Generalized Implicit 12-Lead ECG Reconstruction

Nora Vogt[1], Mattias P Heinrich[2], Jacques Felblinger[1,3], Freddy Odille[1,3]

[1]IADI U1254, Inserm, Université de Lorraine, Nancy, France
[2]Institute of Medical Informatics, University of Luebeck, Luebeck, Germany
[3]CIC-IT 1433, Inserm, Université de Lorraine, CHRU de Nancy, Nancy, France

Abstract

Twelve-lead ECGs are essential for cardiac diagnosis, yet standard electrode placements are often impractical, particularly in MR environments. We propose two reference-free methods for reconstructing 12-lead ECGs from arbitrary sensor configurations: a generalized implicit neural representation (INR) network, and, as a baseline, a linear matrix solution with Laplacian smoothness regularization. The INR is a multilayer perceptron with sinusoidal activations and modulation layers that maps torso surface coordinates, temporal encodings, and subject-specific latent vectors to body surface potential maps (BSPMs), from which standard leads are extracted. Training jointly optimizes for accurate 12-lead reconstruction and BSPM smoothness. At inference, subject-specific latent vectors are optimized to match observed leads (e.g., [III, V1, V5] or MR sensor leads) using a frozen or fine-tuned model. We evaluated both methods on PTB-XL+ normal subjects (8,151 training; 914 test) and an in-house dataset of seven healthy subjects with standard 12-lead ECGs and 12 bipolar MR sensor leads. The INR achieved mean cross-correlation scores of 0.90 versus 0.87 (baseline) on PTB-XL+ and 0.94 versus 0.87 on the in-house dataset, using latent-vector optimization alone. Adapting the model to MR sensor inputs required additional fine-tuning and stronger smoothness regularization, achieving performance comparable to the matrix solution (0.80 vs. 0.79). Further gains are expected from personalized torso geometry and sensor locations.

1. Introduction

Cardiac diagnosis relies on the standardized 12-lead ECG, which uses 10 electrodes placed at predefined positions on the arms, legs, and chest. MRI-compatible systems typically employ reduced lead sets or short electrode cables to minimize heating, enabling cardiac monitoring with signals that differ from diagnostic ECG. Prior work [1] demonstrated that a data-driven reconstruction method can derive accurate 12-lead ECGs from an MR-compatible

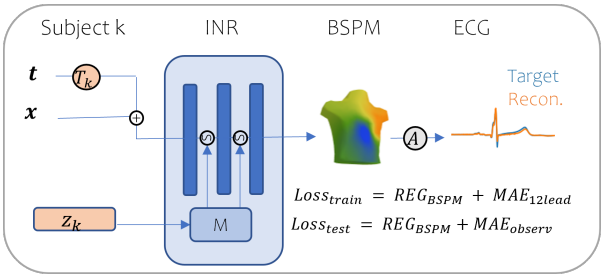

Figure 1: Overview of the proposed generalized implicit neural network. The model maps temporal encodings $\mathbf{t}$, atlas torso surface coordinates $\mathbf{x}$, and subject-specific latent codes $\mathbf{z_k}$ to dense body surface potential maps (BSPMs), from which ECGs are extracted via matrix multiplication.

sensor arrangement consisting of four groups with three bipolar lead measurements each, though requiring calibration using a standard 12-lead ECG acquired outside the MRI environment. While some studies have explored reconstructions from unconventional lead placements (e.g., patch-based devices or wearable ECGs), most related work focuses on reconstructing the 12-lead ECG from subsets of standard leads [2]. Conventional reconstruction methods are typically linear and subject-specific (e.g., matrix-based approaches), whereas machine learning solutions often employ nonlinear models trained on generic population datasets (e.g., long short-term memory networks, convolutional neural networks, or generative adversarial networks)[2]. These approaches are limited by fixed input lead selections and potential poor generalization to external data, motivating patient-specific adaptation to reduce reconstruction errors [3]. ECG view synthesis using implicit neural representations (INR) was introduced by Chen et al. [4] and extended by Zhan et al. [5] for arbitrary-length signal synthesis and electrode placement deviation estimation. However, both approaches represented leads in spherical coordinates, expressing unipolar precordial leads through polar and azimuth angles, limiting applicability to bipolar lead acquisitions. We therefore propose a generalized coordinate-based implicit neural representation (INR)

