

MOX-Report No. 26/2026

Learning geometry-dependent lead-field operators for forward ECG modeling

Dokuchaev, A.; Bonizzoni, F.; Pagani, S.; Regazzoni, F.; Pezzuto, S.

MOX, Dipartimento di Matematica
Politecnico di Milano, Via Bonardi 9 - 20133 Milano (Italy)

mox-dmat@polimi.it

<https://mox.polimi.it>

Learning geometry-dependent lead-field operators for forward ECG modeling

Arsenii Dokuchaev^{*1}, Francesca Bonizzoni², Stefano Pagani², Francesco Regazzoni², and Simone Pezzuto^{1,3}

¹Laboratory of Mathematics for Biology and Medicine, Università di Trento, Italy

²MOX Laboratory, Department of Mathematics, Politecnico di Milano, Italy

³Euler Institute, Università della Svizzera italiana, Switzerland

Abstract

Modern forward electrocardiogram (ECG) computational models rely on an accurate representation of the torso domain. The lead-field method enables fast ECG simulations while preserving full geometric fidelity. Achieving high anatomical accuracy in torso representation is, however, challenging in clinical practice, as imaging protocols are typically focused on the heart and often do not include the entire torso. In addition, the computational cost of the lead-field method scales linearly with the number of electrodes, limiting its applicability in high-density recording settings. To date, no existing approach simultaneously achieves high anatomical fidelity, low data requirements and computational efficiency. In this work, we propose a shape-informed surrogate model of the lead-field operator that serves as a drop-in replacement for the full-order model in forward ECG simulations. The proposed framework consists of two components: a geometry-encoding module that maps anatomical shapes into a low-dimensional latent space, and a geometry-conditioned neural surrogate that predicts lead-field gradients from spatial coordinates, electrode positions and latent codes. The proposed method achieves high accuracy in approximating lead fields both within the torso (mean angular error $<5^\circ$) and inside the heart, resulting in highly accurate ECG simulations (relative mean squared error $<2.5\%$). The surrogate consistently outperforms the widely used pseudo lead-field approximation while preserving negligible inference cost. Owing to its compact latent representation, the method does not require a fully detailed torso segmentation and can therefore be deployed in data-limited settings while preserving high-fidelity ECG simulations.

1 Introduction

In cardiology, electrocardiography records electric potentials originating from the heart over the body surface. A classic example is the 12-lead surface electrocardiogram (ECG or EKG), in which potentials are recorded at standardized chest locations, comprising ten electrode positions, including the ground (reference) electrode. Body Surface Potential Maps (BSPMs) are a spatially denser version of ECG measurements and are typically employed to solve the inverse problem of electrocardiography, which enables non-invasive reconstruction of cardiac potentials or sources [12, 35, 20].

From a modeling perspective, the translation of cardiac sources into an ECG or BSPM is typically achieved using the lead-field method [32]. Lead fields form the basis of efficient forward-modeling workflows in cardiac electrophysiology as well as of the associated inverse problem [30, 16]. Mathematically, the lead field for a single lead corresponds to the solution of the adjoint pseudo-bidomain problem associated with a unit current injection at the electrode pair and is closely related to the Green’s function of the pseudo-bidomain operator. Therefore, it strongly depends on heart shape, heart orientation and position within the torso, electrode locations, and overall anatomy around the heart [50].

Setting up and solving the lead-field problem is, however, non-trivial for several reasons. First, in standard cardiac imaging, high-quality clinical images for torso segmentation and anatomy preparation are often incomplete or missing, because imaging primarily targets the heart. This issue can be partially circumvented by “implanting” patient-specific cardiac anatomy into a template torso model, for instance from a previously segmented case [47, 34], or by using statistical shape models conditioned on available

*Corresponding author: arsenii.dokuchaev@unitn.it

data [21]. Notably, electrode positions must also be segmented to be included in the model; however, electrodes are generally removed during imaging, so precise localization is often difficult [5].

Second, the overall assembly cost of the transfer operator is computationally non-negligible and scales linearly with the number of electrodes (or, more precisely, the number of leads). Computing the lead-field transfer operator requires solving an elliptic partial differential equation (PDE) over the whole-body domain for each independent lead configuration, accounting for electrical conductivities (possibly anisotropic), the shapes of major organs and cavities, and electrode locations [32]. The lead-field problem therefore requires careful segmentation of the entire torso (heart, lungs, major blood cavities, liver, fat, ribs) for electrical conductivities [18] and accurate placement of electrodes on the torso surface.

A typical strategy to circumvent both issues is the use of approximate anatomical models, which relax imaging requirements. The pseudo lead-field approach only requires electrode positions relative to the heart and has a closed-form solution that is very fast to compute [7]. However, torso anatomy is known to affect lead fields (and therefore simulated ECGs), especially for electrodes close to the heart, such as precordial leads. Intermediate approaches exist and are based on surface meshes of the body, heart, and possibly lungs. They usually employ the boundary element method, which does not require meshing the three-dimensional volume. However, this approach cannot be trivially extended to support anisotropy and may suffer from numerical instabilities due to the presence of singular kernels in the boundary integral formulation and the resulting ill-conditioning of the system matrices [48]. Therefore, existing approach cannot guarantee high anatomical fidelity while maintaining low data requirements and computational efficiency.

In this work, we propose a methodology to approximate the lead-field operator efficiently while retaining high accuracy for practical ECG applications. Importantly, the proposed approach does not rely on a detailed volumetric torso segmentation, but instead leverages a compact geometric representation that remains applicable in limited-data settings. Our approach is inspired by neural field methods (or implicit neural representation), which parameterize continuous physical fields defined over general N -dimensional domains using neural networks. Such methods have been widely adopted in 2-D and 3-D image synthesis, reconstruction and rendering tasks [49]. For example, DeepSDF learns a continuous signed distance function (SDF) representation of shapes, enabling reconstruction, interpolation, and completion from partial or noisy 3-D data [29]. Similar neural field approaches have proven effective for the parameterization and reconstruction of cardiac geometries [38]. Specifically, our method has two components: one for geometric encoding and one for lead-field prediction. The first is a DeepSDF model [29, 45], based on an auto-decoder neural network that predicts the Signed Distance Function (SDF) of a shape for a given latent code. During training, the network learns a joint latent representation of torso and heart shapes. We also consider, for comparison, a more standard geometric representation based on Principal Component Analysis (PCA), where the latent representation is the linear subspace spanned by the principal components. The second component is a neural implicit representation of the lead-field function, conditioned on the geometry latent code (either DeepSDF- or PCA-based) and electrode coordinates [26, 36]. The approach is fully data-driven in both components and does not explicitly encode physical constraints in the model architecture or loss formulation. Training requires only a dataset of precomputed lead-field solutions, which are here generated using a high-fidelity torso model for multiple heart-torso configurations. Finally, we compare ECGs computed with this approach against ECGs based on the pseudo lead field.

2 Methods

This section presents the mathematical and computational methods used for preprocessing geometric models, generating training and test samples, and designing and training the proposed neural network architectures. We begin by introducing the mathematical formulation of the lead-field problem and the corresponding forward ECG computation. We then describe the geometric pipeline, including the generation of joint heart-torso anatomies and the sampling of training and test point clouds. Next, we detail the numerical computation of ground-truth lead fields via finite element approximation of the forward problem and the architecture and training of the proposed neural networks. Finally, we provide a description of the pseudo lead-field formulation used for comparison.

An overview of the complete pipeline is shown in Fig. 1. The workflow consists of: (1) generating joint heart-torso geometries and sampling training and test point clouds; (2) defining electrode configurations; (3) computing reference lead fields by solving the forward problem using the finite element method (FEM); (4) training a neural network to learn a compact anatomical representation via DeepSDF; and

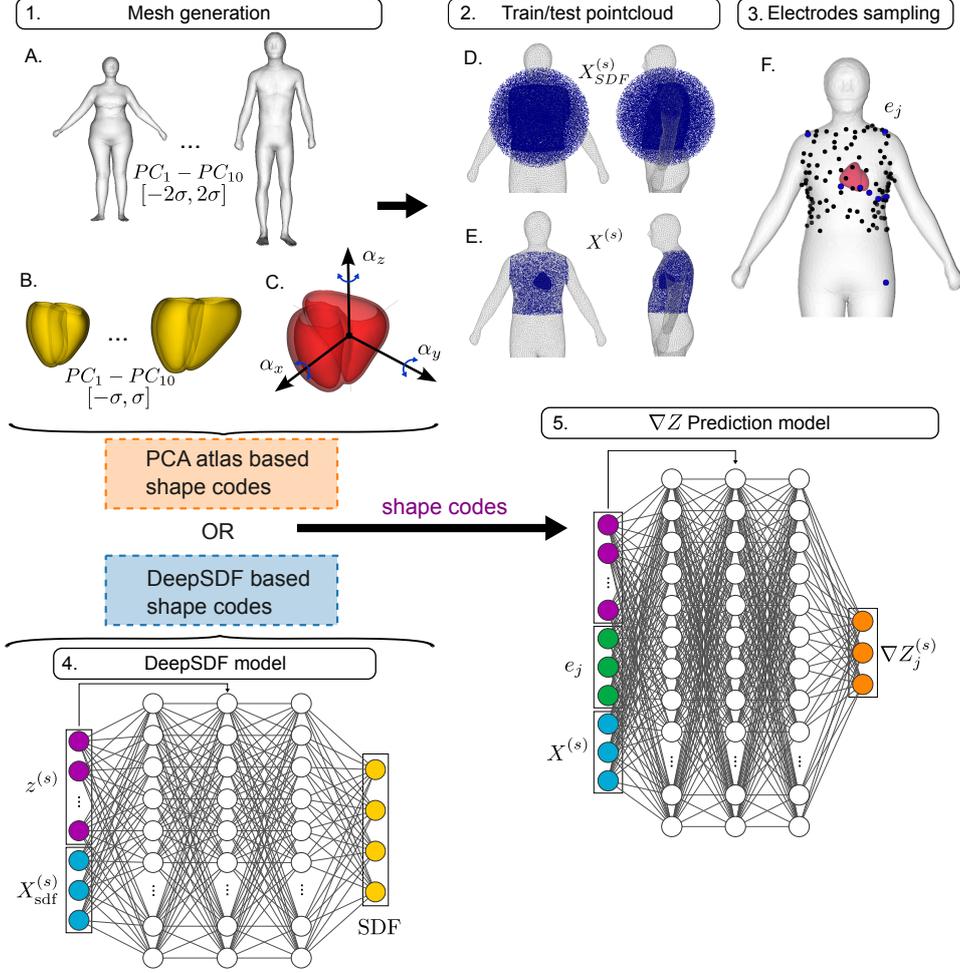


Figure 1: Graphical representation of the developed pipeline. (1.) Mesh generation stage: we used torso and heart models from a statistical PCA atlas. To create joint models, we varied (A.) the first 10 principal components of the torso, (B.) the first 10 principal components of the heart, and (C.) heart rotation angles along anatomical axes: α_x along the LV–RV axis, α_y along the anterior-posterior axis, and α_z along the LV long axis. As a result, a set of geometries was obtained and used to generate training and test point clouds (2.), while a set of features (10 torso modes + 10 heart modes + 3 angles) was used as feature vectors describing each dataset (highlighted with an orange frame). (2.) Train/test point-cloud generation. (D.) Point cloud around and inside the torso used to train/test the DeepSDF model. (E.) Point cloud inside the torso used to train/test the ∇Z prediction model. (3.) Variation of unipolar electrode location on the anterior surface of the torso. (4.) Schematic representation of DeepSDF model used for generation of DeepSDF based shape codes (highlighted with a blue frame). (5.) Schematic representation of lead field gradient ∇Z prediction model.

(5) training a second neural network to predict lead-field gradients from spatial coordinates, electrode positions, and the learned latent codes.

