports hundreds of leads and (ii) a gating Graph Attention (GAT)-based GNN that selects and weights the latent space, as depicted in Figure 1. We argue that such a solution can determine the importance of the different leads in the latent space, while still allowing the retrieval of information from all the leads into the latent space (*e.g.,* to compute shared FMM parameters). Section 3.1 describes our pipeline, including inputs, 1D-convolutional encoder, FMM latent space, 3D-convolutional decoder, and employed loss functions. We describe our GNN gating mechanism in Section 3.2.

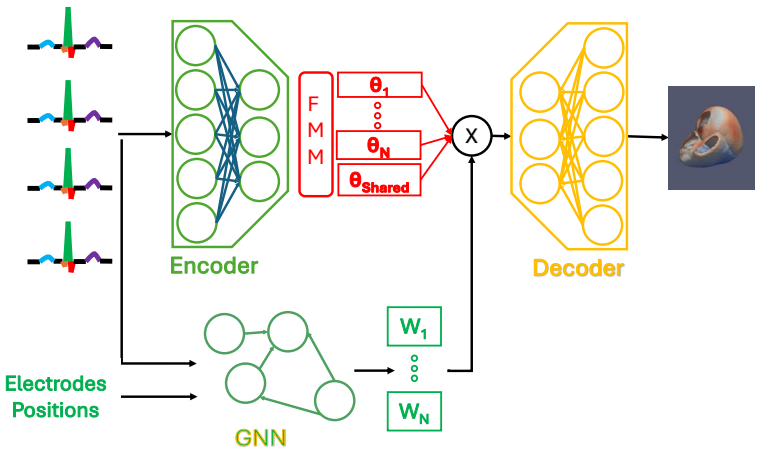

Fig. 1. Model architecture, which includes a 1D convolutional encoder, an FMM latent space for hundreds of leads, a model gating GNN and a 3D convolutional decoder. The FMM latent space includes lead-exclusive and lead-shared θ parameters.

3.1 Encoder-Decoder pipeline

Our input consists of 3D tensors (the cardiac activation map and conduction velocity map obtained from the CT scan) and the measured BSPM data. Our encoder extracts the estimated FMM parameters from the BSPM sensors by means of stacked convolutional blocks with different kernel sizes and number of kernels per layer, in a similar fashion as [7]. The extracted parameters are $\theta = f_{\phi_{enc}}(x_{bspm})$, which includes both lead-shared and not shared parameters:

$$\begin{aligned} &\alpha_j, \omega_j \in \theta_{shared}, \\ &A_{ij}, \beta_{ij} \in \theta_i, \quad i \in \{1, \ldots, N\}, \quad j \in \{1, \ldots, M\}, \end{aligned} \tag{8}$$

where N is the number of leads, M is the number of ECG waves in the BSPM inputs, θ_{shared} represents the parameters shared across all leads, and θ_i represents the parameters specific to the i-th lead. Inspired by [20], we designed our FMM layer to

handle parameters for hundreds of leads. The layer has two outputs: the θ parameters and the reconstructed BSPM signal $\hat{x}_{bspm}$ from θ. We can hence define a loss function in the latent space to constrain θ to be close to the real FMM parameters by minimizing the mean squared error between original and reconstructed BSPM signals of length L:

$$\text{Loss}_{bspm} = \frac{1}{NL} \sum_{i=1}^{N} \sum_{j=1}^{L} (x_{bspm,ij} - \hat{x}_{bspm,ij})^2 \tag{9}$$

Additionally, we constrain the FMM latent space to a meaningful representation by encouraging through a loss term Loss_α the α parameters to be close to the detected peaks in the BSPM leads. By doing so, we avoid running a warm-up regression phase as in [20], which is expensive and requires the ground-truth FMM parameters which are usually not available.

