A Neural Network Architecture for ECG Lead Reconstruction: Separating Shared and Lead-Specific ECG Characteristics

Mohammadsina Hassannia[1], Reza Sameni[1,2]

[1] Emory University, Atlanta, GA, USA.
[2] Georgia Institute of Technology, Atlanta, GA, USA

Abstract

We present a modular neural architecture for multilead ECG lead decomposition that separates inter-lead dependencies of the eight independently collected leads (I, II, V1, V2, V3, V4, V5, and V6), in three successive layers. The first layer generalizes the Dower transform via a linear least squares (LS) approach to approximate each lead from the others. The second layer employs a multilayer perceptron (MLP) to capture nonlinear residuals not captured by the LS layer. Finally, a recurrent neural network (RNN) is used to model temporal structure in the MLP residuals. These layers are additive, with each layer refining the reconstruction of the target lead. Trained and evaluated on the PTB-XL dataset, the model shows consistent improvements in lead reconstruction signal-to-noise ratio (SNR) across all layers. The best reconstruction was achieved for the precordial leads, with V5 showing the highest average SNR of 16.2 dB. The most difficult lead for reconstruction was lead II with lowest average SNR of 7.3 dB. The architecture can serve as a building block for explainable ECG foundation models, automatic diagnostic pipelines, feature encoders, or compression, with applications in noisy/missing lead recovery and low-dimensional ECG reconstruction.

1. Introduction

The 12-lead ECG, the clinical standard for cardiac diagnosis, captures the cardiac electrical activity from different anatomical landmarks, enabling a multi-view of the heart and detection of abnormalities like ischemia and arrhythmia [1–4]. However, the complex interplay between *shared* and *lead-specific* components in multi-lead ECGs remains underexplored. The similarities and differences arise partly from the common cardiac rhythm, while the leads view the heart from different perspectives, resulting in both redundant and distinct patterns across the leads [5, 6]. Hypothetically, disentangling the *shared* information from the *unique*, lead-specific patterns could enable more efficient and interpretable representations for building ECG-specific deep learning architectures, algorithmic diagnosis, waveform recovery, and signal compression.

In this work, we introduce a compact, modular neural network architecture for ECG lead reconstruction, using a three-layer cascade pipeline that decomposes shared and lead-specific information while preserving interpretability. The first layer employs a linear least-squares (LS) model, extending the Dower transform [7], to capture dominant linear correlations across leads. The second layer is a multilayer perceptron (MLP) that estimates nonlinear residuals, modeling complex anatomical and physiological relationships not accounted for by the LS layer. Finally, the third layer applies a recurrent neural network (RNN) to capture temporal dependencies in the residual signal, refining beat-to-beat variations and waveform dynamics. The additive structure of these layers isolates linear, nonlinear, and temporal contributions. The model is trained and evaluated on the PTB-XL dataset [8], which contains 21,837 clinically annotated 12-lead ECGs with different clinical conditions.

Contrasting recent work on lead imputation or generative ECG modeling [6], we focus on both reconstruction accuracy and interpretability through signal decomposition. The results demonstrate consistent improvements in lead reconstruction per layer, confirming the effectiveness of the proposed architecture, as modular blocks for designing explainable foundation models for ECG analysis and algorithmic diagnosis.

2. Methods

An overview of the proposed architecture is shown in Figure 1. The input ECGs are first preprocessed using a zero-phase forward-backward bandpass filter (0.5–80 Hz) to ensure the model is not impacted by baseline wander or high-frequency noises. The preprocessed signals are used to train a three-layer neural network model, where a target ECG lead is reconstructed from the remaining leads. Each layer is trained to progressively refine the estimate of the previous layer by modeling distinct types of inter-lead dependencies: linear spatial (across leads), nonlinear spatial

(across leads), and spatiotemporal (across time and lead). While the pipeline is applicable to arbitrary lead combinations, we focus on the independently collected limb leads (leads I and II) and the six precordial leads (V1, V2, V3, V4, V5, and V6).