The surrogate model is subsequently in-silico validated by comparing predicted lead-field gradients and the resulting ECGs against FEM-based reference solutions for previously unseen geometries and activation patterns.

2.1 Forward ECG modeling and lead-field problem

Let denote by $\Omega \subset \mathbb{R}^3$ the body domain and by $\Sigma = \partial\Omega$ the body surface; the active myocardium of the heart is $\Omega_H \subset \Omega$, while the passive conductor tissue is $\Omega_0 := \Omega \setminus \bar{\Omega}_H$ and $\Gamma = \partial\Omega_H$ is the heart-torso surface. Given the transmembrane potential $V_m(\mathbf{x}, t)$ on Ω_H , the extracellular potential $\phi_0(\mathbf{x}, t)$ and

$\phi_e(\mathbf{x}, t)$, respectively in the passive and active tissue, solves the following pseudo-bidomain equation:

$$\left\{ \begin{array}{ll} -\operatorname{div}(\mathbf{G}\nabla\phi_e) = \operatorname{div}(\mathbf{G}_i\nabla V_m), & \mathbf{x} \in \Omega_H, t \in \mathbb{R}, \\ -\operatorname{div}(\mathbf{G}_0\nabla\phi_0) = 0, & \mathbf{x} \in \Omega_0, t \in \mathbb{R}, \\ \phi_e - \phi_0 = 0, & \mathbf{x} \in \Gamma, t \in \mathbb{R}, \\ \mathbf{G}_0\nabla\phi_0 \cdot \mathbf{n} - \mathbf{G}\nabla\phi_e \cdot \mathbf{n} = \mathbf{G}_i\nabla V_m \cdot \mathbf{n}, & \mathbf{x} \in \Gamma, t \in \mathbb{R}, \\ \mathbf{G}_0\nabla\phi_0 \cdot \mathbf{n} = 0, & \mathbf{x} \in \Sigma, t \in \mathbb{R}, \end{array} \right. \quad (1)$$

where \mathbf{G}_i , \mathbf{G} and \mathbf{G}_0 are respectively the intra-cellular, bulk and torso conductivity tensors, and \mathbf{n} is the outwards normal with respect to the heart-torso surface. In general, the bulk conductivity is $\mathbf{G} = \mathbf{G}_i + \mathbf{G}_e$, where \mathbf{G}_e is the extra-cellular conductivity. The conductivity tensors encode the anisotropy of the tissue, especially for the myocardium. In what follows, we are assuming:

$$\begin{aligned} \mathbf{G}_i &= \sigma_{i,t}\mathbf{I} + (\sigma_{i,f} - \sigma_{i,t})\mathbf{f} \otimes \mathbf{f}, \\ \mathbf{G}_e &= \sigma_{e,t}\mathbf{I} + (\sigma_{e,f} - \sigma_{e,t})\mathbf{f} \otimes \mathbf{f}, \\ \mathbf{G}_0 &= \sigma_0\mathbf{I}, \end{aligned}$$

with $\mathbf{f}(\mathbf{x})$ the local fiber direction in the myocardium, and $\sigma_{i,t}$, $\sigma_{e,t}$, $\sigma_{i,f}$ and $\sigma_{e,f}$ respectively the intra- and extra-cellular fiber and transverse electric conductivities, and σ_0 the torso conductivity. The values used in our simulations are based on [33], specifically: $\sigma_{i,t} = 0.3 \text{ mS cm}^{-1}$, $\sigma_{i,f} = \sigma_{e,f} = 3.0 \text{ mS cm}^{-1}$ and $\sigma_{e,t} = 1.2 \text{ mS cm}^{-1}$. For the torso we considered $\sigma_0 = 0.6 \text{ mS cm}^{-1}$, according to [6].

The solution to the system (1) is defined up to a constant, which can be fixed by setting a reference node or, equivalently, by constraining the solution to have zero average on the body surface:

$$\int_{\Sigma} \phi_0 \, d\mathbf{x} = 0. \quad (2)$$

The ECG is generally formed by a set of leads, each obtained from a linear combination of $\phi_0(\mathbf{x}, t)$ measured at some electrodes locations $\mathbf{e}_j \in \Sigma$. (Under suitable assumptions on the domain, coefficients and V_m , it is possible to prove that the potential is globally $\mathcal{C}(\bar{\Omega})$, thus point-wise evaluation is well-defined. See [13, Thm. 5.1].) For each lead $\ell = 1, \dots, L$, the single lead ECG reads

$$V_\ell(t) = \sum_{j=1}^{n_\ell} \alpha_j \phi_0(\mathbf{e}_j, t),$$

where n_ℓ is number of electrodes for lead ℓ , and α_k are the zero-sum weights. Note that the evaluation of the ECG requires the numerical solution of (1) for each time step of the transmembrane potential.

The lead field method is a natural way to alleviate the computational cost of evaluating the ECG [14, 30, 32]. In fact, it is possible to prove that the solution with zero reference potential of (1) is given by [13, Prop. 5.1]

$$\phi_0(\mathbf{e}_j, t) = - \int_{\Omega_H} \mathbf{G}_i \nabla V_m(\mathbf{x}, t) \cdot \nabla Z(\mathbf{x}; \mathbf{e}_j) \, d\mathbf{x}, \quad (3)$$

where $Z(\mathbf{x}; \mathbf{e}_j)$ is the *lead-field function* with respect to the electrode \mathbf{e}_j , and solves the problem:

$$\left\{ \begin{array}{ll} -\operatorname{div}(\mathbf{G}\nabla Z) = 0, & \mathbf{x} \in \Omega_H, \\ -\operatorname{div}(\mathbf{G}_0\nabla Z_0) = 0, & \mathbf{x} \in \Omega_0, \\ Z - Z_0 = 0, & \mathbf{x} \in \Gamma, \\ \mathbf{G}_0\nabla Z_0 \cdot \mathbf{n} - \mathbf{G}\nabla Z \cdot \mathbf{n} = 0, & \mathbf{x} \in \Gamma, \\ \mathbf{G}_0\nabla Z_0 \cdot \mathbf{n} = \delta_{\mathbf{e}_j}(\mathbf{x}) - |\Sigma|^{-1}, & \mathbf{x} \in \Sigma, \end{array} \right. \quad (4)$$

where $\delta_{\mathbf{e}_j}(\mathbf{x})$ is the Dirac's delta centered at \mathbf{e}_j . Note that this problem does not depend neither on time nor the transmembrane potential. An example of the (gradient of the) lead-field function computed from a single electrode is provided in Fig. 2. With a slight abuse of notation, we will denote simply by $Z(\mathbf{x})$ the lead field in the whole body, without making a distinction between $Z(\mathbf{x})$ and $Z_0(\mathbf{x})$.

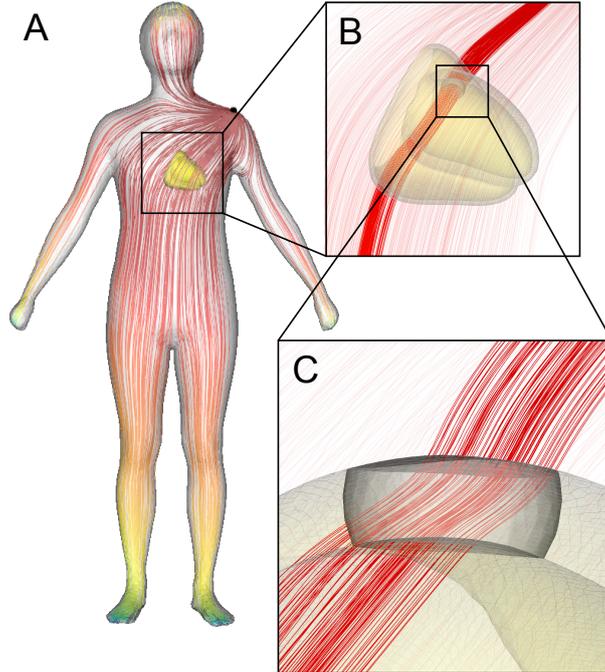


Figure 2: Computed lead field for a unipolar lead (left shoulder), visualized using streamlines representing the direction of ∇Z . (A.) Streamlines shown throughout the entire torso. (B.) Close-up view of the heart region (the heart is shown in yellow); streamlines passing through the free wall of the left ventricle (LV) are highlighted in bold red. (C.) Enlarged view of the LV free wall (represented as a cylinder) with the corresponding gradient streamlines highlighted in bold. Note the change in the slope of the streamlines at the heart-torso interface, reflecting the change in direction of ∇Z across the boundary.

2.2 Data preparation

We used the bi-ventricular statistical shape model (SSM) proposed in [3] as the baseline heart model. This statistical atlas was constructed from MR images of 1093 healthy subjects and provides 100 principal components with associated variances describing the joint distribution of the left ventricular myocardium (LV) and right ventricular blood pool (RV) surfaces. While the original atlas contained separate surface representations, we adopted a modified version described in [40].

To construct a dataset of 100 heart geometries, we varied the weights of the first 10 principal components within the range $[-1, 1]$ standard deviations using Latin hypercube sampling with a uniform distribution. A uniform distribution was chosen instead of a Gaussian one to ensure broad and even coverage of the shape space, including geometries located near the boundaries of the selected range. Latin hypercube sampling guarantees uniform coverage of the resulting 10-dimensional parameter space (see Fig. 3).

For each heart geometry, myocardial fiber orientations were assigned using the rule-based method of [4], with fiber angles $\alpha_{\text{endo}} = 40^\circ$, $\alpha_{\text{epi}} = -50^\circ$, and sheet angles $\beta_{\text{endo}} = -65^\circ$, $\beta_{\text{epi}} = 25^\circ$.

Because the lead field depends not only on electrode configuration but also on the geometry of the computational domain [32], we also varied torso shape. Torso geometries were generated from the MPII Human Shape statistical atlas [31], using a sampling strategy analogous to that employed for the heart.

Even among healthy subjects, the size, position, and orientation of the heart significantly influence the ECG, as demonstrated in clinical studies [17] and computational investigations [27, 24, 50]. To account for this variability, we additionally varied heart translation and rotation within the torso. The baseline heart position was defined by aligning the heart centroid with a reference point inside the torso volume. The baseline orientation was specified by three rotation angles α_X , α_Y , and α_Z describing rotations of the cardiac long axis relative to the anatomical planes, following [28] (see Fig. 1C).

Three-dimensional finite element meshes were generated using Gmsh [15]. The characteristic element sizes were 0.8 mm for the heart and 11.8 mm for the torso. The use of statistical atlases significantly simplified and automated geometry processing, including the identification of subendocardial surfaces, base and apex regions, and other anatomical landmarks required to compute Laplace-Dirichlet scalar fields for myocardial fiber assignment [4]. In addition, atlas-based modeling enabled standardized placement of

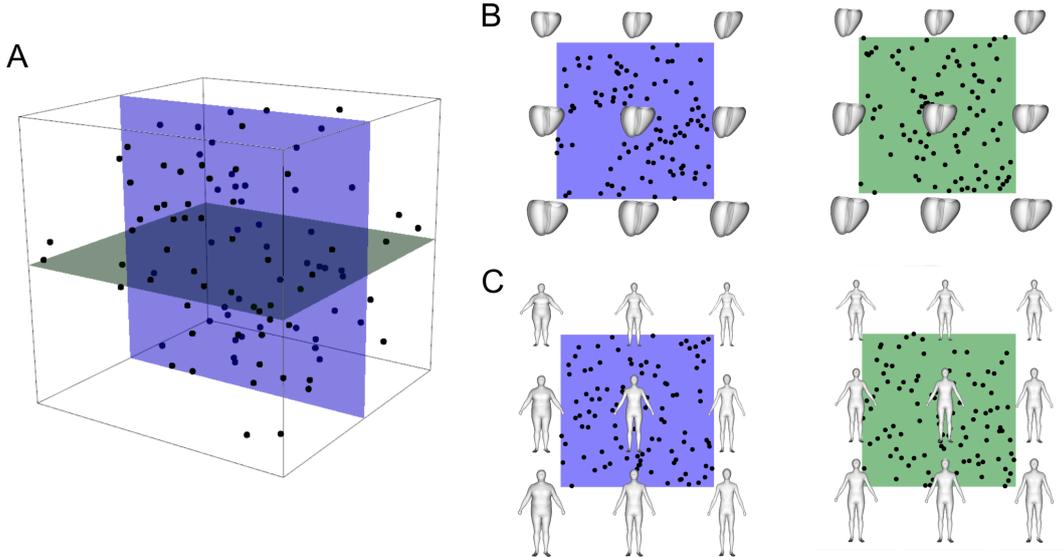


Figure 3: Design of training samples for parametric models of the torso and heart. **A** Projection of a 10-dimensional parametric space onto a 3D cube, obtained by neglecting the last 7 dimensions. Black dots indicate the position of 100 sampling points selected using Latin hypercube sampling from a uniform distribution. **B** and **C** are cross-sections of the parametric space for the heart (**B**) and torso (**C**) showing the variations of the three principal components: PC1, PC2 in blue colours, PC1, PC3 in green colours.

electrodes on the torso surface. On the anterior torso surface of each mesh, 100 unipolar electrodes were uniformly sampled. In addition, the nine independent electrodes of a standard 12-lead ECG configuration were placed for each geometry (see Fig. 1, block 3).