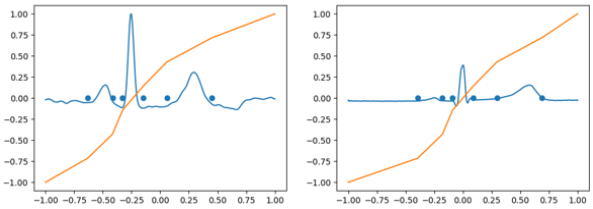

Figure 2: Temporal encoding for two examples. Landmark estimates are indicated by blue dots and resulting temporal encoding are shown in orange.

model that learns a continuous representation of body surface potentials. Inspired by the coordinate-based, latent-conditioned INR framework of Dannecker et al. [6], originally developed for medical image segmentation, we adapt this approach to ECG mapping. Given 3D torso surface coordinates, the model predicts dense body surface potential maps (BSPMs), from which desired ECG leads are extracted by matrix multiplication. While trained on 12-lead ECG reconstruction, the model predicts dense body surface potentials that can be combined through lead extraction matrices to support arbitrary input coordinate and output lead configurations. Subject-specific adaptation is achieved at inference through optimization of latent vectors, with optional model fine-tuning when required.

2. Methods

2.1. The linear matrix solution

As a baseline, we solve 12-lead ECG reconstruction via regularized linear inversion. Given measured leads $\mathbf{M}$, the BSPM reconstruction matrix $\mathbf{R}$ is computed as:

$$\mathbf{R} = \left(\mathbf{T}^\top \mathbf{T} + \lambda \mathbf{L}^\top \mathbf{L}\right)^+ \mathbf{T}^\top, \quad \lambda = \mathbf{1} \tag{1}$$

where $\mathbf{T}$ is the lead extraction matrix for the measured ECG, $\mathbf{L}$ is the Laplacian for smoothness regularization (accounting for torso node distances), and $(\cdot)^+$ denotes pseudoinverse. The 12-lead ECG is then $\mathbf{M_{12}} = \mathbf{A_{12}RM}$, with $\mathbf{A_{12}}$ mapping torso potentials to standard leads.

2.2. The proposed model

The multi-layer perceptron (MLP), visualized in Fig. 1, learns a mapping from 3D atlas torso surface locations $\mathbf{x}$, temporal encodings $\mathbf{t}$, and trainable subject-specific latent codes $\mathbf{z}$ to dense body surface potential maps. Following [6], we use modulated layers conditioned on the latent code $\mathbf{z}$: the scale (ϕ) and shift (ψ) parameters are derived by passing $\mathbf{z}$ through a linear layer, i.e., $(\phi, \psi) = \mathbf{M}\mathbf{z} + \mu$. The output of a modulated layer is then given by $\sin\big(\omega_0 \cdot \phi \cdot (\mathbf{W}\mathbf{x} + \mathbf{b}) + \psi\big)$. We query dense BSPMs at

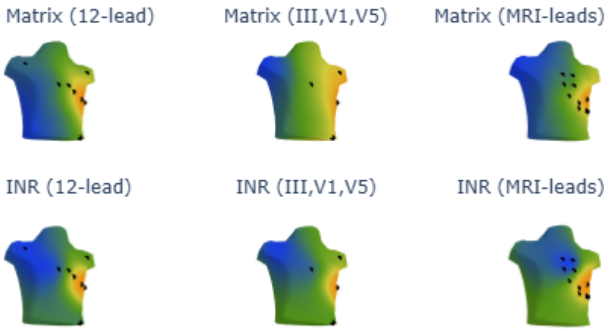

Figure 3: Example BSPM reconstructions with observed electrode locations indicated in black. Top: matrix solution; bottom: INR prediction. From left to right: standard electrodes, lead subset [III, V1, V5], and MR sensors.

