

**PULSE**

**A DL-assisted physics-based approach to the inverse problem of electrocardiography**

Ugurlu, Kutay; Akar, Gozde B.; Dogrusoz, Yesim Serinagaoglu

**DOI**

[10.1109/TBME.2024.3501732](https://doi.org/10.1109/TBME.2024.3501732)

**Publication date**

2025

**Document Version**

Final published version

**Published in**

IEEE Transactions on Biomedical Engineering

**Citation (APA)**

Ugurlu, K., Akar, G. B., & Dogrusoz, Y. S. (2025). PULSE: A DL-assisted physics-based approach to the inverse problem of electrocardiography. *IEEE Transactions on Biomedical Engineering*, 72(4), 1328-1339. <https://doi.org/10.1109/TBME.2024.3501732>

**Important note**

To cite this publication, please use the final published version (if applicable). Please check the document version above.

**Copyright**

Other than for strictly personal use, it is not permitted to download, forward or distribute the text or part of it, without the consent of the author(s) and/or copyright holder(s), unless the work is under an open content license such as Creative Commons.

**Takedown policy**

Please contact us and provide details if you believe this document breaches copyrights. We will remove access to the work immediately and investigate your claim.

***Green Open Access added to TU Delft Institutional Repository***

***'You share, we take care!' - Taverne project***

**<https://www.openaccess.nl/en/you-share-we-take-care>**

Otherwise as indicated in the copyright section: the publisher is the copyright holder of this work and the author uses the Dutch legislation to make this work public.

# PULSE: A DL-Assisted Physics-Based Approach to the Inverse Problem of Electrocardiography

Kutay Ugurlu , Gozde B. Akar , Senior Member, IEEE, and Yesim Serinagaoglu Dogrusoz 

**Abstract**—This study introduces an innovative approach combining deep-learning techniques with classical physics-based electrocardiographic imaging (ECGI) methods. Our objective is to enhance the accuracy and robustness of ECGI reconstructions. We reshape the optimization expression by splitting variables and formulating building blocks based on update expressions. Specifically, we propose a sequential application of analytical solutions and denoiser neural network blocks (PULSE). The denoiser learns the proximal operator associated with the prior distribution of cardiac potentials directly from data, avoiding hand-crafted assumptions about the distribution. The proposed method is compared with zero-order Tikhonov regularization, Bayesian MAP estimation, and an end-to-end learning technique. We achieved more than 10% improvement in all metrics over Bayesian-MAP, end-to-end learning, and Tikhonov solutions. The performance remained consistent throughout cardiac beats, resulting in a 60% reduction in the interquartile ranges of the reconstruction metrics. Geometric variations did not compromise accuracy, with a median localization error consistently below 1cm. Our framework, adaptable to classical methods, augments the clinical pipeline. Improving the accuracy and robustness of pacing site localization holds significant promise for premature ventricular contraction (PVC) research.

**Index Terms**—Deep learning, ECGI, inverse problem, learned priors.

## I. INTRODUCTION

**E**LECTROCARDIOGRAPHIC imaging (ECGI) is a non-invasive imaging modality that aims to diagnose cardiovascular diseases. It provides a trade-off between the convenient but low-resolution nature of 12-lead electrocardiograms (ECG) and the accuracy of invasive intracardiac mapping.

Received 26 June 2024; revised 25 September 2024 and 6 November 2024; accepted 10 November 2024. Date of publication 18 November 2024; date of current version 21 March 2025. This work was supported in part by the Scientific and Technological Research Council of Turkey under Grant 221N175, and in part by the TUBITAK. (Corresponding author: Kutay Ugurlu.)

Kutay Ugurlu is with the Technical University of Delft, 2628, CD Delft, Netherlands, and also with the Electrical and Electronics Engineering, Middle East Technical University, 06800 Ankara, Turkey (e-mail: k.ugurlu@tudelft.nl).

Gozde B. Akar is with the Department of Electrical and Electronics Engineering, Middle East Technical University, Türkiye, and also with the Department of Radiology Weill Cornell Medical School, USA.

Yesim Serinagaoglu Dogrusoz is with the Middle East Technical University, Türkiye.

The code is shared publicly at <https://github.com/kutay-ugurlu/PULSE-Proximal-Unrolling-for-Learning-Solutions-in-ECGI>.

Digital Object Identifier 10.1109/TBME.2024.3501732

In ECGI, cardiac electrical activity is estimated from the body surface potential (BSP) measurements in terms of the intensity of the equivalent cardiac sources. Since BSPs are subject to attenuation and smoothing within the human body, the inverse problem becomes ill-posed when the cardiac sources are modeled as the heart surface potentials. In order to find a solution to this ill-posed problem, regularization should be applied to the solution vector, confining the set of possible solutions to a reasonable space.

Several conventional approaches are employed to solve the inverse problem. Bayesian MAP (MAP) and Tikhonov (Tik) regularization are examples of the approaches that exploit the statistical properties of the solution [1], [2], [3]. On the other hand, approaches like truncated singular value decomposition (TSVD) and truncated total least squares (TTLS), utilize only forward matrix properties to estimate the solution [4], [5], [6].

Recently, learning-based methods have exhibited great success in ECGI reconstructions. These approaches mainly utilize neural networks in different phases of the reconstruction pipeline: I) to estimate the transformation between the torso and heart potentials, II) to find excitation origin directly from measurements, III) to augment the physics-based model [7], [8], [9], [10].

In [8], Bacoyannis et al. used a simulated training set on six cardiac anatomies and on a single pacing site to train the conditional variational autoencoder model that reconstructs the activation pattern throughout the myocardium in a volumetric manner. They used 2D convolutions to capture spatio-temporal correlations in the data. In another study, Pilia et al. generated a simulated dataset to train the Convolutional Neural Networks (CNNs) to localize the ventricular excitation origin in a two-step process [10]. They first reflected the BSPs to a Cartesian grid. Following that, the first network learned the activation time boundaries in temporal direction whereas the networks in the second step learned to output the ventricular coordinates in Cobiveco that correspond to the excitation origin [11].

Tenderini et al. used a partial differential equation (PDE) aware deep learning model for the inverse solution [9]. In this physics-based hybrid approach, physical awareness was pursued by the projection of the epicardial potentials on a subspace spanned by the solutions of the governing PDEs and the inclusion of a tensorial reduced basis solver. To train their networks, they simulated and used a synthetic training dataset using the Aliev-Panfilov model for the EGMs and the corresponding 12-lead ECGs.

Chen et al. proposed an end-to-end training approach to relate the epicardial and torso potentials of the pigs [12]. They proposed 2 ways to utilize the common neural network operations like convolutions and Long Short Term Memory (LSTM) layers. In the first part, they matched the directly recorded 1D torso ECG data and 1D electrograms (EGMs) by training a network that is composed of LSTM and fully connected layers. In the second part, they projected the 3D torso geometry on a 2D plane. They upsampled the 1D ECG data on this 2D plane by bilinear interpolation and used a CNN to match the transformed torso data to the 1D EGMs. They found that CNNs were better at reconstructing the activation time maps. They also pointed out that training data patterns in different cross-validation trials (data coming from different pigs) have shown great variability in performance.

Jiang et al. represented geometric meshes as graphs and redefined convolutional operations on this non-Euclidean geometry framework in [7]. In this spatiotemporal graph CNN, the spatial convolutions are defined with spline-based kernels, whereas the temporal convolutions are kept as the standard 1D convolutions after the spatial convolution. The signals are directly represented on their own domain, that is to say, heart and torso geometries. The physics-based inverse imaging function was designed to be learned by a bipartite graph, assuming that the linearity between heart and torso potentials in the latent space still holds. The linear combination of the vertex voltages of the latent torso geometry expresses the latent heart potentials on the vertices of the latent heart geometry. A later study showed that this methodology showed similar uncertainty with Tikhonov regularization in reconstructed cardiac sources resulting from cardiac and torso geometry variability [13]. It was also concluded that the graph neural network (GNN) based method was less sensitive to the modeled geometric uncertainty. However, the authors stated that the reported generalization metrics were the results of the fine-tuning experiments conducted with a small amount of data. Even with fine-tuning on the geometries that differ from the training data only in terms of rotation patterns, the performance drop is observed in the GNN-based technique [14].

In end-to-end learning-based techniques, the training data is leveraged to learn the conditional probability distribution. This corresponds to reconstructing the heart potentials directly from the measurements. In contrast to this approach, classical methods use training data to estimate the prior distribution of the heart potentials and regularize the solution. However, these methods have limitations. They either necessitate extensive training data, rely on task-specific layers or architectures to comprehend the conditional probability or fall short in generalizing to diverse prior distributions and geometries.