2.1. Linear Least Squares (LS)

Let L denote the number of leads (8 in this work). The first layer uses linear least squares (LS) regression to estimate a target lead y_t as a weighted combination of the remaining $L-1$ leads, which we denote as $\mathbf{x}_t \in \mathbb{R}^{(L-1)\times 1}$. To accomplish this, we solve for $\boldsymbol{\theta}_{\text{LS}} \in \mathbb{R}^{(L-1)\times 1}$ that minimizes

$$\begin{aligned} \boldsymbol{\theta}_{\text{LS}} &= \arg\min_{\boldsymbol{\theta}} \mathbb{E}_t\{(y_t - \boldsymbol{\theta}^\top \mathbf{x}_t)^2\} \\ \hat{y}_t^{(1)} &= \boldsymbol{\theta}_{\text{LS}}^\top \mathbf{x}_t \end{aligned} \tag{1}$$

where $\mathbb{E}_t\{\cdot\}$ denotes averaging over all time indexes t, and $\hat{y}_t^{(1)}$ represents an estimate of the target lead y_t at time t. This layer generalizes the classical Dower Transform [7], which maps vectorcardiographic Frank leads to the 8- or 12-lead system. The layer is designed to capture dominant linear correlations, serving as a first-order linear approximation for the shared information between leads. The weights obtained via LS are subsequently used to initialize a multilayer perceptron (MLP) with an activation function of the form

$$y = \alpha \tanh(x/\alpha) \tag{2}$$

where α is set to a large value greater than the typical ECG amplitude range (10 mV in our results), ensuring that this layer behaves similarly to a linear activation function for normal ECGs while saturating out-of-range ECGs to prevent outliers from biasing model training. The parameter α remains a hyperparameter for further optimization during fine-tuning.

2.2. Nonlinear Least Squares with Multilayer Perceptron (MLP)

While the LS approximation in Layer 1 captures coarse linear inter-lead relationship, it cannot model nonlinear relationships. To address this, the second layer of the proposed architecture is a MLP with a scaled tanh activation similar to (2). For each time point t, the MLP uses $\mathbf{x}_t \in \mathbb{R}^{(L-1)\times 1}$ (i.e., the values of the other leads at the same time step) and estimates a nonlinear correction to the residual $r_t^{(1)} = y_t - \hat{y}_t^{(1)}$ from the first layer. Formally, the MLP learns a function $f_{\text{MLP}}(\cdot) : \mathbb{R}^{L-1} \to \mathbb{R}$ that minimizes the following objective:

$$\begin{aligned} \boldsymbol{\theta}_{\text{MLP}} &= \arg\min_{\theta} \mathbb{E}_t \left\{ \left(y_t - \hat{y}_t^{(1)} - f_{\text{MLP}}(\mathbf{x}_t; \boldsymbol{\theta}) \right)^2 \right\} \\ \hat{y}_t^{(2)} &= \hat{y}_t^{(1)} + f_{\text{MLP}}(\mathbf{x}_t; \boldsymbol{\theta}_{\text{MLP}}) \end{aligned} \tag{3}$$

where $\boldsymbol{\theta}_{\text{MLP}}$ are the MLP parameters. This formulation is structurally analogous to the linear least squares step in Layer 1, but with the added capacity to model nonlinear spatial dependencies between the leads. The output of this layer is an updated estimate $\hat{y}_t^{(2)}$ that incorporates nonlinear refinements to the initial approximation.

2.3. Temporal Modeling using Recurrent Neural Network (RNN)

The third layer models temporal dependencies in $r_t^{(2)} = y_t - \hat{y}_t^{(2)}$, the residual signal that remain after the spatial modeling in Layer 2. To capture temporal dependencies, we employ a recurrent neural network (RNN) that learns to predict a residual correction term based on a temporal window around each time step. While the temporal window may generally be causal or non-causal, we use a non-causal window here: for each target time step t, a window of $2s+1$ input signal $\mathbf{x}_{t-s:t+s} \in \mathbb{R}^{(L-1)\times(2s+1)}$ is extracted and passed to the RNN ($s = 99$ in our results). The network uses this non-causal (bidirectional) window to estimate a correction term for the center point of the window, refining the lead reconstruction at time t. This step captures local temporal dynamics while avoiding information leakage from global sequence-level patterns, beyond the window $2s+1$. Importantly, while the LS and MLP layers were instantaneous and agnostic to the sampling frequency, the RNN layer is sampling rate dependent. Formally, the RNN learns a function $f_{\text{RNN}}(\cdot)$ with parameters $\boldsymbol{\theta}_{\text{RNN}}$ that minimizes the temporal residual loss:

$$\begin{aligned} \boldsymbol{\theta}_{\text{RNN}} &= \arg\min_{\theta} \mathbb{E}_t \left\{ (y_t - \hat{y}_t^{(2)} - f_{\text{RNN}}(\mathbf{x}_{t-s:t+s}; \boldsymbol{\theta}))^2 \right\} \\ \hat{y}_t^{(3)} &= \hat{y}_t^{(2)} + f_{\text{RNN}}(\mathbf{x}_{t-s:t+s}; \boldsymbol{\theta}_{\text{RNN}}) \end{aligned} \tag{4}$$

Combining (1), (3), and (4), the final estimate of the target lead is computed as the cumulative sum of the outputs from all three layers:

$$\hat{y}_t^{(3)} = \boldsymbol{\theta}_{\text{LS}}^\top \mathbf{x}_t + f_{\text{MLP}}(\mathbf{x}_t; \boldsymbol{\theta}_{\text{MLP}}) + f_{\text{RNN}}(\mathbf{x}_{t-s:t+s}; \boldsymbol{\theta}_{\text{RNN}}) \tag{5}$$

Each layer is modular, additive, and designed to capture what the previous step could not recover. This layered lead reconstruction approach enables interpretable decomposition of ECG leads by disentangling the contribution of linear, nonlinear and temporal structure.

3. Results

We evaluate the model on the PTB-XL dataset using 5-fold cross-validation, quantifying reconstruction fidelity of each layer in terms of signal-to-noise ratio (SNR). We report per-lead SNR at each layer. In each fold, all evaluations are conducted on the held-out set. Robustness is assessed by reporting the standard deviation across the folds.

Table 1 and Figure 2 summarize the results. Accordingly, the three-layer architecture yields consistent SNR