After weighting the extracted parameters with the GNN (Section 3.2), θ is given as input to the decoder which produces an estimation of the cardiac activation map. The decoder is a 3D CNN that is conditioned to the BSPM signals and the conduction system velocity (obtained by processing the CT scan). The decoder is taken from [2], with the notable exception of having L1-L2 regularizers for the decoder layers. These regularizers are crucial to avoid θ parameters with low values (*i.e.,* those dropped by the GNN) affecting the activation map reconstruction. Hence, the decoder applies the function: $\hat{x}_{acti} = h_{\phi_{dec}}(\theta, x_{bspm}, x_{conduct})$ to predict the activation map x_{acti}. The loss value for the cardiac activation map reconstruction is computed similarly to [2] by averaging the squared error on the relevant voxels (*i.e.,* on a specific mask) between the reconstruction and the *inverted* original signal, so that the earliest activation points are better identified:

$$\text{Loss}_{acti} = \frac{1}{|\mathcal{M}|} \sum_{(i,j,k) \in \mathcal{M}} \left((1 - x_{acti,ijk}) - \hat{x}_{acti,ijk}\right)^2 \tag{10}$$

where $\mathcal{M}$ is the set of indices defined by the 3D mask x_{mask}. We refer the reader to [2] for further implementation details. The full loss function is therefore:

$$\text{Loss} = \gamma_{bspm} \cdot \text{Loss}_{bspm} + \gamma_\alpha \cdot \text{Loss}_\alpha + \text{Loss}_{acti} \tag{11}$$

where γ_{bspm} and γ_α are loss weight ratios that determine the relative importance of the BSPM loss and α loss compared to the activation loss.

3.2 Graph Neural Network-Gating

We use a GNN to gate the θ_i FMM parameters from the latent space. The GNN includes multiple blocks of GAT attention layers with skip connections and layer normalization.

Our GNN produces as output a set of N w_i scalar weighs, one for each lead. The θ parameters can then be weighted as:

$$\theta'_{ij} = \begin{cases} \theta_{ij}, & \text{if } \theta_{ij} \in \theta_{shared} \\ w_i \cdot \theta_{ij}, & \text{otherwise} \end{cases} \tag{12}$$

where $\mathbf{w} = g_{\phi_{gnn}}(\text{Adj}, x_{bspm})$ and Adj is the adjacency matrix for the electrodes, computed as radial basis function of their positions: $k(x, y) = \exp\left(-\frac{\|x-y\|^2}{2\ell^2}\right)$, ℓ being the length scale parameter. The choice of a suitable ℓ allows a sufficient number of neighboring electrodes to attend to the signals of each electrode. Hence, a node with high importance can reduce the weight of its neighboring nodes, thus removing the redundancy of the available leads and therefore reducing the number of necessary leads.

During testing, the number of leads needed can be reduced by dropping those corresponding to the lowest w_i. The corresponding θ_i parameters are therefore zeroed out, as well as excluding the input leads x_{bspm} from the input to the encoder and the decoder. This approach enables training using *all* the leads, thereby enhancing the robustness of learning to reconstruct the FMM parameters since they are encoded from every lead. Moreover, our model can be dynamically evaluated on *any* number of leads, thus providing a more complete and compact analysis of the impact of lead removal for ECGI compared to full training sessions with a reduced set of leads (which would require $\sum_{k=0}^{N} \binom{N}{k} = 2^N$ different models, one for each possible set of leads).

3.3 Metrics

We evaluate the performance of our method based on two metrics: reconstruction error on the activation map (Mean Absolute Error (MAE)) and onset prediction error. The computation of the MAE follows the same formula of Equation (10), where the absolute value is employed instead of the square function. Hence, the MAE provides a measure of the normalized average temporal distance between ground-truth and predicted cardiac activation map. The onset prediction error is instead trickier to compute, since the identification of an onset depends on the employed policy. One simple policy is the selection of the earliest activation point whose neighbours are activated within a specific time interval. Although such solution is straightforward to implement, it lacks generalizability on multiple onsets (as in the dataset employed in Section 4) and robustness with respect to sparsity of the predicted activation map, which is AI-generated.

For these reasons, we propose to employ clustering algorithms to estimate the onsets on both predicted and ground-truth activation maps. We perform such computation in three steps: (i) we extract the 10% percentile of the activation map, which represents the neighborhood of the activated points at the beginning of the onset, (ii) we weight the extracted points according to their activation time and (iii) we perform K-Means clustering, where K is the number of onsets, which is available during testing. In order to assign the predicted onset to the best estimated onset from the initial points of the ground-truth activation map, we solve a Linear assignment problem (LAP) where the cost matrix is the euclidean distance. The impact of few activated outlier points is thus attenuated by the large number of points used as input to the clustering model.