Contrary to these approaches, we decouple reconstruction and regularization into separate steps while also incorporating the implicit prior distribution by learning the corresponding proximal operator. By employing a hybrid approach that combines the neural network and the physical model, we eliminate the need for the neural network to explicitly learn both the inversion of domain transformation and regularization. In this study, we introduce PULSE: Proximal Unrolling for Learning Solutions in ECGI. The proposed method is mainly composed of two

steps obtained by separating the loss function by half quadratic splitting (HQS), where the first step deals with physics-based estimation of cardiac sources using the forward model, and the second step applies the proximal operator for the implicit regularization that is associated with the implicit prior [15]. The second step enables us to determine the implicit prior by learning the proximal operator [16].

The main contributions of this study can be listed as follows:

- To the best of our knowledge, this is the first study that uses variable splitting in the field of ECGI.
- We propose a hybrid solution that allows the utilization of physical knowledge and a network that can perform proximal updates under geometric variations without significant performance loss.
- Our proposed design is a compatible framework that fits any dataset and forward model, paving the way for other researchers to experiment with the method on larger datasets.

We compare the proposed method with three benchmark methods from the literature: Tikhonov regularization (**Tik**), Bayesian MAP estimation (**MAP**) [17], [18], and an end-to-end learning technique (**E2ET**) proposed in [12]. EGMs with ventricular paced beats measured from seven dog hearts in a torso-tank setup and BSPs simulated from these EGMs are used for training and testing of the methods. Results are assessed in terms of ECGI metrics such as relative error, correlation coefficient, and localization error.

## II. ECGI INVERSE PROBLEM

The ECGI inverse problem can be mathematically formulated in the following form:

$$\mathbf{y}^{(t)} = \mathbf{A}\mathbf{x}^{(t)} + \mathbf{n}^{(t)}, \quad t = 1, 2, 3, \dots, T \quad (1)$$

Here,  $\mathbf{y}^{(t)}$  represents the torso measurements for  $M$  torso electrodes at the time instance  $t$ ,  $\mathbf{A}_{M \times N}$  is the forward model and  $\mathbf{x}^{(t)}$  is the epicardial distribution or cardiac voltage at  $N$  discretized source location at the time instance  $t$ . From now on, the superscript  $(t)$  will be omitted for simplicity in notation.

The solution is

$$\hat{\mathbf{x}} = \arg \min_{\mathbf{x}} \mathbf{J}(\mathbf{x}) \quad (2)$$

where

$$\mathbf{J}(\mathbf{x}) = \underbrace{\|\mathbf{y} - \mathbf{A}\mathbf{x}\|_2^2}_{\text{Data-fidelity term}} + \underbrace{\mu}_{\text{Regularization parameter}} \underbrace{\mathbf{R}(\mathbf{x})}_{\text{Regularization term}} \quad (3)$$

The data-fidelity term penalizes the mismatch between the observations and measurements that the solutions generate through the forward model, whereas the regularization term enforces a condition by adding a cost to the minimization expression utilizing a penalization. We characterize the data-fidelity term by  $\mathcal{L}_2$ -norm due to normally distributed i.i.d. noise in our simulations (see Section III). The regularization parameter  $\mu$  is a trade-off parameter that provides a balance between the residual error and the prior knowledge enforced on the solution.

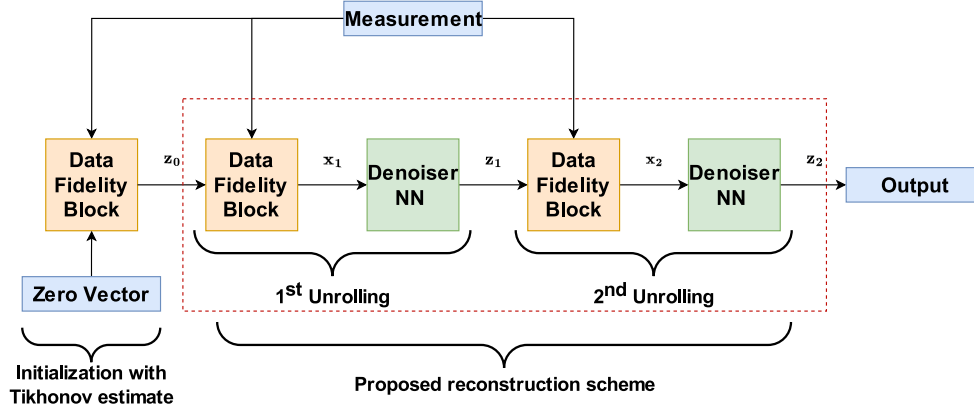


Fig. 1. A sample architecture for the HQS method for two unrolling iterations. The first estimate of the solution is calculated with zeroth order Tikhonov regularization. This estimate is fed to the optimization iterations that start with an analytical solution, followed by the neural network. (This is an illustrative figure of the iteration scheme used. The study utilizes 1 unrolling iteration.)

Employing regularization techniques presents a feasible solution to the issues of uniqueness and ill-posedness inherent in the problem [19]. The regularization is incorporated into a vector-valued cost function  $\mathbf{J}$  that outputs the total penalty of the solution. In the inverse problem setting, the cost function  $\mathbf{J}$  is generally composed of data-fidelity and regularization parts as in (3).

#### A. Variable Splitting With HQS

In the literature, the regularization term is determined based on a hand-crafted prior, such as Gaussian or Laplacian. Instead of minimizing with a manually selected prior, we use a variable splitting approach to decouple the problem. Variable splitting allows us to find a solution by alternatively optimizing the data fidelity term and the regularization term that corresponds to the a priori information (Fig. 1).

In the HQS method, the optimization equation given in (3) is transformed into an optimization problem with the introduction of intermediate variable  $\mathbf{z}$  and regularization parameter  $\mu$  [15]. The new constrained problem can be formulated as follows:

$$\hat{\mathbf{x}} = \arg \min_{\mathbf{x}} (\|\mathbf{y} - \mathbf{A}\mathbf{x}\|_2^2 + \mu\mathbf{R}(\mathbf{z})), \quad \text{s.t. } \mathbf{x} = \mathbf{z} \quad (4)$$

Then, the corresponding unconstrained problem is:

$$\mathbf{J}(\mathbf{x}, \mathbf{z}) = \frac{1}{2}\|\mathbf{y} - \mathbf{A}\mathbf{x}\|_2^2 + \mu\mathbf{R}(\mathbf{z}) + \frac{\lambda}{2}\|\mathbf{x} - \mathbf{z}\|_2^2 \quad (5)$$

where  $\lambda$  is another regularization parameter that penalizes the difference between  $\mathbf{x}$  and  $\mathbf{z}$ . (5) can be split into two equations for alternating optimization on variables  $\mathbf{x}$  and  $\mathbf{z}$  at every iteration  $k$ :

$$\mathbf{x}_{k+1} = \arg \min_{\mathbf{x}} (\|\mathbf{y} - \mathbf{A}\mathbf{x}\|_2^2 + \lambda\|\mathbf{x} - \mathbf{z}_k\|_2^2) \quad (6a)$$

$$\mathbf{z}_{k+1} = \arg \min_{\mathbf{z}} \left( \frac{\lambda}{2}\|\mathbf{z} - \mathbf{x}_{k+1}\|_2^2 + \mu\mathbf{R}(\mathbf{z}) \right) \quad (6b)$$

The expression in (6a) measures how reconstruction matches the observation, hence data-fidelity, whereas (6b) corresponds to a denoising problem defined on  $\mathbf{z}$ .

The solution to (6a) can be found analytically as follows:

$$\mathbf{x}_{k+1} = (\mathbf{A}^T \mathbf{A} + \lambda \mathbf{I})^{-1} (\mathbf{A}^T \mathbf{y} + \lambda \mathbf{z}_k) \quad (7)$$

which is a Tikhonov variant that seeks a solution around a previous estimate.

When the regularization term in (6b) is rewritten as in (8a), the optimization step corresponds to denoising the signal with Gaussian i.i.d. denoiser with noise level  $\sqrt{\frac{\mu}{\lambda}}$  according to Bayesian probability [15]. The Gaussianity is due to the  $\mathcal{L}_2$ -norm-squared data fidelity term with the implicit forward operator identity.

$$\mathbf{z}_{k+1} = \arg \min_{\mathbf{z}} \frac{1}{2 \left(\sqrt{\frac{\mu}{\lambda}}\right)^2} \|\mathbf{z} - \mathbf{x}_{k+1}\|_2^2 + \mu\mathbf{R}(\mathbf{z}) \quad (8a)$$

$$\mathbf{z}_{k+1} = \text{Denoiser} \left( \mathbf{x}_{k+1}, \sqrt{\frac{\mu}{\lambda}} \right) \quad (8b)$$

This step is a proximal update that minimizes the regularization penalty in the proximity of the previous estimate coming from the data-fidelity update. The alternating optimization steps for variables  $\mathbf{z}$  and  $\mathbf{x}$  are applied to incorporate the physics (the forward model) of the problem and the implicit prior. Since the prior is implicit, one of the most robust approaches is to learn the prior from the data using neural networks. The regularization term is going to be kept implicit so that the neural network learns to apply the proximal operator on the subspace that is unknown to us. For comparison, in the zero-order-Tikhonov where the regularization term is  $\mathbf{R}(\mathbf{z}) = \|\mathbf{z}\|_2^2$ , this step would correspond to minimizing the energy of the solution in the proximity of the inverse estimate.