We employed a fixed train-test split at the level of virtual patients, ensuring that no samples from a test geometry were used during training.

2.3 Lead-field neural surrogate

Since the ECG signal $V(t)$ depends on the lead-field gradient ∇Z within the heart region Ω_H , we focus on parameterizing and reconstructing the gradient field $\nabla Z(\mathbf{x})$ throughout the domain. Note that in general ∇Z has a jump discontinuity on the heart-torso interface, due to the different conductivity of the heart and its surrounding tissue.

The problem of learning the lead-field gradient conditioned on anatomical variability can be formulated as approximating a conditional vector-valued function

$$\mathcal{NN}_{\text{LF}}: \mathbb{R}^3 \times \mathbb{R}^3 \times \mathbb{R}^{d_z} \rightarrow \mathbb{R}^3,$$

which maps a spatial point $\mathbf{x} \in \Omega^{(p)}$, an electrode position $\mathbf{e}_j \in \Sigma^{(p)}$, and a latent anatomical code $\mathbf{z}^{(p)} \in \mathbb{R}^{d_z}$ associated with geometry (p) to the corresponding lead-field gradient:

$$\mathcal{NN}_{\text{LF}}(\mathbf{x}, \mathbf{e}_j, \mathbf{z}^{(p)}; \boldsymbol{\theta}_{\text{LF}}) \approx \nabla Z_j^{(p)}(\mathbf{x}),$$

where $Z_j^{(p)}$ is the lead field (with respect to electrode \mathbf{e}_j) for the domain $\Omega^{(p)}$. The latent code will encode for the domain $\Omega^{(p)}$, which includes the heart and other organs, and the electric conductivities, including the fiber direction. The weights of the neural networks are denoted by $\boldsymbol{\theta}_{\text{LF}}$.

We model \mathcal{NN}_{LF} using a fully connected neural network (see Fig. 1, block 5). For each virtual patient (p), we randomly sampled 2^{16} spatial points $\mathbf{x} \in \Omega^{(p)}$ within the torso domain and computed the corresponding gradients $\nabla Z_j^{(p)}(\mathbf{x})$ for each electrode $\mathbf{e}_j^{(p)}$. Sampling was biased toward the torso surface and the heart-torso interface (Fig. 2), where the gradient exhibits sharp spatial variations: approximately 80% of the sampled points lie within 10 mm of these interfaces. The set of sampled points $\mathbf{X}^{(p)}$ differs across patients but is shared among all electrodes of a given patient.

To encode the joint heart-torso geometry, we investigated two strategies:

- A **PCA-based encoding**, with latent anatomical code as:

$$\mathbf{z}^{(p)} = [w_1^{\text{heart}}, \dots, w_{N_h}^{\text{heart}}, w_1^{\text{torso}}, \dots, w_{N_t}^{\text{torso}}, \alpha_x, \alpha_y, \alpha_z]^T,$$

where $\{w_i^{\text{heart}}\}$ and $\{w_i^{\text{torso}}\}$ are the principal component (PC) weights for the heart and torso shapes, and α_x (left-right), α_y (anterior-posterior), and α_z (apex-base) are the heart rotation angles along standard anatomical axes (see Fig. 1, block 1).

- A **DeepSDF-based encoding**, with latent anatomical code corresponding to the learned shape codes obtained from the DeepSDF model (see Sec. 2.5).

In addition to spatial coordinates and latent codes, the electrode position $\mathbf{e}_j^{(p)}$ was included as part of the network input. To enhance generalization across different anatomies, electrode locations were expressed in a normalized torso coordinate system $\tilde{\mathbf{e}}_j = (x_j, y_j, z_j)$, where $x_j \in [-1, 1]$ (left-right direction), $y_j \in [0, 1]$ (posterior-anterior direction), and $z_j \in [0, 1]$ (superior-inferior direction).

Architecturally, the neural network comprises five hidden layers with 256 neurons each and ReLU activation functions. In the third hidden layer, the latent code and electrode coordinates are concatenated to the intermediate feature representation. The final layer is linear with three outputs corresponding to the x , y and z components of the gradient vector.

Training was performed using the ADAM optimizer [19] to minimize the loss function

$$\mathcal{L}(\boldsymbol{\theta}_{\text{LF}}) = \frac{1}{N} \frac{1}{N_{\text{ele}}} \sum_{p=1}^N \sum_{j=1}^{N_{\text{ele}}} \mathbb{E}_{\mathbf{x} \sim \mathbf{X}^{(p)}} \left[\mathcal{L}_{\text{MSE}}(\mathbf{x}, \mathbf{e}_j, \mathbf{z}^{(p)}, \boldsymbol{\theta}_{\text{SDF}}) + \lambda_{\text{cos}} \mathcal{L}_{\text{cos}}(f_{\theta}(\mathbf{x}, \mathbf{e}_j^{(p)}, \mathbf{z}^{(p)}), \nabla Z_j^{(p)}(\mathbf{x})) \right], \quad (5)$$

where

$$\mathcal{L}_{\text{MSE}} = \|\mathcal{N}_{\text{LF}}(\mathbf{x}, \mathbf{e}_j, \mathbf{z}^{(p)}; \boldsymbol{\theta}_{\text{SDF}}) - \nabla Z_j^{(p)}(\mathbf{x})\|_2^2$$

and \mathcal{L}_{cos} is a cosine similarity loss, introduced to improve the fitting of the gradient direction when its magnitude is small (and thus not captured by the first term). The cosine similarity is defined for vectors $\hat{\mathbf{v}}$ and \mathbf{v} as

$$\mathcal{L}_{\text{cos}}(\hat{\mathbf{v}}, \mathbf{v}) = 1 - \frac{\hat{\mathbf{v}} \cdot \mathbf{v}}{\|\hat{\mathbf{v}}\|_2 \|\mathbf{v}\|_2 + \varepsilon}.$$

To better capture high-frequency spatial variations, particularly near the heart-torso interface, we additionally employed Fourier features [43] as a positional encoding mechanism.

2.4 PCA-based geometry encoding

As a low-dimensional representation of the joint heart-torso geometries, we used the same parameter vectors that were employed to generate the shapes. Specifically, these include the weights of the first 10 principal components of the torso, the weights of the first 10 principal components of the heart geometry, and three rotation angles describing the orientation of the heart with respect to the anatomical axes.

The resulting 23-dimensional vector was used as the latent code and provided as input to the neural surrogate \mathcal{N}_{LF} , together with the spatial coordinates of the evaluation points and the electrode positions.

2.5 DeepSDF geometry encoding

To represent the joint anatomical geometries of the torso and heart, we adapted the DeepSDF model with an auto-decoder architecture [29, 45]. This framework encodes complex multi-object geometries into a continuous low-dimensional latent representation. The decoder network is trained to approximate signed distance functions (SDFs) corresponding to four anatomical surfaces: the torso, the left ventricular endocardium (LV), the right ventricular endocardium (RV), and the epicardium.

We denote the DeepSDF decoder by

$$\mathcal{N}_{\text{SDF}}: \mathbb{R}^3 \times \mathbb{R}^{d_z} \rightarrow \mathbb{R}^4,$$

where $\mathbf{x} \in \mathbb{R}^3$ is a spatial coordinate and $\mathbf{z} \in \mathbb{R}^{d_z}$ is a latent code describing the joint heart-torso geometry. The weights of the network are denoted by $\boldsymbol{\theta}_{\text{SDF}}$. The network outputs four signed distance functions, one for each surface, evaluated at location \mathbf{x} for the geometry encoded by \mathbf{z} .

The model consists of five fully connected layers with 256 neurons each. The input is the concatenation of the latent code \mathbf{z} (dimension $d_z = 16$) and the spatial coordinate \mathbf{x} . To preserve spatial information across layers, the coordinate vector is additionally concatenated with the output of the third hidden layer (skip connection). Fourier feature encoding [43] is applied to improve the representation of high-frequency geometric details, and Lipschitz layers with Lipschitz regularization [22] are used to enhance stability and smoothness of the learned latent space. ReLU activations and the Adam optimizer are employed during training.

For each geometry (p), let

$$\mathbf{X}_{\text{SDF}}^{(p)} = \{(\mathbf{x}_n^{(p)}, s_n^{(p)})\}_{n=1}^{N_p}$$

denote sampled spatial points and their corresponding ground-truth SDF values. The training objective jointly optimizes decoder parameters ϕ and latent codes $\mathbf{z}^{(p)}$:

$$\begin{aligned} \mathcal{L}(\boldsymbol{\theta}_{\text{SDF}}, \{\mathbf{z}^{(p)}\}) = & \frac{1}{N} \frac{1}{N_p} \sum_{p=1}^N \sum_{n=1}^{N_p} \left| \mathcal{N}_{\text{SDF}}(\mathbf{x}_n^{(p)}, \mathbf{z}^{(p)}; \boldsymbol{\theta}_{\text{SDF}}) - s_n^{(p)}(\mathbf{x}_n^{(p)}) \right|^2 \\ & + \lambda_{\text{prior}} \frac{1}{N_p} \sum_{n=1}^{N_p} \|\mathbf{z}^{(p)}\|_2^2 + \lambda_{\text{Lip}} \mathcal{L}_{\text{Lip}}(\boldsymbol{\theta}_{\text{SDF}}), \end{aligned} \quad (6)$$

where the first term corresponds to the mean squared reconstruction error between predicted and ground-truth SDF values. The second term imposes a Gaussian prior on the latent codes, and \mathcal{L}_{Lip} denotes the Lipschitz regularization term weighted by λ_{Lip} .

Latent code inference. After training, the decoder parameters $\boldsymbol{\theta}_{\text{SDF}}$ are kept fixed. Given a new joint geometry represented by sampled SDF values $\{(\mathbf{x}_n, s_n)\}_{n=1}^N$ (e.g., derived from a point cloud or segmented MRI), the corresponding latent code is obtained by maximum a posteriori (MAP) estimation:

$$\mathbf{z}^* = \arg \min_{\mathbf{z}} \left(\frac{1}{N} \sum_{n=1}^N |g_{\phi}(\mathbf{x}_n, \mathbf{z}) - s_n|^2 + \lambda_{\text{Maha}} (\mathbf{z} - \boldsymbol{\mu})^{\top} \Sigma^{-1} (\mathbf{z} - \boldsymbol{\mu}) \right), \quad (7)$$

where $\boldsymbol{\mu}$ and Σ denote the empirical mean and covariance of the latent codes estimated from the latent codes obtained after then training phase, and γ is the regularization parameter. The Mahalanobis regularization term enforces anatomical plausibility by penalizing latent codes that are statistically distant from the learned distribution. This prior accounts for correlations between latent dimensions and improves robustness to noise in the input geometry.