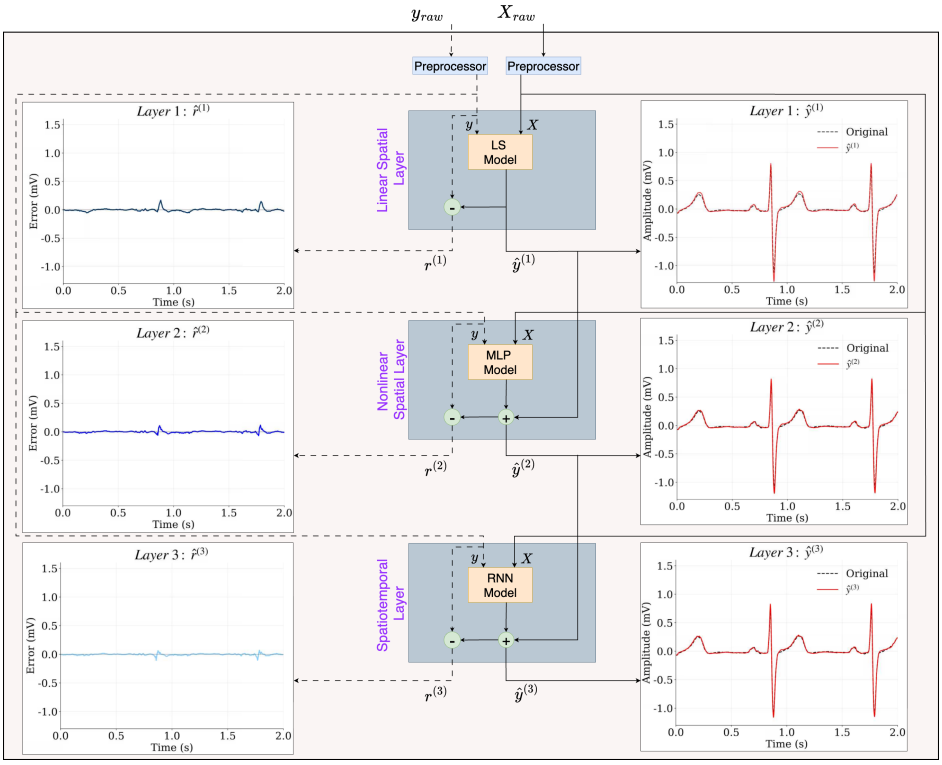


Figure 1. The proposed residual-based iterative model reconstructs a target ECG lead ($\mathbf{y}$) from other leads ($\mathbf{X}$) in three refinement layers. Layer 1 implements a LS model, resulting in $\hat{y}^{(1)}$. Layer 2 utilizes an MLP spatial block to generate a nonlinear approximation $\hat{y}^{(2)}$. Finally, Layer 3 employs an RNN spatio-temporal block to capture temporal and any remaining spatial correlations. These results indicate a progressive improvement in SNR across the layers. y and the dashed line paths are only available in the training phase; solid lines are available during both training and inference.

improvements across all ECG leads. The LS layer captures major inter-lead correlations, the MLP enhances performance by modeling nonlinear spatial dependencies, and the RNN adds temporal refinement. Figure 1 illustrates this progression on a sample target lead: Layer 1 captures the overall waveform but distorts high-frequency components, Layer 2 improves spatial alignment and improves the SNR to 20.99 dB, and Layer 3 further refines temporal details, yielding a final SNR of 22.01 dB. Overall, the model achieves a 22.01 dB in SNR from Layer 1 to Layer 3, highlighting how successive layers refine both morphology and temporal fidelity. Reconstruction accuracy varies across leads, reflecting differences in spatial redundancy. Among the precordial leads, V4 and V5 consistently achieve the highest SNRs, with V5 reaching 16.16 dB after Layer 3. Limb leads I and II are generally more difficult to recover, indicating greater lead-specific information that is harder to infer from others. Lead II obtained the lowest average SNR of 7.33 dB. The Lead II plane is known to have little overlap with the other leads. This confirms our hypothesis that lead-specific components, which cannot be recovered from other views, carry unique diagnostic value, resulting in lower reconstruction SNRs from other leads.

Table 1. SNR Mean$\pm$ STD over 5 folds for reconstructing different measured ECG leads (in dB). The results show progressive improvement from Layer 1 to Layer 3, with varying performance across leads.

ECG Lead	Layer 1 (LS)	Layer 2 (MLP)	Layer 3 (RNN)
I	6.54 ± 0.05	7.82 ± 0.04	8.22 ± 0.04
II	5.32 ± 0.04	6.97 ± 0.05	7.33 ± 0.07
V1	8.74 ± 0.13	9.30 ± 0.07	9.68 ± 0.08
V2	10.09 ± 0.11	11.00 ± 0.08	11.25 ± 0.08
V3	12.17 ± 0.10	12.73 ± 0.05	12.92 ± 0.09
V4	13.71 ± 0.08	14.89 ± 0.09	14.93 ± 0.13
V5	15.09 ± 0.10	16.11 ± 0.07	16.16 ± 0.06
V6	13.17 ± 0.24	14.20 ± 0.05	14.39 ± 0.06

4. Discussion

The results demonstrate that ECG leads can be partially reconstructed from one another using a structured, multi-layer model that progressively captures linear, nonlinear, and temporal dependencies. The consistent improvement