In the literature, several studies utilize neural networks in additive-noise removal tasks [20], [21], [22], [23]. They are also widely used in denoising step of image restoration problems to learn the proximal operator associated with the implicit prior [15], [24], [25], as well as in super-resolution imaging tasks [26]. In these tasks, UNet architecture and Res-Net blocks are highly utilized. In this study, we build a UNet-based architecture with the following modifications:

To accelerate the learning of high-frequency content,

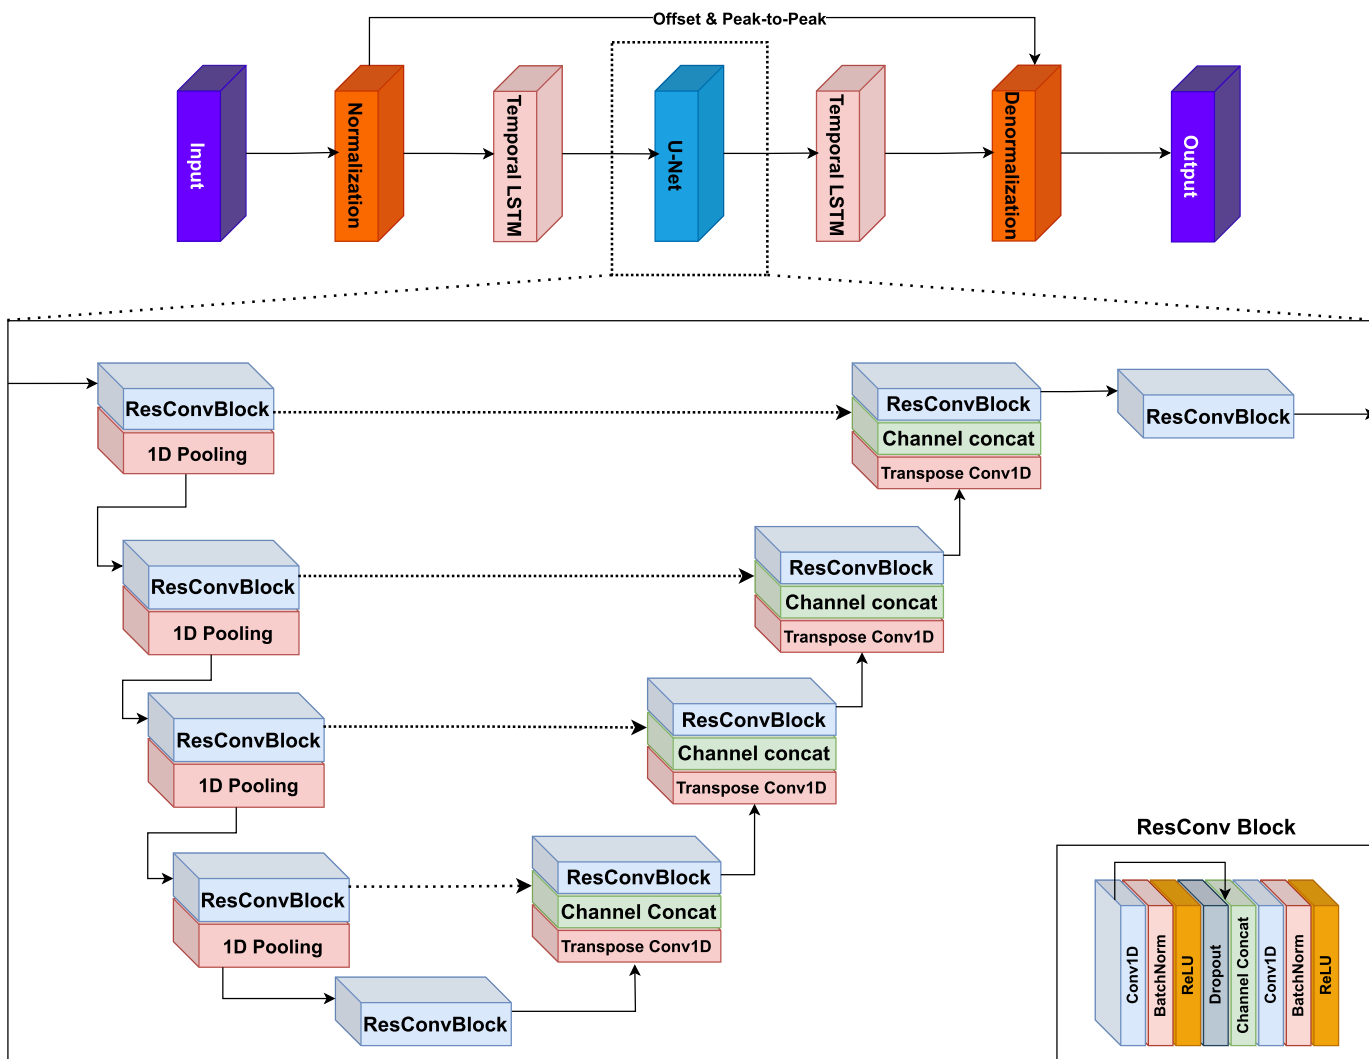


Fig. 2. The proposed architecture where the basic U-Net structure is augmented with LSTM blocks to learn both the spatial and the temporal priors.

- A residual skip connection is added between input and output.
- The encoder and decoder blocks are augmented with residual skip connections as in [24].

In addition, the architecture is prepended and appended to UNet-based structure in Fig. 2 with LSTM blocks to learn the temporal prior.

The neural network model and the regularization parameter of the data-fidelity block are optimized by the Adam optimizer with learning rates  $10^{-4}$  and  $10^{-8}$ , respectively. The learning rate (LR) of the regularization parameter is intentionally set to provide a slow update to avoid the sensitivity of the solution to the regularization parameter. The model parameters are updated after every batch of data is forward-fed to the network, whereas the update period for the regularization parameter is one per every epoch. This approach also guarantees that the sensitive parameter does not get updated with just one batch of data, which may deviate the solution for the next batches further. Instead, the gradient is accumulated, and the mean of the total gradient for every batch is applied to the regularization parameter.

## B. Benchmark ECGI Methods

We compare the reconstructions obtained from **PULSE** with three benchmark methods from the literature: the zero-order Tikhonov regularization (**Tik**) [27], the Bayesian-MAP estimation (**MAP**) method proposed in [17], [18], and an end-to-end trained (**E2ET**) fully connected neural network model proposed in [12].

**Tik** solution is defined by (2) and (3), where the regularization term constrains the energy of the heart potentials. The regularization parameter is selected by the L-curve method (see Section III-G) [28].

**MAP** approach maximizes the a posteriori probability distribution function (pdf) of the heart potentials given the BSP measurements [17], [18]. We assume that the epicardial potentials and BSPs are jointly Gaussian; a priori pdf of the heart potentials is a Gaussian distribution with a well-defined mean and covariance; noise in the measurements are zero-mean Gaussian and uncorrelated with the epicardial potentials.

**E2ET** model is a fully connected neural network model trained with end-to-end approach [12]. It was trained using the Adam optimizer, starting with a learning rate of  $10^{-3}$ .

### III. EXPERIMENTAL DESIGN

#### A. Experimental Data

In this study, data from seven experiments (7 different dog hearts) conducted at the University of Utah are utilized for training and testing the denoising neural network [29]. In these experiments, the isolated heart from one canine, which is perfused by a support dog's circulatory system, is suspended in a torso tank filled with electrolytic solution. All experiments are performed under deep anesthesia via procedures approved by the Institutional Animal Care and Use Committee of the University of Utah, conforming to the Guide for the Care and Use of Laboratory Animals (02-03010). A sock that contains 490 electrodes is placed on the surface of the heart, and the measurements are sampled at a 1kHz sampling frequency.

The hearts are stimulated from different pacing sites that cover the epicardial surface of the ventricles (left (LV) and right (RV)). In total, the complete dataset consists of 380 beats from 7 experiments (3 for training and 4 for testing the methods). After removing the data that contains artifacts, 326 beats are included in the study.

The training data from 3 experiments consists of 309 beats with 185 unique pacing locations. This global training dataset is then split into training and validation sets with a 0.05 ratio for network weight updating and hyperparameter tuning purposes. The same training dataset is used for training and validating the **E2ET** method, and estimating the prior mean and covariance for **MAP**.

From the remaining 4 experiments, 17 beats with unique pacing locations (on LV and RV) are randomly selected to achieve pacing site diversity over the ventricles.

The BSPs are simulated from the experimental epicardial potential measurements by solving the forward problem in an inhomogeneous torso including the lungs (see Section III-B), and distorting the computed BSPs with independent and identically distributed (i.i.d.) noise at 20 dB SNR (the ratio of the signal variance to noise variance [30]).

