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

Arsenii Dokuchaev[1], Francesca Bonizzoni[2], Stefano Pagani[2], Francesco Regazzoni[2], Simone Pezzuto[1,3]

[1] Laboratory of Mathematics for Biology and Medicine, Università di Trento, Italy
[2] MOX Laboratory, Department of Mathematics, Politecnico di Milano, Italy
[3] Euler Institute, Università della Svizzera italiana, Switzerland

Abstract

Forward ECG models based on the lead-field method require an accurate torso representation, which is rarely available from clinical imaging, and their cost scales linearly with the number of electrodes. We propose a shape-informed neural surrogate for the lead-field gradient that removes both limitations. A geometry encoder (DeepSDF or PCA) maps heart-torso anatomy into a compact latent code, and a neural field predicts the lead-field gradient from spatial coordinates, electrode position and the latent code. Trained on 100 finite-element heart-torso geometries sampled from statistical shape atlases, the surrogate reaches a mean angular error below 4 degrees in the heart and a relative ECG error below 2%, clearly outperforming the classical pseudo lead-field approximation, while cutting per-lead inference cost by more than an order of magnitude relative to the finite element method.

1. Introduction

The lead-field method is the standard tool for fast forward simulation of the electrocardiogram (ECG). Lead fields do not depend on cardiac activation and can be precomputed once and reused for any excitation pattern, substantially reducing the cost of repeated ECG simulations [1, 2]. However, the computational cost of the lead-field method scales linearly with the number of electrodes, which becomes a limiting factor in applications like body-surface potential maps (BSPMs) where hundreds of electrodes are required. Another practical bottleneck is that the lead field, in addition to heart, requires a fully segmented torso geometry, which is rarely possible in heart-centered clinical protocols such as cardiac MRI.

We propose a shape informed surrogate model which approximates the lead field gradient as a continuous function of spatial coordinates, electrode position, and a compact encoding of the joint heart-torso anatomy. The model has two blocks: a geometry encoder that compresses the joint heart-torso shape into a low-dimensional shape vector, and the neural field model, which predicts the gradient of lead field at any point, conditioned by the shape vector and the electrode coordinates.

We compared two encoding strategies: a DeepSDF-based auto-decoder [3] and a linear PCA baseline and quantified model accuracy against FEM reference and the pseudo lead-field.

2. Methods

2.1. Forward modeling

For an electrode at $\mathbf{e}_j$, the corresponding lead-field function Z_j satisfies an adjoint elliptic problem over the complete heart-torso domain Ω.

$$\begin{cases} -\nabla \cdot (\mathbf{G}\nabla \mathbf{Z}) = \mathbf{0} & \text{on } \Omega \\ -\mathbf{G}\nabla \mathbf{Z_j} \cdot \mathbf{n} = \sum_{\mathbf{j}} \mathbf{a_j} \delta_{\mathbf{e_j}} & \text{on } \Sigma \end{cases}$$

where where $\mathbf{G} = \mathbf{G_i} + \mathbf{G_e}$ is the bulk conductivity, with intra- and extracellular conductivity tensors $\mathbf{G_i}, \mathbf{G_e}$ respectively and Σ is the body surface.

The ECG potential measured at the electrode e_j at time t can be written as

$$V(t) = \int_{\Omega_H} \nabla Z_j(\mathbf{x}) \mathbf{dx} \cdot \mathbf{G_i} \nabla \mathbf{V_m}(\mathbf{x}, \mathbf{t}).$$

Thus, the ECG depends on the gradient of lead field ∇Z inside the myocardium. Importantly, Z_j does not depend on time or on the particular transmembrane potential.

This observation motivates learning ∇Z_j rather than the Z_j itself. Once the lead field is available, different activation patterns can be evaluated without retraining the surrogate.

2.2. Training data

The entire preprocessing pipeline is depicted in the Fig. 1. We used the bi-ventricular statistical shape model proposed in [4] as the baseline heart model and the MPII Human Shape statistical atlas [5] as the baseline torso model.

To generate a virtual cohort we independently vary the first 10 principal components of both heart and torso shape models within the ranges ± 1 SD for heart and ± 2 SD for torso, using Latin hypercube sampling with uniform distribution (Fig. 1 block 1.). Three heart rotation angles along anatomical axes were also varied to capture inter-subject cardiac orientation variability (Fig. 1 block 2.).

Myocardial fibers were assigned using Laplace Dirichlet Rule based method [6]. This produced 100 training and 10 test virtual geometry sets.

For each joint geometry, a tetrahedral finite-element mesh was constructed using Gmsh [7] and 100 unipolar electrodes were uniformly distributed over the anterior torso surface. A standard 12-lead ECG electrodes were also defined (see Fig. 1 block 3.).