2.6 Evaluation metrics

For the DeepSDF network \mathcal{N}_{SDF} , we evaluated reconstruction quality using the Chamfer distance. For each of the four predicted SDF fields (LV, RV, epicardium, and torso), both the predicted and ground-truth SDF values were thresholded by retaining only points with $\text{SDF} \leq 0$, corresponding to the interior of the surfaces. The resulting point clouds were then compared using the Chamfer distance as a similarity metric.

Specifically, for two point clouds X and Y containing N_X and N_Y points, respectively, the Chamfer distance is defined as

$$\text{CD}(X, Y) = \frac{1}{2} \left(\frac{1}{N_X} \sum_{x \in X} \min_{y \in Y} \|x - y\|_2 + \frac{1}{N_Y} \sum_{y \in Y} \min_{x \in X} \|y - x\|_2 \right). \quad (8)$$

For the neural surrogate \mathcal{N}_{LF} , we assess surrogate performance at two complementary levels.

- **Lead-field (LF) metrics.** We quantify the accuracy of the predicted lead-field gradients by computing the mean angular error and the mean magnitude error over sampled spatial points.

For a given geometry s and electrode j , the angular error at point $\mathbf{x} \in \mathbf{X}^{(s)}$ is defined as

$$\theta_j^{(s)}(\mathbf{x}) = \arccos \left(\frac{\nabla \hat{Z}_j^{(s)}(\mathbf{x}) \cdot \nabla Z_j^{(s)}(\mathbf{x})}{\|\nabla \hat{Z}_j^{(s)}(\mathbf{x})\|_2 \|\nabla Z_j^{(s)}(\mathbf{x})\|_2 + \varepsilon} \right).$$

The mean angular error is obtained by averaging $\theta_j^{(s)}(\mathbf{x})$ over all sampled points $\mathbf{x} \in \mathbf{X}^{(s)}$ and over electrodes. Magnitude errors are computed as the mean Euclidean difference

$$\|\nabla \hat{Z}_j^{(s)}(\mathbf{x}) - \nabla Z_j^{(s)}(\mathbf{x})\|_2,$$

again averaged over spatial sampling points and electrodes.

We additionally report lead-field metrics computed on a restricted subset of electrodes, namely the six precordial leads of the standard 12-lead ECG. In this case, the averaging is performed only over this subset of electrodes.

- **ECG-level metrics.** For each lead, we compare the predicted ECG waveform $\hat{V}(t)$ with the FEM reference solution $V(t)$. We report the relative ℓ_2 error over the time interval $[0, T]$,

$$\frac{\|\hat{V} - V\|_{L^2(0,T)}}{\|V\|_{L^2(0,T)}}.$$

In addition, we evaluate differences in clinically relevant waveform characteristics, including QRS amplitude and QRS duration.

We note that the error in the lead-field gradient and the ECG error are directly related through Eq. (3). Assume that both $V(t)$ and $\hat{V}(t)$ are computed from the same transmembrane potential $V_m(\mathbf{x}, t)$, but using two different lead-field functions, $Z(\mathbf{x})$ and $\hat{Z}(\mathbf{x})$, respectively. Then their difference satisfies

$$V(t) - \hat{V}(t) = - \int_{\Omega_H} \mathbf{G}_i \nabla V_m(\mathbf{x}, t) \cdot \nabla (Z(\mathbf{x}) - \hat{Z}(\mathbf{x})) \, d\mathbf{x}.$$

Applying the Cauchy-Schwarz inequality in $L^2(\Omega_H)$ yields

$$|V(t) - \hat{V}(t)|^2 \leq \|\mathbf{G}_i \nabla V_m(\cdot, t)\|_{L^2(\Omega_H)}^2 \|\nabla (Z - \hat{Z})\|_{L^2(\Omega_H)}^2.$$

Integrating in time over $[0, T]$ gives

$$\|V - \hat{V}\|_{L^2(0,T)}^2 \leq C \|\nabla (Z - \hat{Z})\|_{L^2(\Omega_H)}^2,$$

where the constant

$$C_{V_m} = \int_0^T \|\mathbf{G}_i \nabla V_m(\cdot, t)\|_{L^2(\Omega_H)}^2 \, dt$$

depends only on the transmembrane potential and tissue conductivities. Therefore, controlling the error in the lead-field gradient ∇Z directly yields a bound on the ECG error, through a multiplicative constant determined by the underlying electrophysiological activation.

2.7 Experimental setup

The main goal of this work was to verify the feasibility and evaluate the quality of representing $\nabla Z(x)$ using a neural network, especially when employed to compute the ECG.

To model the transmembrane potential under different physiological and pathological conditions, we considered the following model:

$$V_m(\mathbf{x}, t) = U(t - \tau(\mathbf{x})),$$

where $U(\xi)$ is a template action potential, based on the ten Tusscher-Panfilov model of human ventricular cells [44], and $\tau(\mathbf{x})$ is the activation map satisfying the anisotropic eikonal equation [11, 10]:

$$\begin{cases} \sqrt{\mathbf{V} \nabla \tau \cdot \nabla \tau} = 1, & \mathbf{x} \in \Omega_H, \\ \tau(\mathbf{x}_i) = \tau_i, & i = 1, \dots, N_{\text{stim}}, \end{cases} \quad (9)$$

with $\{(\mathbf{x}_i, t_i)\}$ being the initial set of activation sites and the symmetric positive definite tensor $\mathbf{V}(\mathbf{x})$ encoding for the anisotropic conduction velocity:

$$\mathbf{V}(\mathbf{x}) = v_t^2 \mathbf{I} + (v_f^2 - v_t^2) \mathbf{f} \otimes \mathbf{f} \quad (10)$$

where v_f and v_t represents conduction velocities along the fiber direction and orthogonally to the fibers, respectively.

The eikonal model allows us to simulate different conditions depending on the boundary values $\{\mathbf{x}_i, t_i\}$. Two cases were studied: (i) activation from two pacing sites, with the endocardium of the right ventricle near the apical region and the epicardium of the free wall of the left ventricle, mimicking the cardiac resynchronization therapy, and (ii) simulation of sinus rhythm with the construction of a synthetic conduction system based on rule-based Purkinje network [37, 1].

We examined both the standard clinical electrode configuration for a 12-lead ECG and a denser set of unipolar electrodes distributed across the front of the torso, mimicking BSPMs and the possible uncertainty in the precordial electrodes of a 12-lead ECG. Standard leads allow for direct comparison with clinical ECG morphology, while additional unipolar electrodes provide a wider choice of potential fields and are relevant for applications such as BSPMs and uncertainty in electrode placement.

The results were also compared with pseudo-lead-field, which is obtained from problem (4) assuming $\mathbf{G} = \mathbf{G}_0 = \sigma_0 \mathbf{I}$ and $\Omega = \mathbb{R}^3$ (infinite torso assumption). The exact solution reads as follows:

$$Z^*(\mathbf{x}) = \frac{1}{4\pi\sigma_0\|\mathbf{e}_j - \mathbf{x}\|}, \quad (11)$$

Pseudo lead-field provides a simple approximation that requires no torso information, only the electrode locations.

2.8 Implementation details

All forward problems were solved on tetrahedral meshes using the finite element method implemented in the FEniCS framework [23]. Since the lead-field formulation is defined up to an additive constant, the null space was removed by enforcing a reference potential constraint, see Eq. (2). This ensures uniqueness of the numerical solution and stable computation of the gradient field.

The lead-field neural surrogate \mathcal{NN}_{LF} was trained for 800 epochs using the Adam optimizer with an initial learning rate of 1×10^{-3} . After 400 epochs, the learning rate was reduced by a factor of two. The DeepSDF model $\mathcal{NN}_{\text{SDF}}$ was trained for 2000 epochs using the same optimizer and initial learning rate of 1×10^{-3} . No early stopping strategy was employed. Training was performed on a single NVIDIA A100 GPU. The batch size was set to 10 for \mathcal{NN}_{LF} and 6 for the DeepSDF model. Due to the large number of sampled spatial points, training was computationally intensive: approximately 2 days were required for \mathcal{NN}_{LF} and 1.2 days for the DeepSDF model. Finite element meshes contained approximately 40 000 nodes and 150 000 tetrahedra for the heart domain, and approximately 120 000 nodes and 600 000 tetrahedra for the full torso model.

2.9 Software

Finite element meshes were generated using Gmsh [15], while mesh preprocessing, postprocessing, and visualization were performed using PyVista [42] and VTK [39]. Surface remeshing was carried out using `pyacvd`. Universal ventricular coordinates (UVCs), Laplace–Dirichlet rule-based (LDRB) myocardial fiber assignment, and lead-field computations were implemented in FEniCS. Eikonal simulations were performed using `fim-python`. Purkinje networks were generated using the `fractal-tree` package, and activation times at Purkinje–myocardial junctions (PMJs) were computed using `networkx`. Signed distance functions and point cloud sampling were computed using the `mesh_to_sdf` library, which internally relies on `pyopengl`, `trimesh`, `pyrender`, and `pyDOE`. Latin hypercube sampling (LHS) of geometries and electrode positions was performed explicitly using `pyDOE`. Both the DeepSDF model and the lead-field gradient surrogate were implemented in PyTorch [2]. Training was managed using PyTorch Lightning, with hyperparameter optimization performed via Optuna and experiment tracking conducted using MLflow. Standard scientific computing tasks, including interpolation and nearest-neighbor search (KDTree) for Chamfer distance computation, were carried out using SciPy, NumPy, and scikit-learn. Visualization and plotting were performed using Vedo, PyVista, and Seaborn.

3 Results

3.1 DeepSDF-based geometry encoding

The DeepSDF model provides a compact latent representation of the joint heart-torso anatomy. Reconstruction accuracy was evaluated as follows. Predictions from the DeepSDF decoder were computed on a uniform regular grid of 128^3 spatial points. The predicted SDF values for the LV, RV, epicardium, and

torso were then compared with the corresponding ground-truth SDF values interpolated onto the same 128^3 grid.

A heatmap of the reconstruction errors is shown in Fig. 4. Errors were consistently low across the torso and ventricular surfaces, indicating that the learned latent codes preserve the geometric detail required by the downstream lead-field surrogate. Test errors were comparable to training and validation errors for all surfaces, indicating good generalization.

Slightly larger errors were observed for the RV endocardium and the epicardium, likely due to their higher geometric complexity compared with the LV endocardium and the torso. The maximum error remained below 1.6 mm, while the mean error ranged between 0.75 mm and 1.28 mm across surfaces.

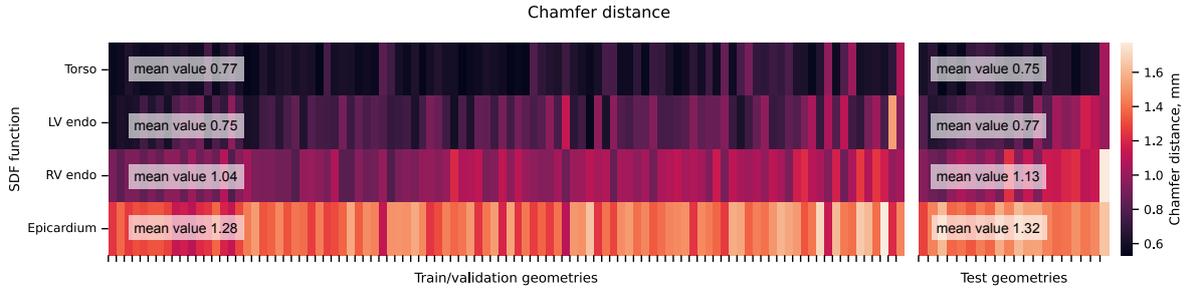


Figure 4: Chamfer distances (in mm) for the four SDF surfaces—torso, LV endocardium, RV endocardium, and epicardium, computed 80/20 training/validation geometries (left panel) and 10 test geometries (right panel). The color scale indicates the Chamfer distance value (in mm), with lighter colors corresponding to larger errors. For clarity, both the SDF surfaces and the geometry sets are sorted by increasing error.