#### B. Forward Problem Solution

The forward problem is solved using the boundary element method (BEM) [31], [32] on a torso mesh with 771 nodes ( $24.63 \pm 6.12$  mm) and a heart mesh containing 490 cardiac nodes ( $5.21 \pm 0.97$  mm). After the transfer matrix is computed, the forward model and torso geometry are reduced such that we have 192 measurement locations proposed by [33].

While an inhomogeneous torso model containing lungs is used for simulating the BSPs, a forward matrix computed from a homogeneous torso model is used in the inverse computations to avoid inverse crime [34].

#### C. Geometric Model Variation

Originally, only one heart-torso geometry is used for generating the training and test datasets. To assess how the

---

#### Algorithm 1: Training Data Creation.

---

**Require:** A training database with  $N$  beats, **batchsize**

- 1: **for** Beat in the training set **do**
- 2:    $T \leftarrow$  time length of the beat
- 3:   **if**  $T < \text{batchsize}$  **then**
- 4:     Repeat the beat in the time dimension  $\lceil \frac{T}{\text{batchsize}} \rceil$
- 5:     **return** Take the first **batchsize** time instances starting from  $t = 0$
- 6:   **else if**  $T = \text{batchsize}$  **then**
- 7:     **return** Beat itself
- 8:   **else**
- 9:     Split the beat into equal pieces of size **batchsize** and save
- 10:    Generate a final overlapping sequence of length **batchsize** from the end of the beat and save
- 11:    Concatenate saved beats in another dimension
- 12:    **return** The concatenated tensor
- 13:   **end if**
- 14: **end for**

---

proximal operator learned by **PULSE** works under different forward/inverse models for the training and test datasets, the heart geometry is later manipulated in the form of rotations for the test data generation. However, the heart geometry is kept as in the original case for the training data generation.

For this, we calculate the maximum allowable rotation of the heart in the inhomogeneous geometry that is utilized to simulate the BSPs. In all three directions, the rotation angles are individually swept from  $-180^\circ$  to  $180^\circ$ , and the angle intervals at which no intersection happens between the heart and lungs are determined as  $[0, -37]$ ,  $[0, -7]$ , and  $[0, 21]$  for the  $x$ ,  $y$  and  $z$  directions, respectively.

Three new heart geometries are generated by applying rotation operators  $\mathcal{R}_x$ ,  $\mathcal{R}_y$ , and  $\mathcal{R}_z$  on the original heart geometry corresponding to the most extreme allowable angles  $-37$ ,  $-7$  and  $21$ , respectively. Modified forward matrices (inhomogeneous and homogeneous) and new test data BSPs are then calculated for these new heart geometries.

Note that in all these scenarios no further geometric error is introduced to the test datasets, *i.e.*, the BSPs are simulated, and the inverse problem is solved on the same heart-torso model except the presence of lungs in the simulation forward model.

#### D. Data Preprocessing

The data from experiments have shape  $[490 \times T_i]$  where  $T_i$  corresponds to the number of time instances EGMs are captured for training data  $i$ . Due to the variations in  $T_i$ , Algorithm 1 is implemented to generate batches of the desired size from the training data. First, a batch size is determined and is given as input to the algorithm. We employ cyclic concatenation to augment data whose time dimension falls short of the desired batch size. If the training data contains more time instances than the desired batch size, it is split into mini-batches of the desired batch size, starting from the first time instance. The last incomplete batch is completed by prepending the previous time instances. These operations are illustrated in Fig. 3.

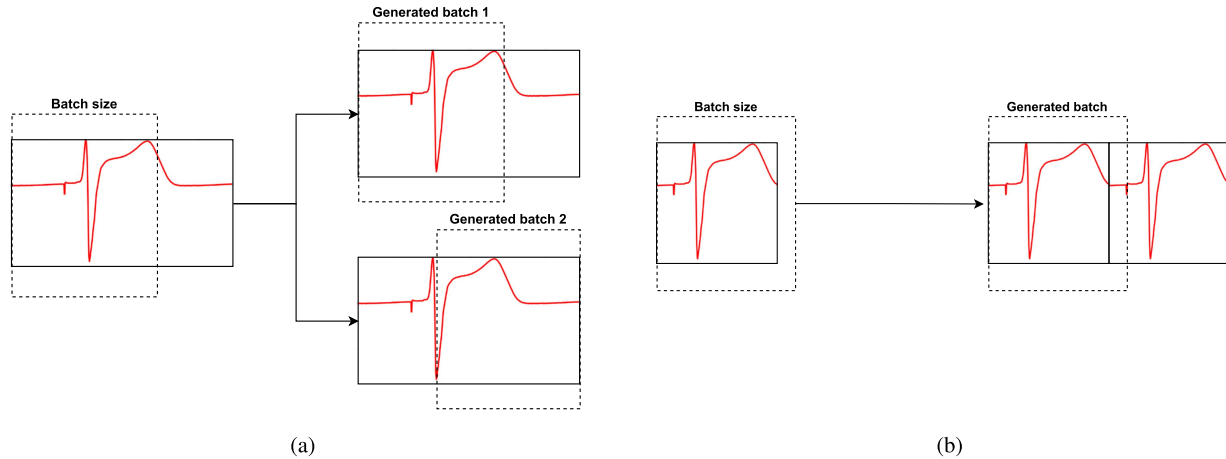


Fig. 3. Illustration of Algorithm 1 for a single EGM. (a) BatchSize <  $T$ . (b) BatchSize >  $T$ .

TABLE I  
PERFORMANCE OF BI-LSTM UNET WITH DIFFERENT BATCH SIZE CONFIGURATIONS. (BEST CASE HIGHLIGHTED WITH **BOLD**)

Batch size	Train CC	Train RE	Validation CC	Validation RE	Train Loss	Validation Loss
16	0.963	0.205	0.848	0.486	<b>3.121</b>	13.203
32	0.965	0.202	0.912	0.349	3.410	7.698
64	<b>0.967</b>	<b>0.195</b>	<b>0.929</b>	<b>0.319</b>	3.552	<b>6.746</b>
128	0.945	0.257	0.888	0.397	4.116	7.807

To determine the batch size, which directly affects the model architecture, a series of experiments are run, and the training and validation results are inspected.

Table I shows that 64 as batch size is the optimal choice to capture the temporal correlations in the training data, hence the batch size parameters of the architectures that learn spatiotemporal prior are fixed to 64.

In addition, to accelerate the neural network training, the inputs are normalized into the range of  $[0, 1]$ . The normalized inputs are fed into the neural network, and the outputs of the network are denormalized using the inverse process, by first amplifying by previously saved peak-to-peak amplitude and then shifting by the minimum of the input. In the inference step, the same normalization and denormalization steps are applied to the test data.

The train-validation split of the data is executed with a fixed random seed to prevent performance fluctuations due to the random split. The train-validation ratio was set to 0.05, hence 15 test beats are used for validation and 294 beats are used for training. The batch generation (Algorithm 1) is applied on these datasets independently to prevent the training data batches from appearing in validation data.

### E. Node sorting

Due to the nature of the coordinate systems on which the signals in these problems are defined, CNNs are not directly applicable on randomly stacked EGMs stored in regular 2D matrices since the measurements and the epicardial signals of the ECGI inverse problem are defined on unstructured grids.

To address this challenge and leverage Euclidean CNNs for denoising, an algorithm that reorders the nodes by sorting numbering is implemented. This deterministic step maximizes spatial proximity in any mesh structure. The number of edges defines the neighborhood order between two nodes in the shortest connecting path on the geometry graph. The algorithm aims to decrease the mean neighborhood order of the nodes within the convolutional neural network's kernel. This allows the kernels, defined in regular grids, to learn features from spatially relevant signals on the geometry. The strategy needs to be implemented only once for each new geometry and can be subsequently used once the nodes are reordered accordingly. A step-by-step explanation of the algorithm can be found in [35].

### F. Comparison Metrics

We compare the reconstructions given by **PULSE**, **E2ET**, **MAP** and **Tik** to the ground truth solutions. In these evaluations, we use relative error and correlation coefficient metrics given in (9), where the ground truth is  $\mathbf{x}$  and the estimated solution is  $\hat{\mathbf{x}}$ .

$$\text{Relative Error (RE): } \frac{\|\mathbf{x} - \hat{\mathbf{x}}\|_2}{\|\mathbf{x}\|_2}$$

$$\text{Pearson's CC: } \frac{(\mathbf{x} - \bar{\mathbf{x}})^T (\hat{\mathbf{x}} - \bar{\hat{\mathbf{x}}})}{\sqrt{(\mathbf{x} - \bar{\mathbf{x}})^T (\mathbf{x} - \bar{\mathbf{x}})} \sqrt{(\hat{\mathbf{x}} - \bar{\hat{\mathbf{x}}})^T (\hat{\mathbf{x}} - \bar{\hat{\mathbf{x}})}}} \quad (9)$$

We use both the spatial and temporal CC and RE values. For the temporal metrics,  $\mathbf{x}$  and  $\hat{\mathbf{x}}$  represent single-lead EGMs.