the torso surface coordinates of a shared atlas and extract the desired leads via multiplication with a lead extraction matrix $\mathbf{A}$. During training, the model is optimized to accurately reconstruct the 12-lead ECGs of the training samples while encouraging smooth body surface maps, using the combined loss:

$$L_{total} = L_{ECG}^{MAE} + \lambda_1 L_{BSPM}^{Laplacian} + \lambda_2 L_{BSPM}^{Init} + \lambda_3 L_{Latent}^{L1}, \tag{2}$$

where L_{ECG}^{MAE} is the mean absolute error between the target and predicted leads (e.g. 12 leads during the training phase), $L_{BSPM}^{Laplacian}$ is a Laplacian smoothness regularization of the BSPM, L_{BSPM}^{Init} is an additional regularization term encouraging BSPM estimates to remain close to the linear matrix solution, and L_{Latent}^{L1} is an L1 regularization on the latent space, encouraging a compact representation. To capture cardiac-cycle features, a subject-specific temporal encoder T_k estimates six landmarks for each subject k, transforming temporal coordinates $\mathbf{t}$ into cardiac phase-aware encodings (Fig. 2). These landmarks are initialized from the training-set-average onset and offset locations of the P wave, QRS complex, and T wave, parameterized as positive increments to ensure monotonic ordering. Given the landmark positions $\{b_i\}_{i=1}^{8}$ and corresponding encoding values $\{y_i\}_{i=1}^{8}$ (evenly distributed over $[-1, 1]$), piecewise linear encodings $t_{\text{enc}} \in [-1, 1]$ are obtained by interpolating between adjacent landmarks:

$$t_{\text{enc}} = y_i + \frac{t - b_i}{b_{i+1} - b_i}\,(y_{i+1} - y_i) \tag{3}$$

where $b_i \le t < b_{i+1}$. No explicit supervision of these landmarks is performed during training. At test time, new subject-specific latent vectors are tuned to reconstruct observed leads (in our experiments, either the lead subset [III, V1, V5] or a set of MR sensor leads) using a frozen or fine-tuned INR model.

Table 1: 12-lead ECG reconstruction performances

Dataset	Cross-correlation (limb, chest)
PTB-XL+ (III, V1, V5)	Matrix: **0.87** (0.86, 0.88)
	INR: **0.90** (0.91, 0.88)
In-house (III, V1, V5)	Matrix: **0.87** (0.85, 0.89)
	INR: **0.94** (0.94, 0.94)
In-house (MR sensors)	Matrix: **0.79** (0.73, 0.85)
	INR: **0.80** (0.73, 0.87)

2.3. Datasets

We considered a fixed atlas torso mesh (PhysioNet 2007 Challenge [7]), comprising 352 torso nodes with annotated standard electrode locations (limb leads placed on the torso). The INR model was trained on the PTB-XL+ [8] dataset, which provides median beat templates for the 12-lead ECGs of the PTB-XL dataset [9] (University of Glasgow). We used the recommended data splits for model evaluation, including only samples with the normal diagnostic superclass (8,151 train and 914 test). We further evaluated the model on in-house data [1], comprising seven healthy subjects for whom both standard 12-lead ECGs and 12 bipolar MR sensor leads were acquired outside the MR environment (ClinicalTrials.gov identifier: NCT02887053). For the in-house dataset, we derived single-beat windows by first bandpass filtering the raw time sequences (4th order, 0.05 Hz-150 Hz), then detecting R-peaks and extracting a median-R-peak-centered window, and finally resampling to a 500 Hz temporal resolution, yielding windows of size 512. For evaluation purposes, the 12-lead and MR-sensor-lead windows were automatically aligned using a search space of [-40, 40] samples to minimize the normalized root mean squared error (NRMSE) between the 12-lead signals and the estimated 12-lead signals provided by the linear matrix solution.