3.2 Lead-field neural surrogate performance

We report quantitative and qualitative results for the lead-field gradient prediction network, as well as for the ECG signals obtained by inserting the predicted ∇Z into Eq. (3). Unless otherwise stated, FEM-based lead-field gradients are used as the reference solution.

For each of the 10 joint geometries from the test sample, SDF functions were assigned and test latent codes were calculated by solving the inference problem in Eq. (7).

3.2.1 Qualitative comparison of gradient fields

The surrogate accurately reproduces the FEM reference ∇Z across the torso volume. Fig. 5 shows streamlines of the lead-field gradient (i.e., curves whose tangents coincide with the direction of ∇Z) for a representative test geometry. Panel A displays the FEM solution (blue), and Panel B shows the surrogate prediction (green).

Visually, the predicted gradients closely match the ground truth in both direction and curvature. In particular, both solutions exhibit a clear change in direction at the heart-torso interface, reflecting the discontinuity in conductivity across tissues. This agreement indicates that the surrogate captures the dominant geometric and boundary-condition effects governing the lead field.

3.2.2 Point-wise error distributions

Fig. 6 shows cumulative distribution functions (CDFs) of the angular error and relative magnitude error over the test cohort. Two evaluation regions are considered: (A) the full torso domain (matching the training sampling distribution), and (B) a restricted region consisting of points inside the heart and within 10 mm of the heart surface.

The heart-focused evaluation is motivated by two considerations. First, the ECG depends only on ∇Z within the heart domain Ω_H (see Eq. (3)). Second, the heart-torso interface is the region where ∇Z changes direction most abruptly, and is therefore potentially most sensitive to prediction errors.

The dashed vertical lines indicate the median (50%) and 95% quantile. For example, in the full-domain setting, 50% of all test points exhibit an angular error below 2.64° for the DeepSDF-based model and below 3.51° for the PCA-based model. Across all configurations, the DeepSDF-based model consistently achieves smaller angular errors, while magnitude errors are comparable between the two

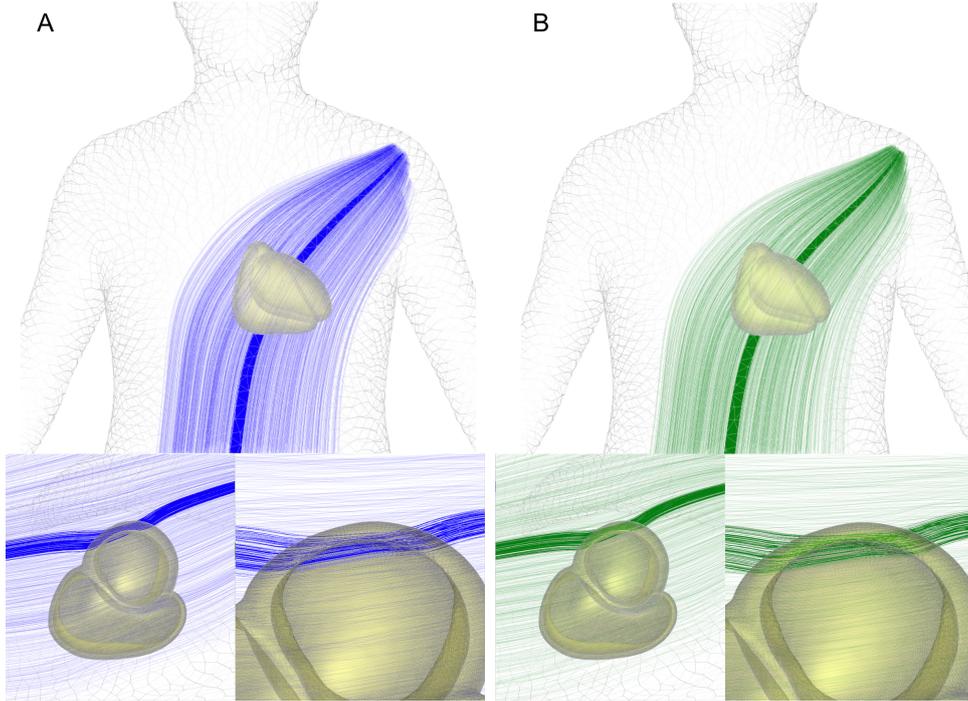


Figure 5: Streamline visualization of the lead-field gradient for a unipolar lead (left shoulder). (A) FEM-based lead-field gradient. (B) Predicted lead-field gradient. Smaller sub-panels show close-up views of the heart and the LV free wall. Streamlines intersecting the LV free wall are highlighted. The bend in the streamlines at the heart-torso interface (highlighted in yellow) reflects the change in conductivity across tissues.

models. The similarity in magnitude errors suggests that both models accurately capture the overall scaling of the gradient field, with differences primarily arising in directional alignment.

3.2.3 Electrode-wise error analysis

Fig. 7 reports electrode-wise average angular errors for both encoding strategies, evaluated on 100 uniformly distributed unipolar electrodes and on the 9 independent electrodes of the standard 12-lead ECG.

The largest errors are observed for the aVF electrode and for precordial leads V1-V3. For precordial electrodes, this can be attributed to their proximity to the heart surface and thus to the heart-torso interface, where the gradient direction varies rapidly and is more challenging to approximate. In these regions, the gradient field exhibits higher curvature and stronger spatial heterogeneity, which amplifies directional errors.

The spatial distribution of angular errors, averaged over 10 test geometries and interpolated onto the torso surface, is shown in Fig. 8A. High-error regions (red) are primarily localized in the central anterior torso, consistent with the proximity to the heart.

3.2.4 Computational efficiency

We compared the computational cost of the proposed surrogate with the FEM-based lead-field computation. On a CPU implementation, solving the forward problem with FEM required approximately 6 s per lead, excluding mesh generation. Mesh construction required an additional 22 s per geometry.

In contrast, evaluation of the trained lead-field surrogate required approximately 250 ms per lead on a CPU (batch size = 1). Thus, even without GPU acceleration or batching, the surrogate provides a 24-fold speedup per lead.

3.2.5 Impact on ECG signals

To assess practical relevance, we computed ECG signals using the predicted lead-field gradients. Fig. 8B-C shows the spatial distribution of the root mean squared error (RMSE) between body-surface potentials obtained with FEM-based and predicted ∇Z for two activation patterns: left bundle branch block

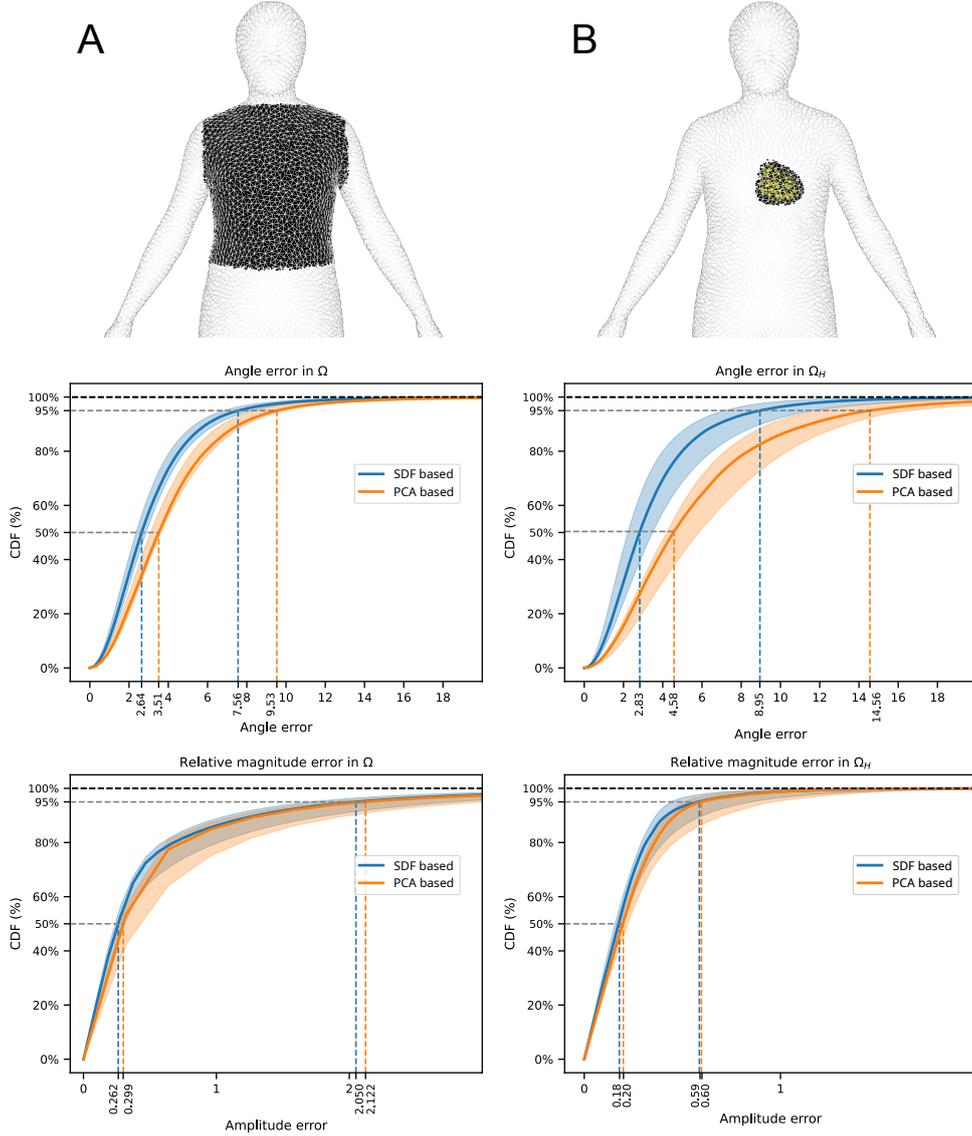


Figure 6: Cumulative distribution functions (CDFs) of prediction errors for ∇Z over the test dataset. Left column: errors over the full torso domain. Right column: errors restricted to points inside the heart and within 10 mm of the heart surface. Top row: angular error (degrees). Bottom row: relative magnitude error. Blue and orange curves correspond to DeepSDF-based and PCA-based models, respectively. Solid lines indicate the mean CDF across 10 test patients; shaded regions show the range (minimum-maximum). Dashed vertical lines mark the median and 95th percentile error values.

(LBBB) and sinus rhythm. Regions of larger ECG error largely coincide with regions of larger angular gradient error, particularly in areas close to the heart.

Despite these localized discrepancies, the resulting ECG waveforms closely match the FEM-based signals (Fig. 9, left), indicating that the surrogate preserves clinically relevant waveform morphology.

Tab. 1 summarizes quantitative performance on the test set. We report both field-level errors (accuracy of ∇Z) and ECG-level errors. We observe that in general the DeepSDF-based approach is always performing better than the PCA-based encoding.