4 Experimental Results

We evaluate our model on the available-on-request dataset from [2,3], which consists of simulated data of source cardiac activation map (generated with the Eikonal model), correspondent measurement on the torso and conduction map for 7 patients (5 for

training and 2 for testing). We train our model with multiple learning rates, early stopping, and learning rate decay for a maximum number of 240 epochs. We then compute the average and standard deviation of the test loss for the activation map reconstruction across the test samples. We evaluate the impact of dropping leads by zeroing out the input leads and the extracted non-shared parameters θ_i in a range between 0 and $N = 260$ with a step size of 20, as explained in Section 3.2. We compare our model performance with the Variational Autoencoders (VAE) from [2]. Details of the training procedure are available in the code repository [19].

Compared to the VAE baseline, our approach achieves lower training, validation, and test loss, thus achieving a more reliable reconstruction of the activation map. Table 1 shows the test and validation loss for different learning rates: (i) for our model with multiple ratios between cardiac activation loss and signal loss and (ii) for our baseline. The results show that the test loss is reduced by up to 6.9% in terms of mean squared error on the prediction of the cardiac activation map. The accuracy of our approach can be visually estimated by considering one example of cardiac activation map ground-truth Figure 5a, prediction Figure 5b, and the absolute difference between them.

Table 1. Comparison of our model with the VAE baseline for multiple hyper-parameters: learning rate, γ_{bspm} and γ_α. The last two are not applicable for the VAE baseline.

Model	**LR**	$\gamma_{\mathbf{bspm}}$	γ_α	**MAE [ms]**	**MSE [$\times 10^{-3}$]**	**Onset Error [mm]**
Ours	1e-5	0.1	1.0	131.7 ± 14.7	27.8 ± 5.6	29.99 ± 11.17
	1e-5	1.0	1.0	128.8 ± 16.3	26.9 ± 6.3	29.76 ± 11.63
	1e-5	1.0	10.0	131.8 ± 14.7	27.8 ± 5.6	29.81 ± 11.11
	1e-4	0.1	1.0	130.2 ± 19.1	27.5 ± 7.4	30.19 ± 12.74
	1e-4	1.0	1.0	**127.1 ± 16.1**	**26.1 ± 6.3**	**27.47 ± 10.66**
	1e-4	1.0	10.0	128.9 ± 15.9	26.9 ± 6.2	28.96 ± 11.94
VAE	1e-4	n/a	n/a	**133.4 ± 16.3**	**28.0 ± 6.7**	**30.36 ± 9.14**
	1e-5	n/a	n/a	145.0 ± 9.7	32.2 ± 4.1	43.88 ± 10.35

Our model achieves the same test reconstruction loss for the cardiac activation map as our baseline when almost 140 leads are dropped. Figure 2b shows the test activation loss of our model as we progressively remove the leads with the lowest weight w_i from the inputs. Although the impact on the loss is dramatic when a large number of leads are ignored, our model remains robust to lead removal up to almost 50 leads. Moreover, our model achieves a lower test loss than the baseline, even when 134 leads are removed. Hence, we argue that NN-based learning of important leads is a feasible technique to reduce the burden of placing hundreds of electrodes for ECGI.

Our model achieves lower prediction error on the estimated onsets compared to the baseline. Table 1 provides the average and standard deviation on the error on the position of the onsets estimated by clustering on the firstly activated points, as defined in Section 3.3. Our results prove that our model is able to better detect the predicted position of 9.5% compared to our baseline. However, the high standard deviation indicates high variability between different test samples. Indeed, as shown in Figure 3, test samples with a single onset demonstrate much lower prediction error ($17.9\,\mathrm{mm}$)