For each virtual patient, 2^{16} spatial points were sampled within computational domain, with sampling density increased near the heart-torso interface. The corresponding signed-distance functions and reference lead-fields were interpolated from the finite element representations onto these points and together with the coordinates of these points were used as training and test data for both DeepSDF auto-decoder and the geometry-conditioned neural surrogate.

2.3. Geometry encoding

Two geometry encoding methods were investigated, namely PCA and DeepSDF encoding.

For PCA encoding we used the same parameters that were used for joint geometries generation: 10 torso and 10 heart PCA weights, and 3 heart rotation angles in radians resulting in a 23-dimensional latent vector.

For DeepSDF encoding, we followed the [8] approach, in particular, by training auto-decoder model which represents the torso, LV and RV endocardium and epicardium using shared 16-dimensional latent code. Neural network was composed of 5 linear fully-connected Lipschitz-normalized layers, with 256 neurons each. Fourier feature encoding [9] was applied to improve the representation of high-frequency geometric details and ReLU activations and the Adam optimizer were used during training. Latent codes for new anatomies were inferred by MAP optimization with a Mahalanobis prior.

2.4. Geometry-conditioned lead-field surrogate

The neural surrogate was modeled by fully connected network with 5 linear layers with 256 neurons each, taking as input spatial coordinates, normalized electrode position and the latent code and outputting 3 components of ∇Z vector. Spatial coordinates were also encoded using Fourier features encoding. The network is trained using a combination of vector-valued mean-squared error and cosine similarity $Loss = \mathcal{L}_{\text{MSE}} + \lambda \mathcal{L}_{\cos}$ where the first term controls the magnitude of the predicted gradient and the second explicitly penalizes errors in its direction. Training used Adam, 800 epochs and ≈ 2 GPU-days on an NVIDIA A100.

To evaluate the practical impact of lead-field errors, two activation patterns were considered: sinus rhythm with a rule-based Purkinje activation model and a pathological activation mimicking left bundle branch block (LBBB). The resulting transmembrane potentials were inserted into the lead-field formulation using either FEM or the predicted lead fields.

3. Results

The DeepSDF decoder reconstructs the four anatomical surfaces with sub-mm mean Chamfer distance (0.75-1.28 mm) and comparable train/test error, indicating good generalization of the latent space to unseen geometries.

Encoding	LF angle [deg]	ECG-lead angle [deg]	ECG error [%]
PCA	5.22 ± 0.61	5.93 ± 1.26	2.4 ± 1.3
DeepSDF	$\mathbf{3.89 \pm 0.51}$	$\mathbf{4.64 \pm 0.99}$	$\mathbf{1.8 \pm 1.0}$

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.

More importantly, the latent representation retained the geometric information required for the lead-field prediction task.

The neural surrogate reproduced the FEM lead-field gradients throughout the torso and inside the myocardium. As summarized in Table 1, the DeepSDF representation consistently outperforms PCA at both gradient and ECG level. The median angular error over the full torso was 2.64 degrees for DeepSDF compared with 3.51 degrees for PCA. Magnitude errors were similar for the two representations, suggesting that the main advantage of DeepSDF was improved directional alignment of the gradient field.

The largest angular errors occurred close to the heart, especially on precordial electrodes (Fig. 3). These regions