4 Discussion

In this work, we proposed a shape-informed neural surrogate model for the lead-field gradient arising in the coupled heart-torso problem, designed as a drop-in replacement for the full-order model in forward ECG simulations. Importantly, we did not surrogate the mapping from the transmembrane potential V_m

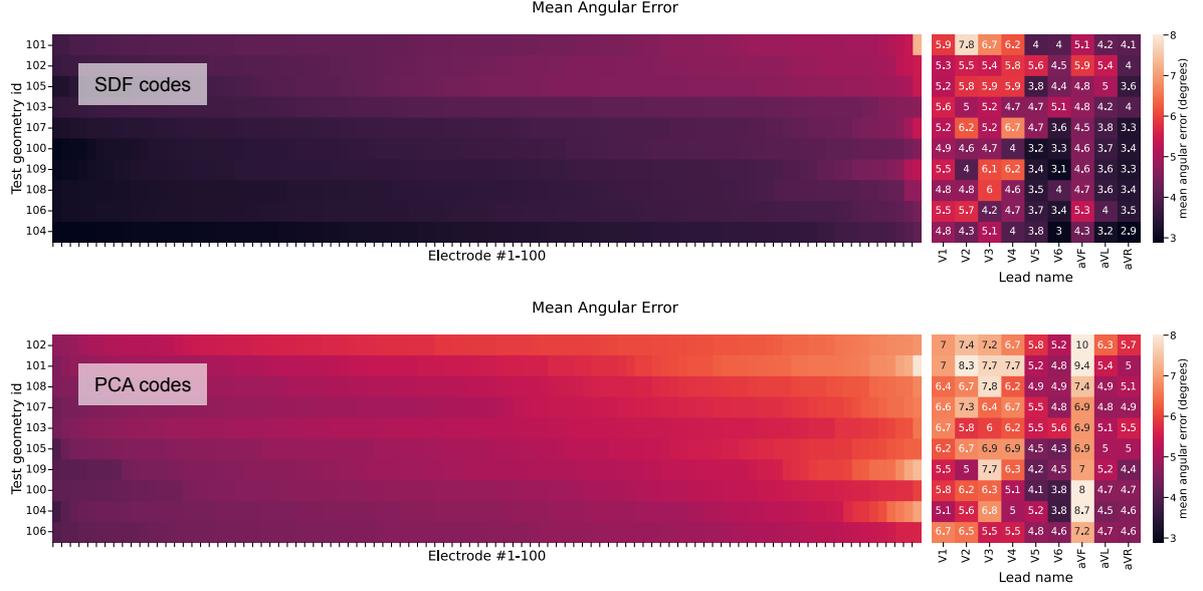


Figure 7: Average angular error (degrees) in predicting ∇Z using DeepSDF-based (top row) and PCA-based (bottom row) encodings. Left column: errors across 100 uniformly distributed unipolar electrodes. Right column: errors across the 9 independent electrodes of the standard 12-lead ECG. Electrodes and patients are sorted by increasing error for visualization clarity.

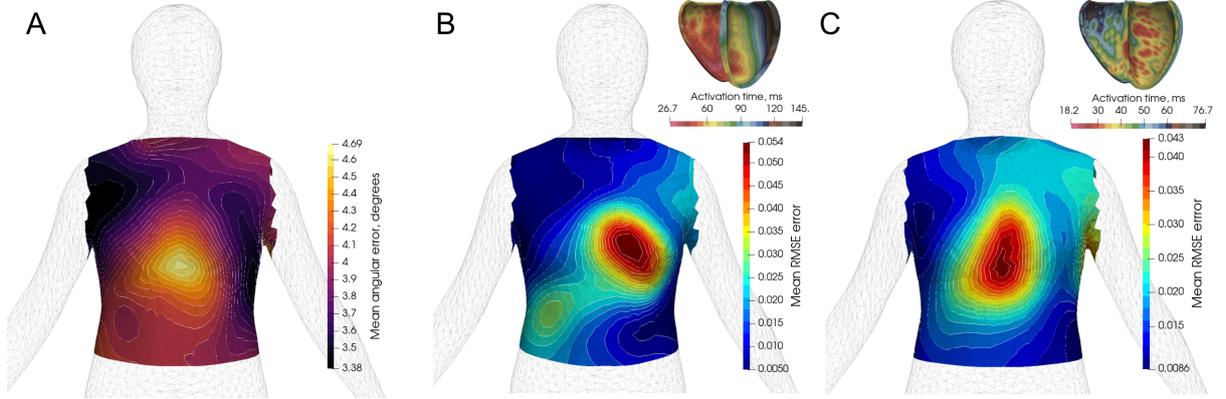


Figure 8: Spatial distribution of errors projected onto the torso surface and averaged over 10 test cases. (A) Mean angular error (degrees) between DeepSDF-based predicted and FEM-based ∇Z , computed over 100 electrodes. (B–C) Mean RMSE of unipolar ECG signals for LBBB (B) and sinus rhythm (C), computed using FEM-based and predicted ∇Z . Errors are interpolated onto the torso surface; blue indicates smaller error. Insets show representative ventricular activation maps used for ECG computation.

Geometry encoding	Angular error	Angular error (ECG leads)	ECG rel. ℓ_2 error
PCA-based	$5.22^\circ \pm 0.61^\circ$	$5.93^\circ \pm 1.26^\circ$	0.024 ± 0.013
DeepSDF-based	$3.89^\circ \pm 0.51^\circ$	$4.64^\circ \pm 0.99^\circ$	0.018 ± 0.010

Table 1: Average angular error (degrees) between FEM-based and predicted lead-field gradients ∇Z for points within the heart, and corresponding ECG relative ℓ_2 error. The third column reports angular error restricted to precordial ECG leads. Values are given as mean \pm standard deviation across test geometries.

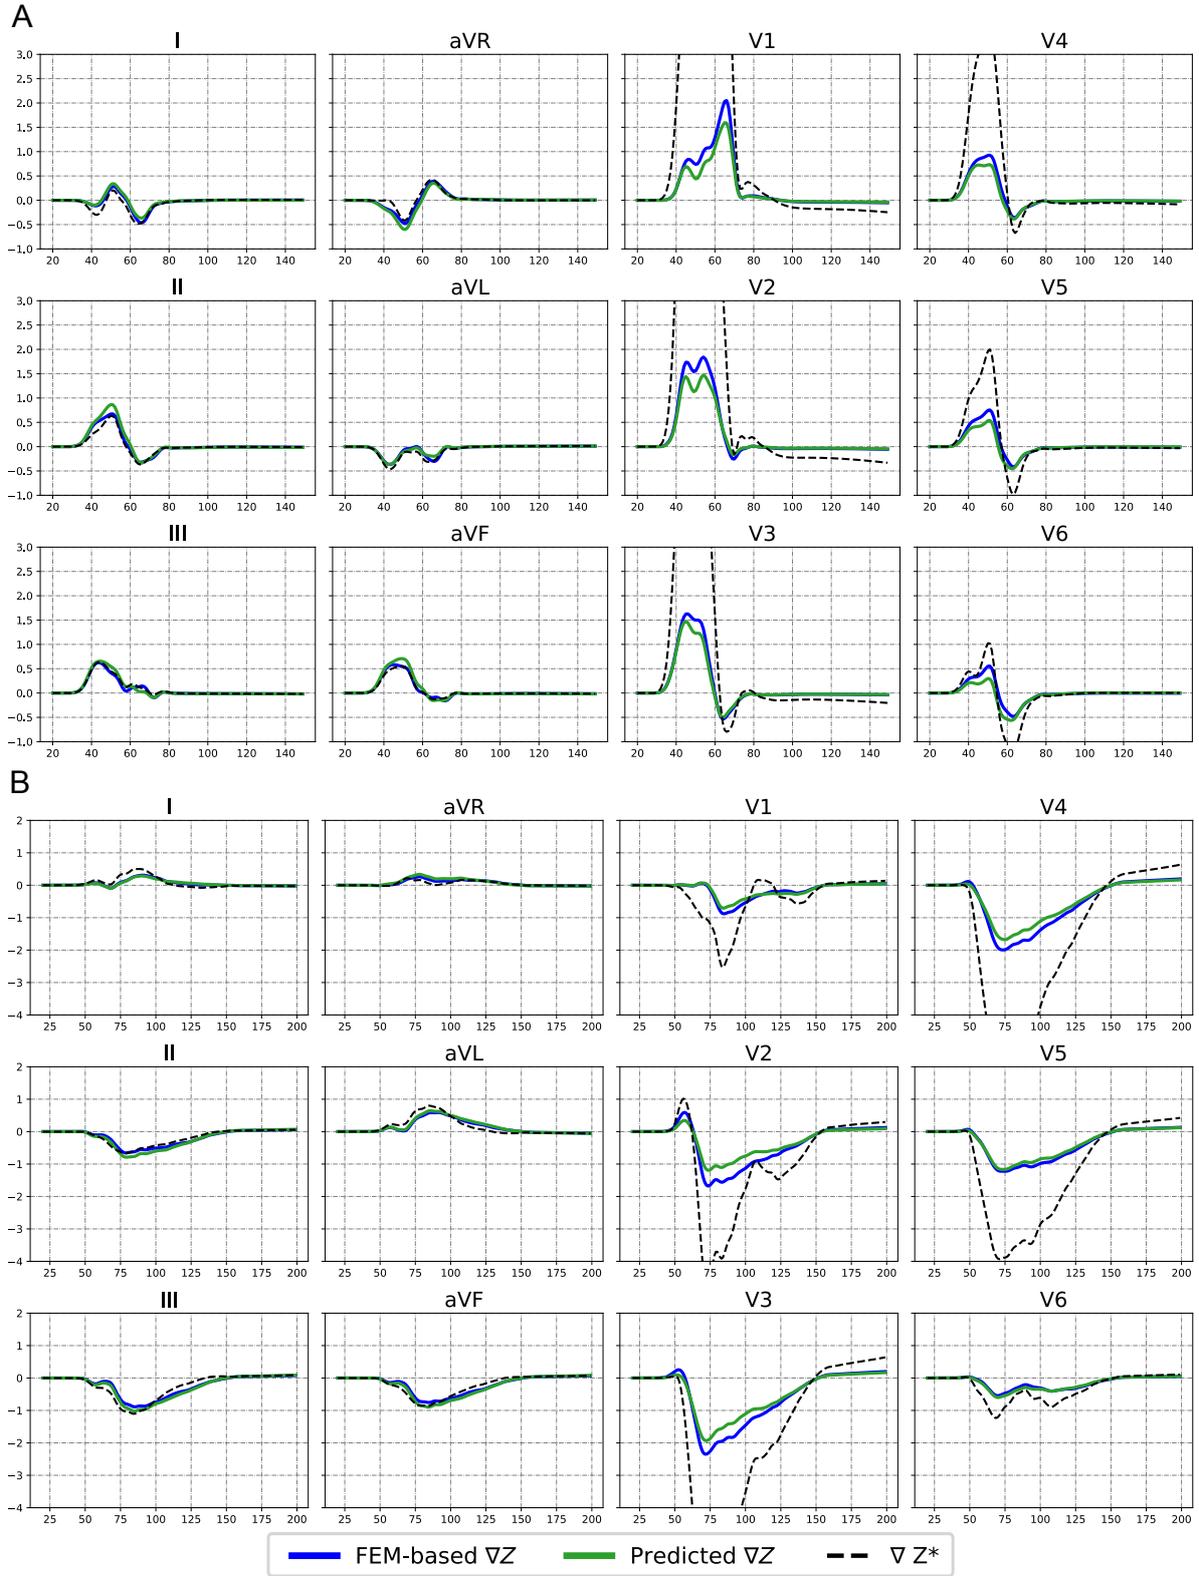


Figure 9: Forward ECG simulation for sinus rhythm (A) and LBBB (B) cases, computed using FEM-based (ground truth, blue), predicted ∇Z (green) and using the pseudo lead-field formulation ∇Z^* .

to the ECG signal, i.e., the solution operator of the pseudo-bidomain problem (1). Instead, we focused exclusively on approximating the lead-field gradient. This design choice offers several advantages. The ECG can be expressed explicitly via Eq. (3), which defines a linear functional acting on V_m . Once the lead field is available, ECG computation is computationally inexpensive and remains fully general with respect to the underlying transmembrane potential and the intracellular conductivity tensor \mathbf{G}_i . The latter plays a crucial role in determining ECG amplitude and morphology [30]. By surrogating only the lead-field gradient, we preserve this generality while replacing only the geometrically dependent component of the model, namely the mapping that embeds torso conductivity, electrode configuration, and anatomical variability.