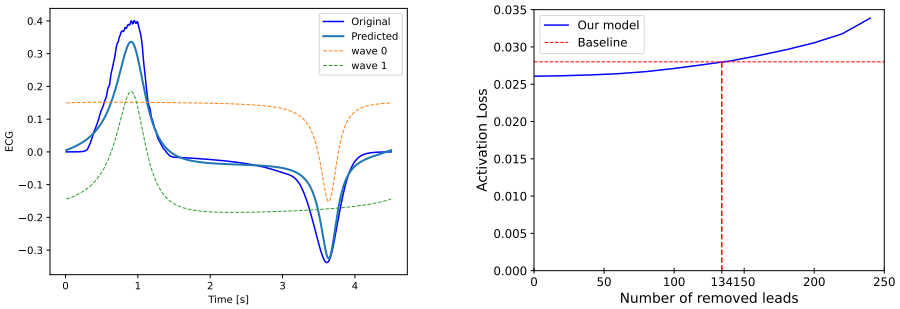

(a) Latent FMM reconstruction (b) Test loss vs. number of removed leads

Fig. 2. Results for our model with learning rate $1e{-}4$ and loss weight ratio 1.0: (a) shows the FMM reconstruction in the latent space for lead 100, proving that our training procedure can extract meaningful FMM parameters (b) shows the increase in test activation loss when leads are removed, demonstrating that (i) up to 50 leads can be removed with negligible performance degradation (ii) up to 134 leads can be removed while having lower loss than the baseline.

compared to two ($30.9\,\mathrm{mm}$) and three onsets ($29.9\,\mathrm{mm}$). Hence, finding multiple onsets is the most challenging task. Figure 3a shows that we can remove up to 150 leads without affecting the estimated onset prediction error for the one onset case. Interestingly, the performance of the model increases when few leads are removed. Such leads have almost zero weight and do not contribute to the activation map reconstruction. The specific sample depicted in Figure 4 graphically shows that removing 150 leads reduces the onset prediction by only $7.5\,\mathrm{mm}$. Instead, spatial resolution is significantly affected by lead removal in the two- and three-onset cases, which are more challenging due to the superimposition of propagation from multiple pacing sites.

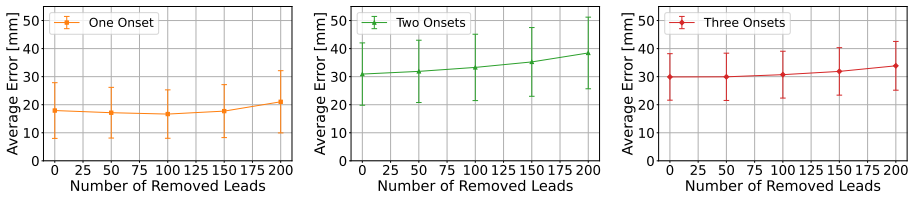

(a) One onset (b) Two onsets (c) Three onsets

Fig. 3. Prediction error for the pacing site when varying the number of removed leads for our model with learning rate $1e{-}4$, $\gamma_{\text{bspm}} = \gamma_{\alpha} = 1.0$. Figures (a), (b), (c) depicts the average error for test samples with one, two and three onsets respectively.

Our model achieves a meaningful, explainable latent space by means of FMM coefficients estimation. Figure 2a shows the reconstruction obtained from the esti-

Fig. 4. Results for our model with learning rate $1e{-}4$ and $\gamma_{\text{bspm}} = \gamma_{\alpha} = 1.0$: (a) shows the full heart with the predicted onsets when multiple leads are removed by means of our GNN-gating mechanism while (b) zooms-in on the predicted onsets. The black numbers above the points indicate the number of zeroed-out leads. For the sample shown, if up to 150 leads are removed, the distance to the predicted pacing site with all the leads is within $7.5\,\text{mm}$ (green sphere)

mated FMM parameters in the latent space, which closely resembles the real BSPM signals. Hence, we argue that our latent space can indeed be used as a compressed version of the inputs, thus allowing θ_i to be correctly gated. Compared to [20], our model does not need expensive warm-up regression to achieve a meaningful latent space. That is due to the fact that: (i) our dataset only includes the QRS complex and T wave, which can be both represented with a single FMM wave from 8 and (ii) having more leads as input is beneficial to the accuracy of the estimation, as pointed out by [16].

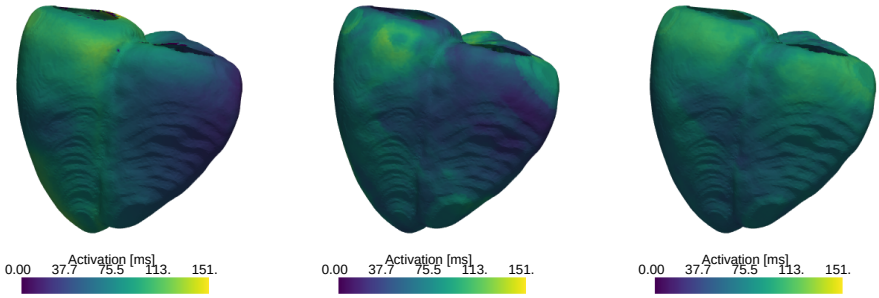

(a) Ground-truth (b) Predicted (Ours) (c) Predicted (VAE)

Fig. 5. Example cardiac activation map: (a) ground-truth, (b) prediction with our approach, (c) prediction with VAE baseline. The reconstruction capability depends on the sample, patient, and training hyper-parameters. Table 1 provides a quantitative analysis.