Furthermore, the activation time is computed to compare the performance of the proposed method to the benchmark methods. We implement the spatially coherent activation time (AT) method on the reconstructed EGMs [36]. Then, we mark the node having minimum AT as the beat's excitation origin. We calculate the Localization Error (LE) as the Euclidean distance between the predicted origin of excitation and the actual pacing site. Even though we have obtained both spatial and temporal RE and CC, we demonstrate the methodological differences using only temporal RE and spatial CC since they are representative enough to provide information about how well the reconstructed

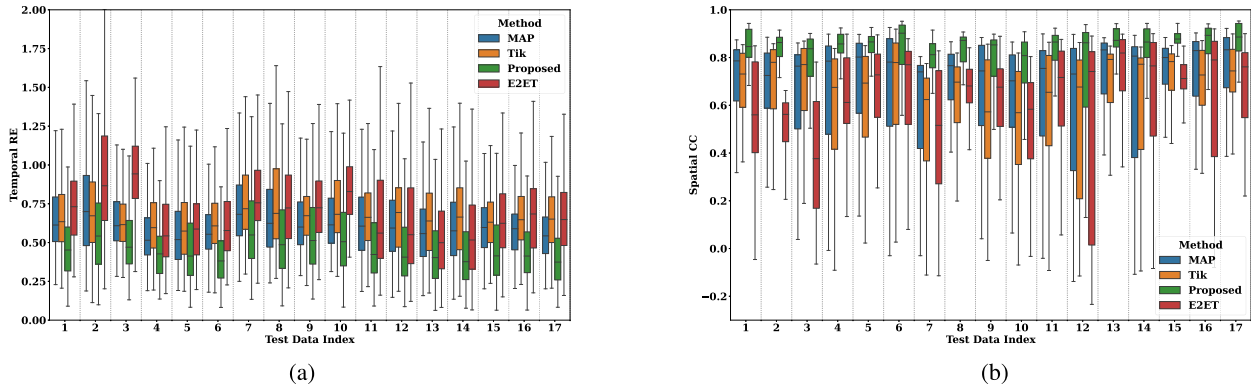


Fig. 4. Comparison of the methods in terms of spatial CC and temporal RE. This indicates that the spatial map matches are much more consistent throughout the beat due to the utilization of the temporal prior. Furthermore, distributions coming from **PULSE** are shifted towards the improvement direction (towards 0 in RE and towards 1 in CC) when compared to the others, implying overall better performance. (a) Temporal RE distribution comparison. (b) Spatial CC distribution comparison.

TABLE II  
RECONSTRUCTION METRICS REGARDING THE WHOLE DATASET  
ILLUSTRATED IN MEDIAN(IQR) REPRESENTATION FOR THE TRAINING  
GEOMETRY

Test Metric	Value			
	Proposed	MAP	Tikhonov	E2ET[12]
Temporal CC	<b>0.90(0.04)</b>	0.79(0.06)	0.79(0.05)	0.76(0.15)
Temporal RE	<b>0.42(0.09)</b>	0.60(0.05)	0.66(0.05)	0.64(0.17)
Spatial CC	<b>0.87(0.02)</b>	0.77(0.06)	0.71(0.10)	0.71(0.16)
Spatial RE	<b>0.50(0.04)</b>	0.62(0.07)	0.68(0.10)	0.71(0.12)

EGMs fit the ground truth in terms of signal amplitudes and isopotential maps on the heart surface.

### G. Regularization Parameter Selection

The L-curve technique, which is highly effective for determining the regularization parameter in BEM-formulated problems, is used to find the regularization parameter  $\lambda$  in (6a). In this method, the norm of the constraint and the residual are plotted on a logarithmic scale, and the point with the highest curvature is selected as the optimal point. In this study, an exponential sweep covers a broad range of regularization parameters initially, followed by a linear sweep focusing around the L-corner. The optimal parameter corresponding to the L-corner is found as approximately  $1.75 \times 10^{-4}$ .

During the training phase, the regularization parameter was initialized with this value, with minor adjustments permitted during the training process. When the training was completed, we observed that the regularization parameter eventually converged to the value initially determined by the L-curve. We directly use this value for every inverse test solution without any modification.

## IV. RESULTS

### A. EGM Reconstruction Performance

In this section, the proposed method's performance is compared with **Tik**, **MAP** and **E2ET** on the original training geometry. The EGM reconstruction accuracy for all test data in

terms of RE and CC, in median (IQR) representation, is collected under Table II. The subject-specific boxplot distributions of the temporal RE and spatial CC metrics are given in Fig. 4.

In all the test data, the proposed method outperformed the remaining methods. On average, **PULSE** improves the reconstructions of **Tik** by 24%, **E2ET** by 22%, and **MAP** by 18% in terms of temporal RE. The temporal CC given by **PULSE** is observed to be the highest for every test beat, and it resulted in an improvement of 11%-14%.

When spatial metrics are analyzed, it is observed that the proposed method provides more stable and consistent reconstruction throughout different time instances of the beat. This effect can be observed in the decreased IQR by 80% and 60% in the spatial CC and RE, respectively. This can also be observed in Fig. 4 (panel (b)), where the spatial CC is plotted in box plots for all test beats. Moreover, the spatial CC distributions provided by **PULSE** are much more compact compared to the remaining methods. In panel (a), we observe that **PULSE** provides the lowest temporal RE for every test beat. These observations together demonstrate the benefit of learning the temporal component of the prior: **PULSE** provides spatially and temporally better and more consistent reconstructions. It is important to note that **E2ET** ended up having the highest IQR in the reconstruction metrics. In terms of the spatial metrics, this means that the performance is not temporally consistent throughout the beat, with more temporal fluctuations. For the temporal metrics, a low IQR means consistent reconstruction among different nodes on the heart. Due to a lack of spatiotemporal regularization, the **E2ET** model failed to reconstruct the inherent patterns of the out-of-distribution test data.

We also observed that **E2ET** performance varies across different test data. It performed best on test data 13, both in terms of RE and CC, and showed more consistent reconstructions compared to test data 8 and 15 in terms of spatial CC. However, it performed worse on most of the test data, which are out-of-distribution by design, as they do not share a common experiment subject or beat excitation origin.

**PULSE** outperformed **MAP** and **Tik** in terms of median AT-CC and LE values (see **Training** scenario in Fig. 7). It

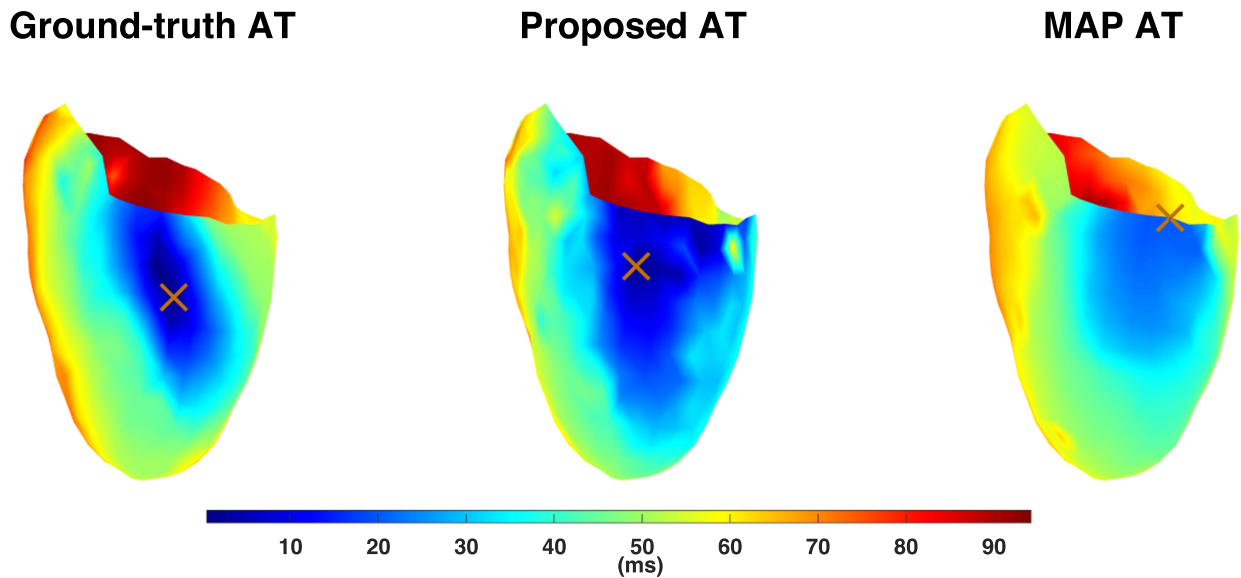


Fig. 5. The activation time maps computed on the reconstructions given by **MAP** and **PULSE** are compared to the ground truth. The brown cross represents the estimated pacing location by each method.