A key strength of the proposed framework is that it is fully shape-informed. Still, for a new patient, only limited geometric information is required: once a surface representation is available (e.g., from segmented imaging data or sparse point clouds), inference of the latent code can be performed efficiently using either PCA projection or an inexpensive DeepSDF inference step [45]. Such a mechanism is particularly attractive in the solution of the inverse problem of electrocardiography [47, 25] and in time-dependent settings, where cardiac geometry may evolve over time (e.g., during cardiac cycles or remodeling). Tracking geometric changes and updating the latent code enables rapid evaluation of the resulting impact on the ECG without repeated full-order simulations [51]. The proposed framework is also particularly attractive in settings requiring repeated forward simulations with varying electrode configurations. For example, in electro-anatomical mapping procedures during catheter ablation [46], the recording electrode is continuously repositioned, potentially acquiring thousands of electrograms. In such scenarios, the FEM-based lead-field computation is impractical in real time, whereas surrogate inference enables near-interactive evaluation. More broadly, the method is well suited for scenarios involving dense electrode arrays or iterative inverse procedures.

In constructing the surrogate, we adopted a segregated approach that decouples geometry encoding from neural field regression. As shown in recent related work [8], such a strategy is particularly suitable in data-scarce biomedical settings. The geometric encoder (either PCA-based or DeepSDF-based) serves as a generative shape model that provides a compact, low-dimensional representation of joint heart-torso anatomy. This representation simplifies the input design of the neural surrogate and enables data augmentation through sampling in the latent space. In turn, this improves robustness and generalization when only a limited number of high-fidelity simulations are available.

From a quantitative standpoint, the surrogate achieves mean angular errors for the DeepSDF-based (resp. PCA-based) encoding below 4° (resp. slightly above 4°) and relative ECG errors below 2% (resp. 2.5%) on the test set. While such errors would not be considered negligible from a purely numerical analysis perspective, they are small in the context of forward ECG modeling. As shown in Sec. 2.6, the ECG error can be bounded in terms of the L^2 error of ∇Z , with a multiplicative constant depending only on the electrophysiological activation. Therefore, controlling the gradient error directly limits the induced ECG error. Angular errors are localized primarily near the heart-torso interface and, despite these localized discrepancies, the resulting ECG waveforms preserve clinically relevant morphology and amplitude. Compared with the classical pseudo lead-field approximation, the proposed method achieves substantially lower angular and ECG errors while retaining fast inference time. Both PCA-based and DeepSDF-based encodings yield accurate lead-field surrogates; however, quantitative results (Table 1) indicate a consistent advantage of the DeepSDF representation. Specifically, the DeepSDF-based model achieves a lower mean angular error ($3.89^\circ \pm 0.51^\circ$ versus $5.22^\circ \pm 0.61^\circ$ for PCA) and a lower ECG relative ℓ_2 error (1.8% versus 2.4%). The improvement is particularly evident for precordial leads, where accurate representation of localized geometric detail near the heart-torso interface is critical.

In its current CPU-based implementation, the surrogate inference per lead is roughly $25\times$ faster than the FEM computation, even without GPU acceleration. This gap becomes more pronounced in configurations with many electrodes, such as BSPMs comprising hundreds of leads.

In the present study, myocardial fiber orientations were assigned using a rule-based method and kept fixed once generated from the anatomical geometry. Similarly, tissue conductivities were assumed fixed across geometries. While intracellular conductivity \mathbf{G}_i influences ECG amplitude through Eq. (3), its impact on the lead-field solution itself is expected to be limited, as the lead field primarily depends on torso geometry and electrode configuration. Nevertheless, extending the surrogate to include additional input parameters, such as conductivity variations or fiber anisotropy, is conceptually straightforward: these parameters could be appended to the latent code. Such extensions would, however, increase the dimensionality of the input space and the associated increased requirement of training data. Similarly, anatomical structures such as lungs and blood cavities introduce additional conductivity heterogeneities, and could be incorporated into the geometric encoding and surrogate framework.

Beyond the specific application to forward ECG simulation, this work contributes to a broader class of geometry-conditioned surrogate models for PDE operators [9, 41]. In line with recent developments in shape-informed surrogate modeling for cardiac mechanics [8], we demonstrate that the action of a parameter-dependent solution operator can be efficiently approximated by decoupling geometric encoding from field approximation. The idea of approximating geometry-dependent Green functions or solution kernels extends naturally beyond cardiac electrophysiology. In the present case, the lead-field function plays the role of a Green-type kernel for the pseudo-bidomain formulation, mapping cardiac source terms to body-surface potentials. Learning this operator in a geometry-aware manner enables efficient reuse across different activation patterns and electrophysiological states, without retraining or restricting the space of admissible transmembrane potentials. Lead fields also arise in electroencephalography (EEG) and other bioelectric inverse problems, where accurate yet computationally efficient forward models are essential. More generally, many elliptic and parabolic PDEs admit integral representations in terms of Green’s functions whose structure depends strongly on domain geometry. The proposed framework suggests a pathway toward learning such geometry-dependent kernels in a data-efficient manner.

Acknowledgements

This work has been supported by the project PRIN2022 (MUR, Italy, 2023-2025, no. P2022N5ZNP) “SIDDMS: shape-informed data-driven models for parametrized PDEs, with application to computational cardiology”, funded by the European Union (Next Generation EU, Mission 4 Component 2). F.B., S. Pagani and F.R. acknowledge the grant Dipartimento di Eccellenza 2023-2027, funded by MUR, Italy. F.B., S.Pagani, S.Pezzuto and F.R. are members of GNCS, “Gruppo Nazionale per il Calcolo Scientifico” (National Group for Scientific Computing) of INdAM (Istituto Nazionale di Alta Matematica). F.B. acknowledges the “INdAM - GNCS Project”, codice CUP E53C24001950001. S.Pezzuto acknowledges the support of the CSCS-Swiss National Supercomputing Centre project no. lp100 and the SNSF-FWF project “CardioTwin” (no. 214817).

References

- [1] F Álvarez-Barrientos, M Salinas-Camus, S Pezzuto, and F Sahli Costabal. “Probabilistic learning of the Purkinje network from the electrocardiogram”. In: *Medical Image Analysis* 101 (Apr. 2025), p. 103460. DOI: 10.1016/j.media.2025.103460. arXiv: 2312.09887.
- [2] J Ansel, E Yang, H He, et al. “PyTorch 2: Faster Machine Learning Through Dynamic Python Bytecode Transformation and Graph Compilation”. In: *29th ACM International Conference on Architectural Support for Programming Languages and Operating Systems, Volume 2 (ASPLOS ’24)*. ACM, Apr. 2024. DOI: 10.1145/3620665.3640366. URL: <https://docs.pytorch.org/assets/pytorch2-2.pdf>.
- [3] W Bai, W Shi, A de Marvao, TJW Dawes, DP O’Regan, SA Cook, and D Rueckert. “A biventricular cardiac atlas built from 1000+ high resolution MR images of healthy subjects and an analysis of shape and motion”. en. In: *Med. Image Anal.* 26.1 (Dec. 2015), pp. 133–145. DOI: 10.1016/j.media.2015.08.009.
- [4] JD Bayer, RC Blake, G Plank, and NA Trayanova. “A novel rule-based algorithm for assigning myocardial fiber orientation to computational heart models”. en. In: *Ann. Biomed. Eng.* 40.10 (Oct. 2012), pp. 2243–2254. DOI: 10.1007/s10439-012-0593-5.
- [5] JA Bergquist, B Zenger, LC Rupp, A Busatto, J Tate, DH Brooks, A Narayan, and RS MacLeod. “Uncertainty quantification of the effect of cardiac position variability in the inverse problem of electrocardiographic imaging”. In: *Physiological Measurement* 44.10 (2023), p. 105003.
- [6] M Boulakia, S Cazeau, MA Fernández, JF Gerbeau, and N Zemzemi. “Mathematical modeling of electrocardiograms: a numerical study”. en. In: *Ann. Biomed. Eng.* 38.3 (Mar. 2010), pp. 1071–1097. DOI: 10.1007/s10439-009-9873-0.
- [7] J Camps, ZJ Wang, R Doste, LA Berg, M Holmes, B Lawson, J Tomek, K Burrage, A Bueno-Orovio, and B Rodriguez. “Harnessing 12-lead ECG and MRI data to personalise repolarisation profiles in cardiac digital twin models for enhanced virtual drug testing”. In: *Medical image analysis* 100 (2025), p. 103361.

- [8] D Carrara, M Hirschvogel, F Bonizzoni, S Pagani, S Pezzuto, and F Regazzoni. “Shape-informed Cardiac Mechanics Surrogates in Data-Scarce Regimes via Geometric Encoding and Generative Augmentation”. Preprint. 2026. arXiv: 2602.20306.
- [9] G Catalani, J Fesquet, X Bertrand, F Tost, M Bauerheim, and J Morlier. “Towards scalable surrogate models based on neural fields for large scale aerodynamic simulations”. In: *Computers & Fluids* 306 (2026), p. 106929. DOI: 10.1016/j.compfluid.2025.106929.
- [10] P Colli Franzone and L Guerri. “Spreading of excitation in 3-D models of the anisotropic cardiac tissue. I. Validation of the eikonal model”. en. In: *Math. Biosci.* 113.2 (Feb. 1993), pp. 145–209. DOI: 10.1016/0025-5564(93)90001-q.
- [11] P Colli Franzone, L Guerri, and S Rovida. “Wavefront propagation in an activation model of the anisotropic cardiac tissue: asymptotic analysis and numerical simulations”. en. In: *J. Math. Biol.* 28.2 (1990), pp. 121–176. DOI: 10.1007/bf00163143.
- [12] P Colli Franzone, B Taccardi, and C Viganotti. “An approach to inverse calculation of epicardial potentials from body surface maps.” In: *Advances in cardiology* 21 (1978), pp. 50–54.
- [13] P Colli Franzone, LF Pavarino, and S Scacchi. *Mathematical cardiac electrophysiology*. Vol. 13. Cham: Springer, 2014.
- [14] DB Geselowitz. “On the theory of the electrocardiogram”. In: *Proc. IEEE Inst. Electr. Electron. Eng.* 77.6 (June 1989), pp. 857–876. DOI: 10.1109/5.29327.
- [15] C Geuzaine and JF Remacle. “Gmsh: a three-dimensional finite element mesh generator with built-in pre- and post-processing facilities”. In: *Int. J. Numer. Methods Eng.* 79.11 (Sept. 2009), pp. 1309–1331. DOI: 10.1002/nme.2579.
- [16] T Grandits, K Gillette, G Plank, and S Pezzuto. “Accurate and Efficient Cardiac Digital Twin from surface ECGs: Insights into Identifiability of Ventricular Conduction System”. In: *Medical Image Analysis* 105 (Oct. 2025), p. 103641. DOI: 10.1016/j.media.2025.103641. arXiv: 2411.00165.
- [17] R Hoekema, GJ Uijen, L van Erning, and A van Oosterom. “Interindividual variability of multilead electrocardiographic recordings: influence of heart position”. en. In: *J. Electrocardiol.* 32.2 (Apr. 1999), pp. 137–148. DOI: 10.1016/s0022-0736(99)90092-4.
- [18] DU Keller, FM Weber, G Seemann, and O Dössel. “Ranking the influence of tissue conductivities on forward-calculated ECGs”. In: *IEEE Transactions on Biomedical Engineering* 57.7 (2010), pp. 1568–1576.
- [19] DP Kingma and J Ba. “Adam: A Method for Stochastic Optimization”. In: *Proceedings of the International Conference on Learning Representations (ICLR)*. 2015.
- [20] L Li, J Camps, B Rodriguez, and V Grau. “Solving the inverse problem of electrocardiography for cardiac digital twins: A survey”. In: *IEEE Reviews in Biomedical Engineering* 18 (2024), pp. 316–336.
- [21] L Li, H Smith, Y Lyu, J Camps, S Qian, B Rodriguez, A Banerjee, and V Grau. “Personalized topology-informed localization of standard 12-lead ECG electrode placement from incomplete cardiac MRIs for efficient cardiac digital twins”. In: *Medical Image Analysis* 101 (2025), p. 103472.
- [22] HTD Liu, F Williams, A Jacobson, S Fidler, and O Litany. “Learning smooth neural functions via Lipschitz regularization”. In: *Special Interest Group on Computer Graphics and Interactive Techniques Conference Proceedings*. Vol. 1. New York, NY, USA: ACM, Aug. 2022. DOI: 10.1145/3528233.3530713.
- [23] A Logg, KA Mardal, and G Wells. *Automated solution of differential equations by the finite element method: The FEniCS book*. Vol. 84. Springer Science & Business Media, 2012.
- [24] A Mincholé, E Zacur, R Ariga, V Grau, and B Rodriguez. “MRI-based computational torso / biventricular multiscale models to investigate the impact of anatomical variability on the ECG QRS complex”. In: *Front. Physiol.* 10 (Aug. 2019), p. 1103. DOI: 10.3389/fphys.2019.01103.
- [25] R Molero, A González-Ascaso, AM Climent, and MS Guillem. “Robustness of imageless electrocardiographic imaging against uncertainty in atrial morphology and location”. In: *Journal of Electrocardiology* 77 (2023), pp. 58–61.
- [26] C Nagel, S Schuler, O Dössel, and A Loewe. “A bi-atrial statistical shape model for large-scale in silico studies of human atria: model development and application to ECG simulations”. In: *Medical Image Analysis* 74 (2021), p. 102210.