5 Conclusions

In conclusion, this study has investigated the possible reduction of the number required leads when performing ECGI tasks. Our GNN-based approach to lead-gating on FMM-enabled latent space allows reducing the number of employed leads by 50 with only a slight increase in the error for cardiac activation map reconstruction. Our findings also indicate that it is possible to encode ECG data from BSPM to a FMM representation with a convolutional model without requiring computing the ground-truth FMM coefficients. As future work we will investigate how the different components of our model affect the reconstruction capabilities. Moreover, we will train our models for a finer grid of hyper-parameters to show the robustness of our approach.

Acknowledgements This work was supported by the Swedish Research Council (project "Scalable Federated Learning", no. 2021-04610) and by the King Abdullah University of Science and Technology (KAUST) Office of Research Administration (ORA) under Award No. ORA-CRG2021-4699. Daniel F. Perez-Ramirez is supported by the Swedish Foundation for Strategic Research (SSF). We acknowledge Euro HPC Joint Undertaking for awarding us access to the Leonardo supercomputer located at the CINECA data center in Bologna, Italy.

References

1. Adler, J., Öktem, O.: Solving ill-posed inverse problems using iterative deep neural networks. Inverse Problems **33**(12), 124007 (Nov 2017). https://doi.org/10.1088/1361-6420/aa9581
2. Bacoyannis, T., Krebs, J., Cedilnik, N., Cochet, H., Sermesant, M.: Deep learning formulation of ecgi for data-driven integration of spatiotemporal correlations and imaging information. In: Coudière, Y., Ozenne, V., Vigmond, E., Zemzemi, N. (eds.) Functional Imaging and Modeling of the Heart. pp. 20–28. Springer International Publishing, Cham (2019). https://doi.org/10.1007/978-3-030-21949-9_3
3. Bacoyannis, T., Ly, B., Cedilnik, N., Cochet, H., Sermesant, M.: Deep learning formulation of electrocardiographic imaging integrating image and signal information with data-driven regularization. EP Europace **23**(Supplement_1), i55–i62 (03 2021). https://doi.org/10.1093/europace/euaa391
4. Chen, K.W., Bear, L., Lin, C.W.: Solving inverse electrocardiographic mapping using machine learning and deep learning frameworks. Sensors **22**(6) (2022). https://doi.org/10.3390/s22062331
5. Ghimire, S., Dhamala, J., Gyawali, P.K., Sapp, J.L., Horacek, M., Wang, L.: Generative modeling and inverse imaging of cardiac transmembrane potential. In: Frangi, A.F., Schnabel, J.A., Davatzikos, C., Alberola-López, C., Fichtinger, G. (eds.) Medical Image Computing and Computer Assisted Intervention – MICCAI 2018. pp. 508–516. Springer International Publishing, Cham (2018). https://doi.org/10.1007/978-3-030-00934-2_57
6. Gilmer, J., Schoenholz, S.S., Riley, P.F., Vinyals, O., Dahl, G.E.: Neural Message Passing for Quantum Chemistry. Proc. 34th Int. Conf. Mach. Learn. (ICML) **3**, 2053–2070 (apr 2017)
7. Jang, J.H., Kim, T.Y., Lim, H.S., Yoon, D.: Unsupervised feature learning for electrocardiogram data using the convolutional variational autoencoder. PLOS ONE **16**(12), 1–16 (12 2021). https://doi.org/10.1371/journal.pone.0260612
8. Jiang, X., Missel, R., Toloubidokhti, M., Gillette, K., Prassl, A., Plank, G., et al.: Hybrid neural state-space modeling for supervised and unsupervised electrocardiographic imaging. IEEE Transactions on Medical Imaging **PP** (2024). https://doi.org/10.1109/TMI.2024.3377094, epub ahead of print