Fig. 6. EGM reconstructions at the ground truth excitation origin by different approaches in test data 3. The estimated activation time for each of the method is marked with a dot of same color on the regarding EGM curve.

provided a more compact distribution in both AT-CC and LE, with fewer outliers. The median AT-CC was improved by 3% and 6% compared to **MAP** and **Tik**, respectively, and LE was reduced by approximately 50% compared to the other methods.

To further investigate the single cases, we took a closer look at one of the test data in Fig. 5 where the LE was higher than 5 mm. **Tik** had the poorest performance in recreating the data. **MAP** caused the highest amount of latency in the reconstructions since it produced a lighter region of blue in the earliest activated region, which means delayed activation times. It is expected because **MAP** only regularizes the solution spatially, and this spatial smoothing ends up concentrating the EGMs in the middle of the reconstruction timeline. With the spatiotemporal regularization enabled by **PULSE**, we managed to obtain an AT map that carries more detailed information. However, it produced a wider earliest-activated region compared

to the test AT. In Fig. 6, we plotted the EGMs reconstructed in the ground truth excitation origin to investigate where this difference comes from. We observe that **MAP** reconstruction was not able to filter out the noise sufficiently due to the lack of temporal prior knowledge. For the same reason, we observe noisy reconstructions from **E2ET**. Furthermore, the fully connected network in **E2ET** created a non-sensible EGM pattern as a prediction for an out-of-distribution test sample. On the other hand, we see smoothed reconstructions given by **PULSE** in the temporal direction. However, we see a latency in the QRS region in this reconstructed EGM.

We did not have any training data sample that has the same excitation origin as this test data. Hence, in this scenario, **PULSE** tends to reconstruct with latency for the specific node. From 24ms onwards, EGM reconstructed by **PULSE** follows the ground truth successfully. We attribute this improvement to the learned temporal a priori knowledge about the data.

### B. Generalization Under Geometric Variation

In this section, we evaluate the performance of the trained model under different rotational variations. Here, we keep the geometric model used for the training data the same as before, but we applied the rotations described in Section III-C to obtain three new test datasets with heart geometries different than the training data. In these evaluations, we excluded **E2ET** from the methods to focus more on the performance differences between similarly performing methods, considering that it performed poorer in reconstruction with training geometry.

In Table III, the median and IQR values of the spatial and temporal EGM metrics (CC and RE) are reported for the rotation scenarios,  $\mathcal{R}_x$ ,  $\mathcal{R}_y$ , and  $\mathcal{R}_z$ . Here, ‘Training’ denotes that the test dataset was constructed using the same heart-torso model as the training dataset. Note that **Tik** and **MAP** suffer no significant performance loss since they do not need the information from the

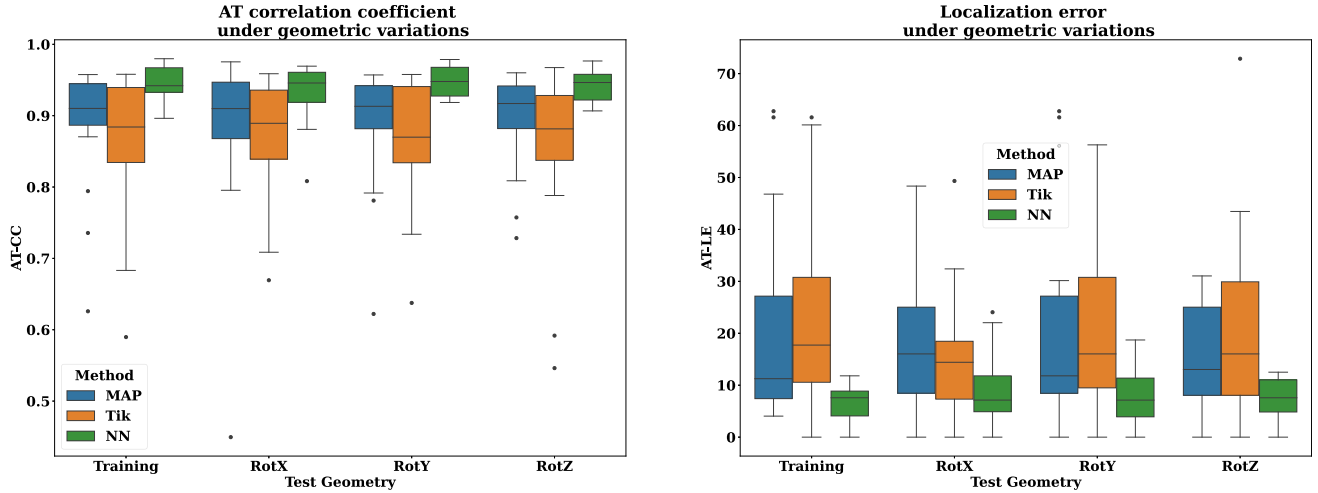


Fig. 7. AT correlation coefficient and localization error distributions for 4 different forward models.

TABLE III  
AVERAGE RELATIVE ERROR IN MEDIAN (IQR) REPRESENTATION UNDER DIFFERENT ROTATION SCENARIOS (BEST HIGHLIGHTED IN BOLD)

Test Geometry	Spatial RE			Temporal RE		
	Proposed	MAP	Tik	Proposed	MAP	Tik
Training	<b>0.50(0.04)</b>	0.62(0.07)	0.68(0.10)	<b>0.42(0.09)</b>	0.60(0.05)	0.66(0.05)
$\mathcal{R}_x$	<b>0.54(0.06)</b>	0.60(0.07)	0.69(0.08)	<b>0.48(0.08)</b>	0.58(0.05)	0.65(0.05)
$\mathcal{R}_y$	<b>0.51(0.04)</b>	0.62(0.07)	0.67(0.10)	<b>0.44(0.08)</b>	0.59(0.06)	0.66(0.06)
$\mathcal{R}_z$	<b>0.54(0.05)</b>	0.63(0.08)	0.70(0.09)	<b>0.46(0.09)</b>	0.57(0.07)	0.63(0.07)
	Spatial CC			Temporal CC		
	Proposed	MAP	Tik	Proposed	MAP	Tik
Training	<b>0.87(0.02)</b>	0.77(0.06)	0.71(0.10)	<b>0.90(0.04)</b>	0.79(0.06)	0.79(0.05)
$\mathcal{R}_x$	<b>0.84(0.05)</b>	0.79(0.06)	0.71(0.08)	<b>0.87(0.05)</b>	0.78(0.05)	0.78(0.05)
$\mathcal{R}_y$	<b>0.86(0.02)</b>	0.78(0.07)	0.71(0.10)	<b>0.90(0.04)</b>	0.79(0.06)	0.79(0.06)
$\mathcal{R}_z$	<b>0.84(0.04)</b>	0.76(0.07)	0.73(0.08)	<b>0.89(0.05)</b>	0.80(0.05)	0.79(0.05)

heart-torso model of the training dataset. Nevertheless, **PULSE** could still generalize to the maximum allowable geometric variation between the test and training datasets and produced superior results than **MAP** and **Tik** in 3 different rotations; it still outperforms **MAP** and **Tik** by at least 12% in RE and 10% in CC. However, it exhibited slight performance loss due to its training being dependent on the geometry. In  $\mathcal{R}_x$  and  $\mathcal{R}_z$  cases, the consistency (IQR) is deteriorated compared to the ‘Training’ geometry.  $\mathcal{R}_y$  reconstructions have not been affected significantly since the maximum allowable rotation was calculated as  $7^\circ$ .

We present the distributions of the AT metrics for the 3 extra test data sets in box plots of Fig. 7. **PULSE** continues to outperform the others under geometric variation in terms of both performance and consistency (median and IQR).

### C. Generalization Under Measurement Noise Variation

Lastly, we evaluated our original trained model (at 20 dB SNR) with test data BSPs corrupted with 30 dB SNR noise to see whether it only performs as reported under the trained noise level. The results presented in Table IV show that the generalization capability extends to a different noise level, still outperforming other benchmarks.

TABLE IV  
RECONSTRUCTION METRICS REGARDING THE WHOLE DATASET ILLUSTRATED IN MEDIAN(IQR) REPRESENTATION FOR THE TRAINING GEOMETRY UNDER 30 Db SNR NOISE FOR THE TEST DATA

Test Metric	Value		
	Proposed	MAP	Tikhonov
Temporal CC	<b>0.91(0.04)</b>	0.83(0.04)	0.82(0.03)
Temporal RE	<b>0.42(0.09)</b>	0.55(0.06)	0.63(0.04)
Spatial CC	<b>0.87(0.03)</b>	0.81(0.03)	0.76(0.10)
Spatial RE	<b>0.50(0.05)</b>	0.59(0.02)	0.66(0.10)

We compared these results with those in Table II. Even though the performance of other benchmark methods increased more than that of **PULSE** under 30 dB SNR noise, it is observed to be more robust to different levels of noise. This result is expected for methods that utilize relatively weak priors such as i.i.d Gaussian or correlated Gaussian.