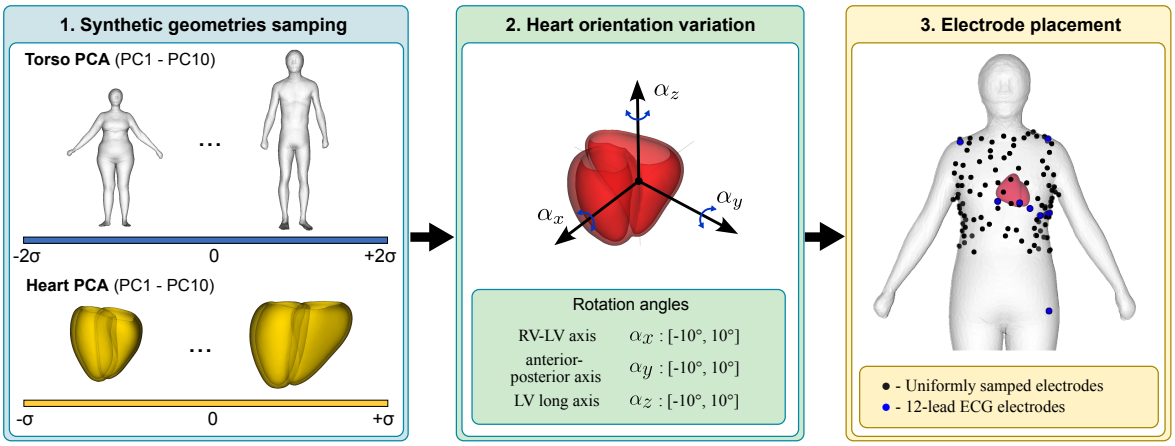

Figure 1. Graphical representation of the developed pipeline. (1.) Mesh generation stage: we used torso and heart models from a statistical PCA atlas. To create joint models, we varied the first 10 principal components of the torso, the first 10 principal components of the heart, and (2.) heart rotation angles along LV–RV axis (α_x), anterior-posterior axis (α_y) and LV long axis (α_z) anatomical axes. (3.) Variation of unipolar electrode location on the anterior surface of the torso.

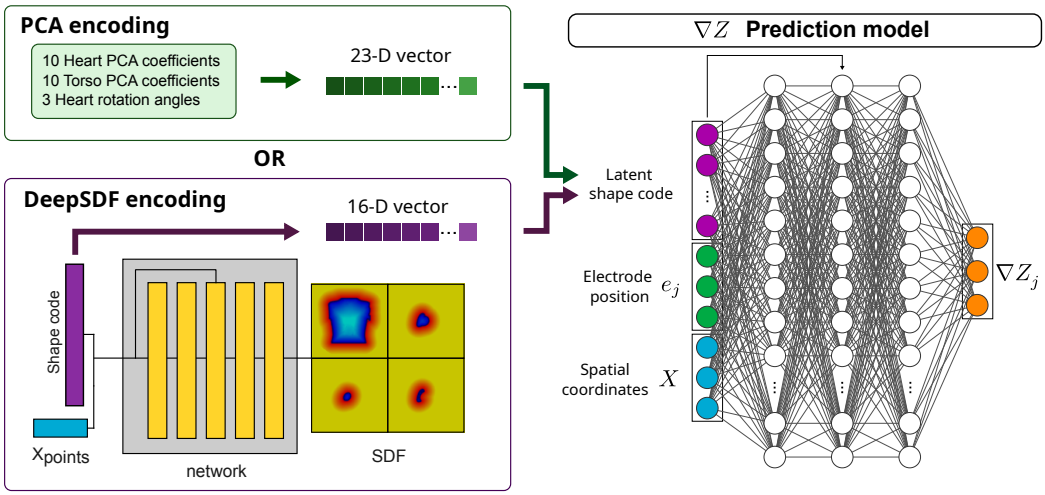

Figure 2. Schematic representation of the ∇Z prediction framework with PCA- and DeepSDF-based geometry encodings

correspond to rapid spatial variations of the lead field and are therefore the most challenging part of the domain to approximate.

For the nine independent electrodes of the standard 12-lead ECG, the angular error was $4.64 \pm 0.99°$ for DeepSDF and $5.93 \pm 1.26°$ for PCA. The improvement was particularly relevant for the precordial leads, where the torso geometry has a strong influence on the lead field.

Despite localized errors near the heart-torso interface, ECG waveforms obtained from the predicted lead fields closely matched the FEM reference for both tested sinus rhythm and LBBB activation patterns.

Concerning computational efficiency, on CPU, a FEM lead-field solve required approximately 6s per lead, excluding mesh generation. In comparison, evaluation of the trained neural surrogate required approximately 250ms per lead. Thus, the proposed method provided a 24-fold speedup without GPU acceleration or batching.

The advantage becomes increasingly relevant as the number of electrodes grows. For dense BSPM configurations, the surrogate replaces repeated PDE solves with independent neural-field evaluations.

4. Discussion and conclusion

We presented a shape-informed neural surrogate for the geometry-dependent lead-field gradient, trained on FEM solutions over a population of statistical heart-torso anatomies. On test geometries, the surrogate achieves a mean angular error below 4 degrees and a relative ECG error below 2%, outperforming the pseudo lead-field approx-