9. Kipf, T.N., Welling, M.: Semi-Supervised Classification with Graph Convolutional Networks. In: Proc. 5th Int. Conf. Learn. Representations (ICLR). ICLR (Apr 2017), https://openreview.net/forum?id=SJU4ayYgl
10. Lecun, Y., Bottou, L., Bengio, Y., Haffner, P.: Gradient-based learning applied to document recognition. Proceedings of the IEEE **86**(11), 2278–2324 (1998). https://doi.org/10.1109/5.726791
11. Misra, S., van Dam, P., Chrispin, J., Assis, F., Keramati, A., Kolandaivelu, A., Berger, R., Tandri, H.: Initial validation of a novel ecgi system for localization of premature ventricular contractions and ventricular tachycardia in structurally normal and abnormal hearts. Journal of Electrocardiology **51**(5), 801–808 (2018). https://doi.org/10.1016/j.jelectrocard.2018.05.018
12. Ongie, G., Jalal, A., Metzler, C.A., Baraniuk, R.G., Dimakis, A.G., Willett, R.: Deep learning techniques for inverse problems in imaging. IEEE Journal on Selected Areas in Information Theory **1**(1), 39–56 (2020). https://doi.org/10.1109/JSAIT.2020.2991563
13. Parreira, L., Carmo, P., Adragao, P., Nunes, S., Soares, A., Marinheiro, R., et al.: Electrocardiographic imaging (ecgi): What is the minimal number of leads needed to obtain a good spatial resolution? Journal of Electrocardiology **62**, 86–93 (2020). https://doi.org/10.1016/j.jelectrocard.2020.07.004
14. Potyagaylo, D., Chmelevsky, M., Kalinin, A.: Single-layer based algorithms for solving the inverse problem of ecg. In: 2019 Computing in Cardiology (CinC). pp. Page 1–Page 4 (2019). https://doi.org/10.22489/CinC.2019.408
15. Rueda, C., Larriba, Y., Peddada, S.D.: Frequency modulated möbius model accurately predicts rhythmic signals in biological and physical sciences. Scientific Reports **9**, 18701 (2019). https://doi.org/10.1038/s41598-019-54569-1
16. Rueda, C., Rodríguez-Collado, A., Fernández, I., Canedo, C., Ugarte, M.D., Larriba, Y.: A unique cardiac electrocardiographic 3D model. toward interpretable AI diagnosis. iScience **25**(12), 105617 (2022). https://doi.org/10.1016/j.isci.2022.105617
17. Scarselli, F., Gori, M., Tsoi, A.C., Hagenbuchner, M., Monfardini, G.: The graph neural network model. IEEE Trans. Neural Netw. **20**(1), 61–80 (jan 2009). https://doi.org/10.1109/TNN.2008.2005605
18. Veličković, P., Casanova, A., Liò, P., Cucurull, G., Romero, A., Bengio, Y.: Graph attention networks. In: Proc. 6th Int. Conf. Learn. Representations (ICLR). ICLR (Oct 2018). https://doi.org/10.17863/CAM.48429
19. Verardo, G.: ECGI Repository (2024), https://github.com/giacomoverardo/ecg-imaging, accessed: 2024-08-15
20. Verardo, G., Boman, M., Bruchfeld, S., Chiesa, M., Koch, S., Maguire Jr., G.Q., Kostic, D.: FMM-Head: enhancing autoencoder-based ecg anomaly detection with prior knowledge (2023). https://doi.org/https://doi.org/10.48550/arXiv.2310.05848
21. Webber, M., Joy, G., Bennett, J., Chan, F., Falconer, D., Shiwani, H., et al.: Technical development and feasibility of a reusable vest to integrate cardiovascular magnetic resonance with electrocardiographic imaging. Journal of Cardiovascular Magnetic Resonance **25**(1), 73 (2023). https://doi.org/10.1186/s12968-023-00980-7
22. Wu, Z., Pan, S., Chen, F., Long, G., Zhang, C., Yu, P.S.: A comprehensive survey on graph neural networks. IEEE Trans. Neural Netw. **32**(1), 4–24 (2021). https://doi.org/10.1109/TNNLS.2020.2978386
23. Yang, Y., Rocher, M., Moceri, P., Sermesant, M.: Explainable electrocardiogram analysis with wave decomposition: Application to myocardial infarction detection. Statistical Atlases and Computational Models of the Heart. Regular and CMRxMotion Challenge Papers pp. 221–232 (2022). https://doi.org/10.1007/978-3-031-23443-9_21