2.4. Experimental setup

The network was configured with three hidden layers of size 128 ($\omega = 20$), modulation layers of size 256 (after the first and third hidden layers), and a latent code size of $z = 128$ (initialized as $\mathbf{z} \sim \mathcal{N}(\mathbf{0}, \mathbf{10^{-2}})$). The model was trained for 2 epochs using a batch size of 1024 (full torso grid, random t) and the AdamW optimizer, with cosine learning-rate scheduling and initial learning rates of $1e{-}4$ for the INR, $5e{-}4$ for the latent codes, and $5e{-}4$ for the temporal encoding. Loss function weights were $\lambda_1 = 10$, $\lambda_2 = 0.001$, and $\lambda_3 = 0$ during training. Test-time tuning varied across the three datasets. For the PTB-XL+ test set, we kept the same hyperparameters as during training, except for enabling latent-space regularization ($\lambda_3 = 10$). Fine-tuning was performed jointly across all samples for 20 epochs (batch size 1024), freezing the INR model and tuning only the subject-specific temporal encodings and latent codes; the latter were initialized using the latent code of the closest training sample, measured by Pearson correlation over the observed lead subset. For the in-house 12-lead ECG dataset, latent codes and temporal encodings were tuned per subject for 800 epochs (batch size 512, corresponding to all timestamps), with all other hyperparameters as for PTB-XL+ test. Given the domain shift of the MR-sensor lead dataset, we applied low-rank adaptation (LoRA) [10] to better fit the observed leads and prevent underfitting from tuning latent codes and temporal encodings alone. Here, parameters were trained for 1000 epochs (batch size 512) using LoRA finetuning (r=8, alpha=8) with a learning rate of $1e{-}5$ and an increased BSPM prior consistency weight ($\lambda_2 = 0.01$) to reduce overfitting. Latent codes were initialized using the latent code of the closest PTB-XL+ training sample, with simulated training sample MR-sensor signals derived from the fitted model.

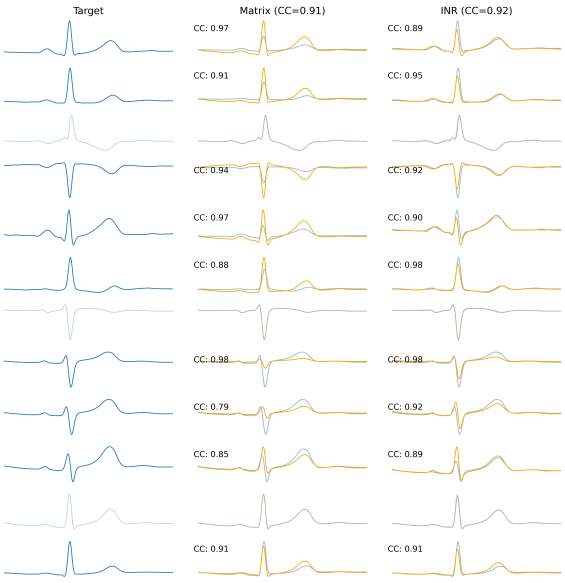

Figure 4: Example PTB-XL+ 12-lead prediction for a sample with median performance (matrix solution CC=0.91, INR CC=0.92).

3. Results and Discussion

Cross-correlation (CC) performance is summarized in Table 1. For the sub-lead experiments, the INR model outperformed the linear solution, achieving average cross-correlation scores of 0.90 vs. 0.87 and 0.94 vs. 0.87 for

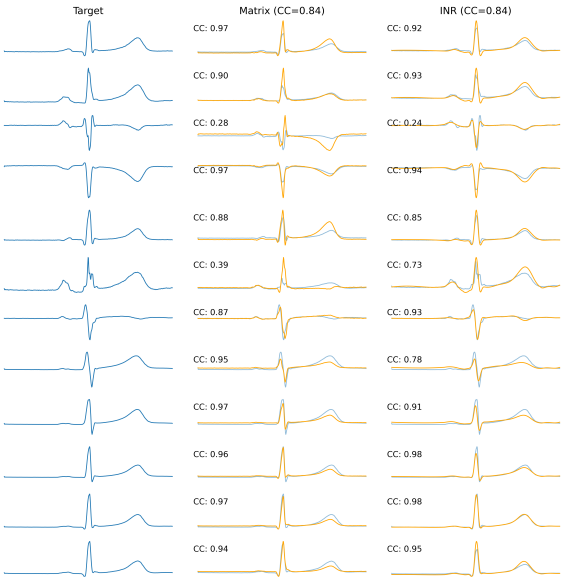

Figure 5: Example MR sensor 12-lead prediction for a sample with medium performance (matrix solution CC=0.84, INR CC=84).