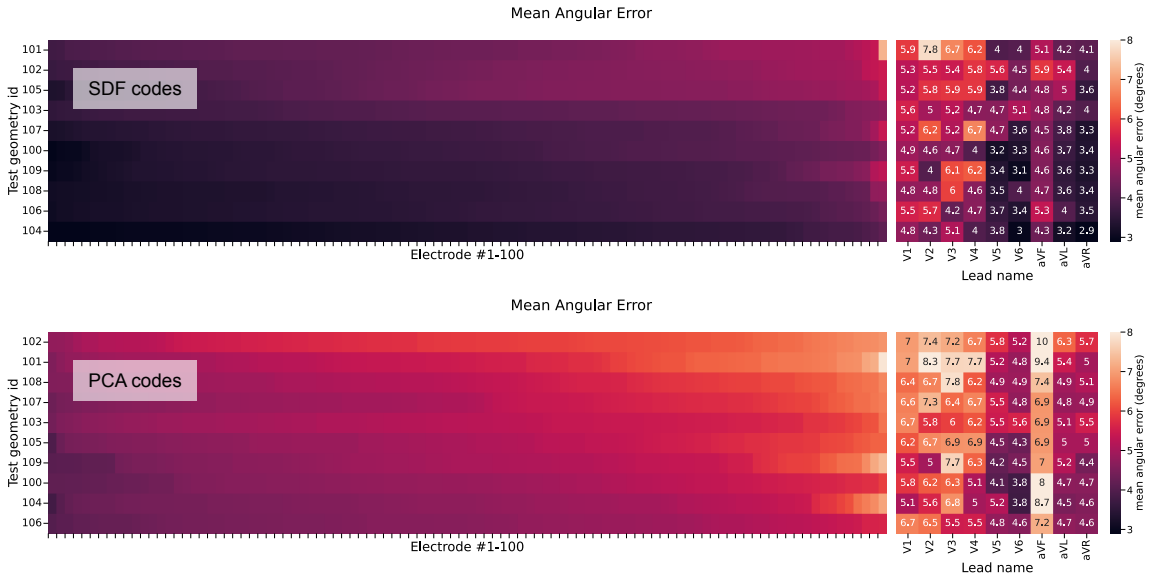

Figure 3. Average angular error (degrees) in predicting ∇Z using DeepSDF-based (top-row) and PCA-based (bottom row) representations. Left column indicate errors across 100 uniformly distributed unipolar electrodes. Right column: errors across the 9 electrodes of standard 12-lead ECG. Electrodes and patients are sorted by increasing errors.

imation while reducing per-lead inference time by more than 24 times relative to FEM. The DeepSDF encoding consistently outperforms the PCA baseline, particularly for precordial leads near the heart-torso interface, and requires only a compact latent shape code at inference time — making the approach well suited to cardiac digital twins, ECG inverse problems, and settings with repeated forward solves under varying electrode configurations.

Acknowledgments

This project has received funding from the European Research Council (ERC) under the European Union's Horizon Europe research and innovation programme (Grant agreement No. 101230597, awarded to SP). We also acknowledge support from the PRIN2022 project (MUR, Italy, 2023–2025, no. P2022N5ZNP) "SIDDMs", funded by the European Union (NextGenerationEU, Mission 4, Component 2); the CSCS Swiss National Supercomputing Centre (project no. lp157); and the SNSF–FWF project "CardioTwin" (no. 214817)

References

[1] Geselowitz DB. On the theory of the electrocardiogram. Proc IEEE Inst Electr Electron Eng 06 1989;77(6):857–876.

[2] Potse M. Scalable and accurate ECG simulation for reaction-diffusion models of the human heart. Front Physiol 04 2018; 9:370.

[3] Park JJ, Florence P, Straub J, Newcombe R, Lovegrove S. Deepsdf: Learning continuous signed distance functions for shape representation. In Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. 2019; 165–174.

[4] Schuler S, Loewe A. Biventricular statistical shape model of the human heart adapted for computer simulations, January 2021. URL https://doi.org/10.5281/zenodo.4506463.

[5] Pishchulin L, Wuhrer S, Helten T, Theobalt C, Schiele B. Building statistical shape spaces for 3d human modeling. Pattern Recognition 2017;.

[6] Bayer JD, Blake RC, Plank G, Trayanova NA. A novel rule-based algorithm for assigning myocardial fiber orientation to computational heart models. Ann Biomed Eng October 2012;40(10):2243–2254.

[7] Geuzaine C, Remacle JF. Gmsh: a three-dimensional finite element mesh generator with built-in pre- and post-processing facilities. Int J Numer Methods Eng 09 2009; 79(11):1309–1331.

[8] Verhülsdonk J, Grandits T, Costabal FS, Pinetz T, Krause R, Auricchio A, Haase G, Pezzuto S, Effland A. Shape of my heart: Cardiac models through learned signed distance functions. 2023; .

[9] Tancik M, Srinivasan PP, Mildenhall B, Fridovich-Keil S, Raghavan N, Singhal U, Ramamoorthi R, Barron JT, Ng R. Fourier features let networks learn high frequency functions in low dimensional domains. arXiv csCV 06 2020;.

Address for correspondence:

Arsenii Dokuchaev
Trento, Italy
arsenii.dokuchaev@unitn.it, zodelheim@gmail.com