- [27] UC Nguyễn, M Potse, F Regoli, et al. “An in-silico analysis of the effect of heart position and orientation on the ECG morphology and vectorcardiogram parameters in patients with heart failure and intraventricular conduction defects”. en. In: *J. Electrocardiol.* 48.4 (Sept. 2015), pp. 617–625. DOI: 10.1016/j.jelectrocard.2015.05.004.
- [28] F Odille, S Liu, P van Dam, and J Felblinger. “Statistical variations of heart orientation in healthy adults”. In: *2017 Computing in Cardiology Conference (CinC)*. Computing in Cardiology, July 2017. DOI: 10.22489/cinc.2017.225-058.
- [29] JJ Park, P Florence, J Straub, R Newcombe, and S Lovegrove. “DeepSDF: Learning continuous signed distance functions for shape representation”. In: *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*. 2019, pp. 165–174.
- [30] S Pezzuto, P Kaľavský, M Potse, FW Prinzen, A Auricchio, and R Krause. “Evaluation of a Rapid Anisotropic Model for ECG Simulation”. In: *Frontiers in Physiology* 8 (May 2017), p. 265. DOI: 10.3389/fphys.2017.00265.
- [31] L Pishchulin, S Wuhrer, T Helten, C Theobalt, and B Schiele. “Building Statistical Shape Spaces for 3D Human Modeling”. In: *Pattern Recognition* (2017).
- [32] M Potse. “Scalable and accurate ECG simulation for reaction-diffusion models of the human heart”. In: *Front. Physiol.* 9 (Apr. 2018), p. 370. DOI: 10.3389/fphys.2018.00370.
- [33] M Potse, B Dubé, J Richer, A Vinet, and RM Gulrajani. “A comparison of monodomain and bidomain reaction-diffusion models for action potential propagation in the human heart”. en. In: *IEEE Trans. Biomed. Eng.* 53.12 Pt 1 (Dec. 2006), pp. 2425–2435. DOI: 10.1109/TBME.2006.880875.
- [34] S Qian, D Ugurlu, E Fairweather, LD Toso, Y Deng, M Strocchi, L Cicci, RE Jones, H Zaidi, S Prasad, et al. “Developing cardiac digital twin populations powered by machine learning provides electrophysiological insights in conduction and repolarization”. In: *Nature Cardiovascular Research* 4.5 (2025), pp. 624–636.
- [35] C Ramanathan, RN Ghanem, P Jia, K Ryu, and Y Rudy. “Noninvasive electrocardiographic imaging for cardiac electrophysiology and arrhythmia”. In: *Nature medicine* 10.4 (2004), pp. 422–428.
- [36] C Rodero, M Strocchi, M Marciniak, S Longobardi, J Whitaker, MD O’Neill, K Gillette, C Augustin, G Plank, EJ Vigmond, et al. “Linking statistical shape models and simulated function in the healthy adult human heart”. In: *PLoS computational biology* 17.4 (2021), e1008851.
- [37] F Sahli Costabal, DE Hurtado, and E Kuhl. “Generating Purkinje networks in the human heart”. en. In: *J. Biomech.* 49.12 (Aug. 2016), pp. 2455–2465. DOI: 10.1016/j.jbiomech.2015.12.025.
- [38] J Sander, BD de Vos, S Bruns, N Planken, MA Viergever, T Leiner, and I Išgum. “Reconstruction and completion of high-resolution 3D cardiac shapes using anisotropic CMRI segmentations and continuous implicit neural representations”. In: *Computers in Biology and Medicine* 164 (2023), p. 107266.
- [39] W Schroeder, K Martin, and B Lorensen. *The Visualization Toolkit (4th ed.)* Kitware, 2006.
- [40] S Schuler and A Loewe. *Biventricular statistical shape model of the human heart adapted for computer simulations*. Zenodo, Jan. 2021. DOI: 10.5281/zenodo.4506463. URL: <https://doi.org/10.5281/zenodo.4506463>.
- [41] L Serrano, TX Wang, E Le Naour, JN Vittaut, and P Gallinari. “AROMA: Preserving Spatial Structure for Latent PDE Modeling with Local Neural Fields”. In: *Advances in Neural Information Processing Systems*. Ed. by A Globerson, L Mackey, D Belgrave, A Fan, U Paquet, J Tomczak, and C Zhang. Vol. 37. Curran Associates, Inc., 2024, pp. 13489–13521. DOI: 10.52202/079017-0431.
- [42] B Sullivan and A Kaszynski. “PyVista: 3D plotting and mesh analysis through a streamlined interface for the Visualization Toolkit (VTK)”. In: *Journal of Open Source Software* 4.37 (May 2019), p. 1450. DOI: 10.21105/joss.01450. URL: <https://doi.org/10.21105/joss.01450>.
- [43] M Tancik, PP Srinivasan, B Mildenhall, S Fridovich-Keil, N Raghavan, U Singhal, R Ramamoorthi, JT Barron, and R Ng. “Fourier features let networks learn high frequency functions in low dimensional domains”. In: *arXiv [cs.CV]* (June 2020).
- [44] KHWJ ten Tusscher and AV Panfilov. “Alternans and spiral breakup in a human ventricular tissue model”. en. In: *Am. J. Physiol. Heart Circ. Physiol.* 291.3 (Sept. 2006), H1088–100. DOI: 10.1152/ajpheart.00109.2006.

- [45] J Verhulsdonk, T Grandits, F Sahli Costabal, T Pinetz, R Krause, A Auricchio, G Haase, S Pezzuto, and A Effland. “Shape of my heart: Cardiac models through learned signed distance functions”. In: *Proceedings of The 7nd International Conference on Medical Imaging with Deep Learning*. Ed. by N Burgos, C Petitjean, M Vakalopoulou, S Christodoulidis, P Coupe, H Delingette, C Lartizien, and D Mateus. Vol. 250. Proceedings of Machine Learning Research. PMLR, 2024, pp. 1584–1605. arXiv: 2308.16568. URL: <https://proceedings.mlr.press/v250/verhulsdonk24a.html>.
- [46] A Verma, Cy Jiang, TR Betts, J Chen, I Deisenhofer, R Mantovan, L Macle, CA Morillo, W Haverkamp, R Weerasooriya, et al. “Approaches to catheter ablation for persistent atrial fibrillation”. In: *New England Journal of Medicine* 372.19 (2015), pp. 1812–1822.
- [47] J Vicente-Puig, J Chamorro-Servent, E Zacur, et al. “Volumetric non-invasive cardiac mapping for accessible global arrhythmia characterization”. en. In: *Commun. Med. (Lond.)* (2026). DOI: 10.1038/s43856-025-01332-5.
- [48] Y Wang and Y Rudy. “Application of the method of fundamental solutions to potential-based inverse electrocardiography”. In: *Annals of biomedical engineering* 34.8 (2006), pp. 1272–1288.
- [49] Y Xie, T Takikawa, S Saito, O Litany, S Yan, N Khan, F Tombari, J Tompkin, V Sitzmann, and S Sridhar. “Neural fields in visual computing and beyond”. en. In: *Comput. Graph. Forum* 41.2 (May 2022), pp. 641–676. DOI: 10.1111/cgf.14505.
- [50] E Zappon, MAF Gsell, K Gillette, and G Plank. “Quantifying anatomically-based in-silico electrocardiogram variability for cardiac digital twins”. In: *Comput. Biol. Med.* 189.109930 (Mar. 2025), p. 109930. DOI: 10.1016/j.combiomed.2025.109930.
- [51] E Zappon, M Salvador, R Piersanti, F Regazzoni, A Quarteroni, et al. “An integrated heart-torso electromechanical model for the simulation of electrophysiological outputs accounting for myocardial deformation”. In: *Computer Methods in Applied Mechanics and Engineering* 427 (2024), p. 117077.

MOX Technical Reports, last issues

Dipartimento di Matematica
Politecnico di Milano, Via Bonardi 9 - 20133 Milano (Italy)

- 25/2026** Carrara, D.; Hirschvogel, M.; Bonizzoni, F.; Pagani, S.; Pezzuto, S.; Regazzoni, F.
Shape-informed cardiac mechanics surrogates in data-scarce regimes via geometric encoding and generative augmentation
- 23/2026** Ballini, E.; Muscarnera, L.; Fumagalli, A.; Scotti, A.; Regazzoni, F.
Elimination-compensation pruning for fully-connected neural networks
- 21/2026** Bottacini, G.; Torzoni, M.; Manzoni, A.
Neural Markov chain Monte Carlo: Bayesian inversion via normalizing flows and variational autoencoders
- 17/2026** Caldera, L.; Bottacini, G.; Cavinato, L.
MAGIC-Flow: multiscale adaptive conditional flows for generation and interpretable classification
- 20/2026** Caldera, L.; Cavinato, L.; Cirone, A.; Cama, I.; Garbarino, S.; Lodi, R.; Tagliavini, F.; Nigri, A.; De Francesco, S.; Cappozzo, A.; Piana, M.; Ieva, F.;
DISARM++: Beyond scanner-free harmonization
- 19/2026** Caldera, L.; Cavinato, L.; Ieva, F.
Scanner-agnostic MRI harmonization via SSIM-guided disentanglement
- 15/2026** Zecchi, A. A.; Sanavio, C.; Cappelli, L.; Perotto, S.; Roggero, A.; Succi, S.
Block encoding of sparse matrices with a periodic diagonal structure
- 14/2026** Agasisti, T.; Cannistrà, M.; Paganoni, A.M.
Nudging communication for students at risk: experimental evidence from an Italian university
- 13/2026** Dimola, N.; Coclite, A.; Zunino, P.
Neural Preconditioning via Krylov Subspace Geometry
- 12/2026** Corbetta A; Logan K.M.; Ferro M; Zuccolo L; Perola M.; Ganna A.; Di Angelantonio E.;Ieva F.
Longitudinal patterns of statin adherence and factors associated with decline in over one million individuals in Finland and Italy