### V. DISCUSSION

In this study, we proposed a modular physics-based neural network method for reconstructing epicardial signals from the BSP measurements. By leveraging the physical model information, we projected the measurements onto the inverse model range space, ensuring the solution is physically sensible. The performance of this method was evaluated and compared with

three benchmark methods from the literature: the widely used Tikhonov regularization (**Tik**), Bayesian MAP estimation that uses a training data of EGMs for estimating the prior pdf (**MAP**) [17], [18], and an end-to-end learning technique (**E2ET**) proposed in [12] that uses the same EGM-BSP training data pairs as **PULSE**.

We evaluated the performances in terms of well-known ECGI metrics such as spatial and temporal CC and RE, AT-CC, and LE. First, a thorough investigation was carried out in a simple scenario where the test and training data were generated on a single heart-torso geometry. Then, **PULSE** was evaluated for the model mismatch between the training and test data geometries. For this, three new heart-torso geometries different from the training geometry were obtained by rotating the heart in the  $x$ ,  $y$ , and  $z$  directions. New test data were simulated in these modified geometries and the network learned in the original geometry was applied directly to these new test data. Finally, **PULSE** was assessed for the measurement noise SNR mismatch, where the network trained at 20 dB SNR was used on test data simulated at 30 dB SNR.

The main take-home messages of this study, which could guide future physics-based AI approaches, are:

- Learning the implicit prior distribution from the data led to improved estimation performance compared to using fixed or assumed priors. Additionally, the learned spatiotemporal priors facilitated the reconstruction of temporally consistent patterns.
- Including the underlying physics of the inverse problem allowed using neural networks for denoising. This brings 3 improvements:
  - 1) Leveraging neural networks to impose signal structure (regularization via denoising subproblem) produced more consistent results than inverting the forward model and regularizing the solution together.
  - 2) Excluding domain transformation from the neural network training process prevents it from learning artificial patterns during the reconstruction process.
  - 3) Learning a simpler task in neural network training resulted in a network with fewer parameters, thereby reducing the risk of overfitting and the necessity for larger training datasets.

### A. Performance With Respect to the Benchmark Methods

**PULSE** outperformed all three benchmark methods in terms of the employed ECGI metrics, even **MAP** and **E2ET** that employ the same training data with **PULSE**. The results showed that once the associated prior is learned, the learned proximal operator improves the solutions given by the hand-crafted regularization priors. In **PULSE**, we break down the problem into two steps: a physics-based (data-fidelity) component and the inherent characteristics of the oracle (proximal update). Consequently, the proximal network can generalize across diverse geometries with similar cardiac source distributions, assuming an accurately generated forward model from medical scans. In addition, the auto-encoding nature of the network eliminates the

need for re-implementation to reconstruct a different number of cardiac sources. It can be trained using a multi-resolution dataset in which consistent proximity is enforced in node ordering.

When the training and validation performances of **PULSE** and **E2ET** were compared on non-normalized data, **E2ET** model achieved an MSE training loss of 1.13 and a validation loss of 4.47 (Average temporal metrics: CC = 0.99, RE = 0.15). On the other hand, **PULSE** model achieved an MSE training loss of 1.14 and a validation loss of 6 (Average temporal metrics: CC = 0.94, RE = 0.29). According to these training results, **E2ET** shows a better data fitting compared to our model. However, the test results presented in Section IV-A show that **PULSE** outperforms **E2ET** when the test and training data come from different animals, even with a perfect match between the test and training geometric models. These test results show that encoding domain transformation and regularization within a network can lead to poor generalization with limited training data for the **E2ET** model. In **PULSE**, using a physical model estimate and regularizing it with a proximal network helps identify artifacts in imperfect data. Since training is based on denoising samples, artifacts are not learned as patterns. In contrast, end-to-end training strategies learn artifacts as patterns, affecting performance during inference. Fitting an end-to-end neural network to the forward model transformation alone would yield one of many possible representations, as the forward model is highly underdetermined.

Contrary to the claims in [14] and [37], we have successfully utilized Euclidean neural networks by excluding domain transformations during training. Leveraging the decoupled problem, we employ neural networks solely as proximal operators. We enable their effective application in learning across diverse geometries through decoupling and careful architectural design. Hence, we replace the necessity for complex architectures and computationally intensive operations such as graph convolutions with splines. Instead, we employ a physics-based forward model in conjunction with lightweight neural networks that impose minimal computational overhead.

An additional factor enhancing the effectiveness of Euclidean networks in recognizing EGM patterns is the reordering algorithm. This greedy pre-processing step sorts spatially correlated signals after the mesh creation, facilitating the use of Euclidean convolutional kernels on the data for capturing local features in denoising.

### B. Robustness to Test-Training Geometry Mismatch

A potential disadvantage of **PULSE** is its tendency to capture noisy patterns originating from a specific forward model. However, we observed that the method can maintain its performance within the physically possible rotation ranges of cardiac movements inside the torso. This provides a degree of flexibility during inference when different forward models are used, as the effect of the forward model is removed from the solution by the two-step design after the model-measurement update step. **E2ET** approaches, on the other hand, learn the forward

model during the training process, making them more prone to overfitting and limited by the training data.

In principle, it is possible to learn both the regularization and the transformation with a larger neural network and avoid overfitting with weight regularization or dropout techniques. However, as suggested by Meng et al., the hybrid deep unfolding approach offers a desirable trade-off between computational time cost and accuracy compared to end-to-end training strategies [38]. This approach provides interpretability and flexibility by clearly describing optimization steps and utilizing a relatively simple neural network architecture that imposes signal structure through proximal denoising. While end-to-end training strategies capture the measurement constraint and regularization simultaneously, they result in only a slightly better estimate but with a significantly increased computational burden.

### C. Robustness to Test-Training Noise SNR Mismatch

The hybrid reconstruction scheme's training process enables the neural network to inherently assess noise levels in classical estimates. Although the noise level in measurements remains constant, the noise in estimates varies based on the forward model's noise-filtering capability. Consequently, the neural network learns to perform appropriate proximal updates, allowing it to function under varying noise levels. In contrast, end-to-end trained models lack an explicit mechanism to prevent them from learning and encoding measurement noise as reconstruction features. Our results show that this approach effectively filters out noise from 20 dB SNR torso measurements, reflecting a practical scenario where measurements undergo several stages of noise filtering. Furthermore, the model generalizes to different noise levels due to prior learning via its proximal operator.

### D. Limitations and Future Work

Despite demonstrating the strengths of the proposed physics-based neural network approach, this study has some limitations.

We only used simulated BSPs from experimental EGMs with single-site ventricularly paced data mimicking PVCs as the measurement data. This dataset included data from 7 experiments (3 for training and 4 for testing), and model variations between the training and test sets were introduced in terms of heart rotations in the test data geometry and measurement noise SNR. As such, the Utah dataset captures patient variety (in terms of priors with different parameters) and maximizes electrophysiological diversity to some extent. However, the model's performance with different experimental or clinical data on drastically different geometries and a variety of arrhythmias needs further evaluation to test generalization across subjects. The proposed network's architecture supports multi-resolution learning due to its auto-encoding nature, thus enabling these evaluations with a larger dataset.

This study tested the effects of geometric variations between the training and testing data geometries to evaluate how the NN-based proximal operator generalizes across different heart-torso pairs. However, these results were obtained using an accurately generated forward model for the test data and do not include robustness tests against measurement-model mismatches, where the forward model is not precisely quantified. Future work will

expand the variety of test data and heart-torso geometries to enhance the evaluation of the proposed method's performance against such geometric model mismatches.

As a future study, we will also compare our approach with GNNs. With limited data, they yield less sensitivity to the modeled geometric uncertainty. However, more experiments have to be done with large-scale data while also considering the computational burden and the need for defining specific task-based operations on graphs.

## VI. CONCLUSION

In this study, we propose a modular and lightweight framework for accurate electrogram reconstructions. Unlike recent studies in ECGI, we adopt a hybrid approach that maximizes available information while avoiding the need for training large and complex neural networks. Our approach is highly scalable, requiring no data preprocessing and enabling direct implementation on any dataset. Additionally, we design a compact neural network with predefined layers, eliminating the need for task-specific layers or domain-specific architectures. Notably, Euclidean networks demonstrate effectiveness on irregular heart meshes when proximity is enforced through proper node ordering. Importantly, we find that this improvement remains consistent even under geometric variations, such as rotation. Overall, our approach enhances both reconstruction accuracy and pacing site localization robustness. We envision seamless integration of our method into clinical ECGI pipelines to enhance PVC localization.