PTB-XL+ and the in-house 12-lead ECG dataset, respectively. Further improvements are especially needed for adaptation to MR-compatible sensor placements without standard limb leads. Despite model finetuning, results remained close to those of the linear matrix solution (0.80 vs. 0.79). Example BSPM predictions and 12-lead ECG reconstructions are visualized in Fig. 3 and Figs. 4–5, respectively. A limitation of the proposed model is the use of a shared atlas torso with predefined electrode locations rather than subject-specific geometries. This is particularly relevant for limb leads, which are currently modeled using torso-based electrode positions, whereas PTB-XL recordings and in-house 12-lead ECG used standard clinical placements on the limbs. Further improvements may be achieved by incorporating personalized torso models learned from paired MRI and ECG data, together with improved electrode localization. Additional limitations include the restriction to single-beat reconstruction and evaluation exclusively in healthy subjects. Future extensions toward continuous ECG reconstruction will require temporal encoding mechanisms capable of modeling long-range dynamics.

4. Conclusion

Further improvements are needed to enhance MR-sensor reconstructions, particularly for the limb leads. Future work will focus on personalized torso models, advanced electrode localization, and evaluating model performance under magnetohydrodynamic distortions present in in-MR acquisitions.

Acknowledgments

This research was funded by the ANR grant SWEET-HEART (ANR-23-CE19-0030) from the national call on ”Technologies pour la santé”. It was further supported by the Programme Hubert Curien (PHC) Procope 2025 (MICROVOLT 53020YJ) and the CIC-IT of Nancy.

References

[1] Dos Reis JE, Soullié P, Oster J, Palmero Soler E, Petitmangin G, Felblinger J, Odille F. Reconstruction of the 12-lead ECG using a novel mr-compatible ECG sensor network. Magnetic resonance in medicine 2019;82(5):1929–1945.

[2] Obianom EN, Ng GA, Li X. Reconstruction of 12-lead ECG: a review of algorithms. Frontiers in Physiology 2025; 16:1532284.

[3] Obianom EN, Koya A, Feng F, Ng GA, Li X. Transfer learning is the electrocardiogram reconstruction capstone. Computers in Biology and Medicine 2026;207:111644.

[4] Chen J, Zheng X, Yu H, Chen DZ, Wu J. Electrocardio panorama: synthesizing new ECG views with self-supervision. arXiv preprint arXiv210506293 2021;.

[5] Zhan Z, Hu Y, Zhan J, Lian W, Wu W, Chen J. NEF-NET+: Adapting electrocardio panorama in the wild. arXiv preprint arXiv251102880 2025;.

[6] Dannecker M, Kyriakopoulou V, Cordero-Grande L, Price AN, Hajnal JV, Rueckert D. Cina: Conditional implicit neural atlas for spatio-temporal representation of fetal brains. In International Conference on Medical Image Computing and Computer-Assisted Intervention. Springer, 2024; 181–191.

[7] Goldberger AL, Amaral LA, Glass L, Hausdorff JM, Ivanov PC, Mark RG, Mietus JE, Moody GB, Peng CK, Stanley HE. Physiobank, physiotoolkit, and physionet: components of a new research resource for complex physiologic signals. circulation 2000;101(23):e215–e220.

[8] Strodthoff N, Mehari T, Nagel C, Aston PJ, Sundar A, Graff C, Kanters JK, Haverkamp W, Dössel O, Loewe A, et al. PTB-XL+, a comprehensive electrocardiographic feature dataset. Scientific data 2023;10(1):279.

[9] Wagner P, Strodthoff N, Bousseljot RD, Kreiseler D, Lunze FI, Samek W, Schaeffter T. Ptb-xl, a large publicly available electrocardiography dataset. Scientific data 2020;7(1):154.

[10] Hu EJ, Shen Y, Wallis P, Allen-Zhu Z, Li Y, Wang S, Wang L, Chen W, et al. Lora: Low-rank adaptation of large language models. Iclr 2022;1(2):3.

Address for correspondence:

Nora Vogt
IADI U1254, Bâtiment Recherche, CHRU de Nancy Brabois, Rue du Morvan, 54511 Nancy, France
nora.vogt@univ-lorraine.fr