## REFERENCES

- [1] E. Ozkoc et al., "Prior model selection in Bayesian MAP estimation-based ECG reconstruction," in *Proc. 13th IEEE Int. Conf. Meas.*, 2021, pp. 142–145.
- [2] Y. S. Dogrusoz, "Statistical estimation applied to electrocardiographic imaging," in *Proc. 12th IEEE Int. Conf. Meas.*, 2019, pp. 2–9.
- [3] O. Onak and Y. S. Dogrusoz, "Performance of Tikhonov regularization with different constraints in electrocardiographic imaging," in *Proc. 28th IEEE Signal Process. Commun. Appl. Conf.*, 2020, pp. 1–4.
- [4] P. C. Hansen, *Rank-Deficient and Discrete ILL-Posed Problems: Numerical Aspects of Linear Inversion*. Philadelphia, PA, USA: SIAM, 1998.
- [5] G. Shou et al., "Truncated total least squares: A new regularization method for the solution of ECG inverse problems," *IEEE Trans. Biomed. Eng.*, vol. 55, no. 4, pp. 1327–1335, Apr. 2008.
- [6] M. Jiang et al., "Combination of the LSQR method and a genetic algorithm for solving the electrocardiography inverse problem," *Phys. Med. Biol.*, vol. 52, no. 5, 2007, Art. no. 1277, doi: 10.1088/0031-9155/52/5/005.
- [7] X. Jiang et al., "Learning geometry-dependent and physics-based inverse image reconstruction," in *Proc. Int. Conf. Med. Image Comput. Comput. Assist. Interv.*, 2020, pp. 487–496.
- [8] T. Bacoyannis et al., "Deep learning formulation of electrocardiographic imaging integrating image and signal information with data-driven regularization," *EP Europace*, vol. 23, no. Supplement\_1, pp. i55–i62, 2021.
- [9] R. Tenderini et al., "PDE-Aware deep learning for inverse problems in cardiac electrophysiology," *SIAM J. Sci. Comput.*, vol. 44, no. 3, pp. B605–B639, 2022. [Online]. Available: <https://pubs.siam.org/doi/abs/10.1137/21M1438529>
- [10] N. Pilia et al., "Non-invasive localization of the ventricular excitation origin without patient-specific geometries using deep learning," *Artif. Intell. Med.*, vol. 143, 2023, Art. no. 102619. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S0933365723001331>
- [11] S. Schuler et al., "Cobiveco: Consistent biventricular coordinates for precise and intuitive description of position in the heart—with MATLAB implementation," *Med. Image Anal.*, vol. 74, 2021, Art. no. 102247. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S1361841521002929>

- [12] K.-W. Chen, L. Bear, and C.-W. Lin, "Solving inverse electrocardiographic mapping using machine learning and deep learning frameworks," *Sensors*, vol. 22, no. 6, 2022, Art. no. 2331. [Online]. Available: <https://www.mdpi.com/1424-8220/22/6/2331>
- [13] X. Jiang et al., "Uncertainty quantification of cardiac position on deep graph network ECGI," in *Proc. IEEE Comput. Cardiol.*, 2022, vol. 498, pp. 1–4.
- [14] X. Jiang et al., "Improving generalization by learning geometry-dependent and physics-based reconstruction of image sequences," *IEEE Trans. Med. Imag.*, vol. 42, no. 2, pp. 403–415, Feb. 2023.
- [15] K. Zhang et al., "Learning deep CNN denoiser prior for image restoration," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2017, pp. 3929–3938. [Online]. Available: [https://openaccess.thecvf.com/content\\_cvpr\\_2017/html/Zhang\\_Learning\\_Deep\\_CNN\\_CVPR\\_2017\\_paper.html](https://openaccess.thecvf.com/content_cvpr_2017/html/Zhang_Learning_Deep_CNN_CVPR_2017_paper.html)
- [16] Z. Fang, S. Buchanan, and J. Sulam, "What's in a Prior? Learned proximal networks for inverse problems," in *Proc. 12th Int. Conf. Learn. Representations*, 2024.
- [17] Y. Serinagaoglu, D. H. Brooks, and R. S. MacLeod, "Bayesian solutions and performance analysis in bioelectric inverse problems," *IEEE Trans. Biomed. Eng.*, vol. 52, no. 6, pp. 1009–1020, Jun. 2005.
- [18] Y. Serinagaoglu, D. H. Brooks, and R. S. MacLeod, "Improved performance of Bayesian solutions for inverse electrocardiography using multiple information sources," *IEEE Trans. Biomed. Eng.*, vol. 53, no. 10, pp. 2024–2034, Oct. 2006.
- [19] O. Bouhamama et al., "A patchwork method to improve the performance of current methods for solving the inverse problem of electrocardiography," *IEEE Trans. Biomed. Eng.*, vol. 70, no. 1, pp. 55–66, Jan. 2023.
- [20] M. Zheng et al., "A hybrid CNN for image denoising," *J. Artif. Intell. Technol.*, vol. 2, no. 3, pp. 93–99, 2022. [Online]. Available: <https://ojs.istp-press.com/jait/article/view/101>
- [21] A. E. Ilesanmi and T. O. Ilesanmi, "Methods for image denoising using convolutional neural network: A review," *Complex Intell. Syst.*, vol. 7, no. 5, pp. 2179–2198, 2021, doi: [10.1007/s40747-021-00428-4](https://doi.org/10.1007/s40747-021-00428-4).
- [22] C. Tian et al., "Enhanced CNN for image denoising," *CAAI Trans. Intell. Technol.*, vol. 4, no. 1, pp. 17–23, 2019. [Online]. Available: <https://onlinelibrary.wiley.com/doi/abs/10.1049/trit.2018.1054>
- [23] K. Zhang et al., "Beyond a Gaussian denoiser: Residual learning of deep CNN for image denoising," *IEEE Trans. Image Process.*, vol. 26, no. 7, pp. 3142–3155, Jul. 2017.
- [24] K. Zhang et al., "Plug-and-play image restoration with deep denoiser prior," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 44, no. 10, pp. 6360–6376, Oct. 2022.
- [25] J. Fu et al., "Rotation equivariant proximal operator for deep unfolding methods in image restoration," *IEEE Tran. Pattern Anal. Mach. Intell.*, vol. 46, no. 10, pp. 6577–6593, Oct. 2024.
- [26] P. Lei et al., "Decomposition-based variational network for multi-contrast mri super-resolution and reconstruction," in *Proc. IEEE/CVF Int. Conf. Comput. Vis.*, 2023, pp. 21296–21306.
- [27] A. N. Tikhonov and V. I. A. Arsenin, *Solutions of ILL-Posed Problems*. Ultimo, NSW, Australia: Halsted Press, 1977.
- [28] P. C. Hansen, "Analysis of discrete ILL-posed problems by means of the L-curve," *SIAM Rev.*, vol. 34, no. 4, pp. 561–580, 1992. [Online]. Available: <http://epubs.siam.org/doi/10.1137/1034115>
- [29] R. MacLeod, B. Taccardi, and R. Lux, "Electrocardiographic mapping in a realistic torso tank preparation," in *Proc. 17th IEEE Int. Conf. Eng. Med. Biol. Soc.*, 1995, vol. 1, pp. 245–246.
- [30] F. Aldemir and Y. S. Dogrusoz, "Compensation of model errors in electrocardiographic imaging using Bayesian estimation," in *Proc. IEEE Comput. Cardiol.*, 2021, vol. 48, pp. 1–4.
- [31] R. C. Barr, M. Ramsey, and M. S. Spach, "Relating epicardial to body surface potential distributions by means of transfer coefficients based on geometry measurements," *IEEE Trans. Biomed. Eng.*, vol. BME-24, no. 1, pp. 1–11, Jan. 1977.
- [32] P. C. Stanley, T. C. Pilkington, and M. N. Morrow, "The effects of thoracic inhomogeneities on the relationship between epicardial and torso potentials," *IEEE Trans. Biomed. Eng.*, vol. BME-33, no. 3, pp. 273–284, Mar. 1986.
- [33] R. L. Lux et al., "Limited lead selection for estimation of body surface potential maps in electrocardiography," *IEEE Trans. Biomed. Eng.*, vol. BME-25, no. 3, pp. 270–276, May 1978.
- [34] A. Wirgin, "The inverse crime," 2004. *arXiv preprint math-ph/0401050*.
- [35] K. Uğurlu, "Improvement of electrocardiographic imaging reconstructions: A physics-guided AI approach and an efficient method for training data reduction," M.S. thesis, Middle East Tech. Univ., Çankaya/Ankara, Türkiye, 2023.
- [36] J. Duchateau, M. Potse, and R. Dubois, "Spatially coherent activation maps for electrocardiographic imaging," *IEEE Trans. Biomed. Eng.*, vol. 64, no. 5, pp. 1149–1156, May 2017.
- [37] S. Ghimire et al., "Improving generalization of deep networks for inverse reconstruction of image sequences," in *Proc. Conf. Inf. Process. Med. Imag.*, 2019, pp. 153–166.
- [38] Z. Meng, X. Yuan, and S. Jalali, "Deep unfolding for snapshot compressive imaging," *Int. J. Comput. Vis.*, vol. 131, no. 11, pp. 2933–2958, 2023.