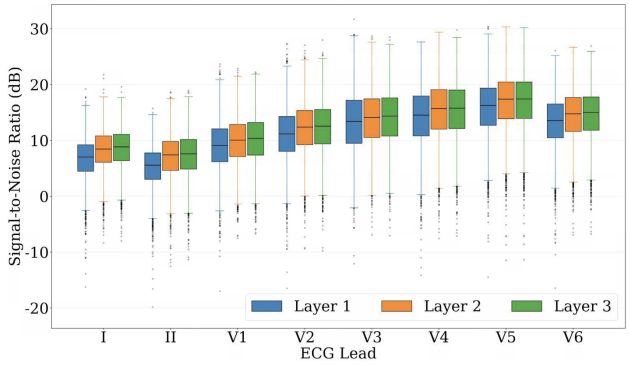


Figure 2. Signal-to-Noise Ratio (SNR) for lead reconstruction across the eight directly measured ECG leads (I, II, V1–V6) at three layers of the pipeline: Layer 1 (LS), Layer 2 (LS + MLP), and Layer 3 (LS + MLP + RNN). Boxplots summarize the SNR variability across the left-out validation set.

in SNR across layers confirms the significance of the modular design: each layer contributes unique explanatory power by modeling what the previous layer fails to capture. Even the final RNN layer, despite its simplicity, adds temporal refinement, which are useful for capturing temporal dynamics.

From a machine and deep learning standpoint, the key significance of the proposed model is the decomposition of ECG leads into shared and lead-specific components, enabling its usage as a signal decomposition block in more sophisticated architectures such as ECG foundation models used for ECG analysis and algorithmic diagnosis. The results confirm that certain leads (e.g., V4, V5) are more redundant and predictable from others, while the limb leads (I and II) carry more unique information. These differences align with anatomical expectations: mid-precordial leads lie closer to the heart and often capture "near-field" signal structure, whereas limb leads capture "far-field" characteristics. This approach to ECG decomposition provides a principled way to reason about lead importance and opens the door to lead selection/prioritization and feature extraction, in diagnostic pipelines.

These findings have multiple implications. First, identifying leads with low redundancy can guide lead selection for wearable or mobile ECG devices, where hardware constraints prevent full 12-lead acquisition. Second, shared and lead-specific components may respond differently to pathological changes, providing new waveform features for disease detection and interpretation. Finally, by quantifying inter-lead dependencies, the model can support more efficient ECG compression strategies. The proposed architecture may also be combined with more generic neural network architectures, or used as building blocks in ECG foundation models. This makes the model a versatile building block of larger architectures, extending its application beyond lead reconstruction.

5. Conclusion

We proposed a modular, three-layer neural network for ECG lead reconstruction and shared versus lead-specific information separation, capturing linear, nonlinear, and temporal inter-lead dependencies through residual learning. This approach decomposes leads into common and unique components, helping with lead reconstruction and interpretability. Experiments on the PTB-XL dataset showed consistent SNR improvements, confirming lead redundancy and uniqueness, with applications in ECG compression, lead selection for hardware-constrained devices, and interpretable representation learning. Future work may extend this model to diagnostic tasks and link lead-specific residuals to disease biomarkers.

References

[1] Clifford GD, Azuaje F, McSharry PE. Advanced methods and tools for ecg data analysis. Artech House 2006;.

[2] Goldberger AL. Clinical Electrocardiography: A Simplified Approach. 7th edition. Philadelphia, PA: Elsevier Health Sciences, 2006.

[3] Ribeiro AH, Ribeiro MH, Paixão GMM, Oliveira DM, Gomes PR, Canazart JA, Ferreira MP, Andersson CR, Macfarlane PW, Meira W, Schön TB, Ribeiro AL. Automatic diagnosis of the 12-lead ECG using a deep neural network. Nat Commun Apr. 2020;11:1760.

[4] Zhang D, Yang S, Yuan X, Zhang P. Interpretable deep learning for automatic diagnosis of 12-lead electrocardiogram. iScience Apr. 2021;24(4):102373.

[5] Ramirez E, Ruiperez-Campillo S, Casado-Arroyo R, Merino JL, Vogt JE, Castells F, Millet J. The art of selecting the ECG input in neural networks to classify heart diseases: a dual focus on maximizing information and reducing redundancy. npj Digital Medicine Jul. 2024;7(1):123. PMCID: PMC11491564, PMID: 39434723.

[6] Moghaddam DEP, Banta A, Post A, Razavi M, Aazhang B. Reconstructing 12-lead ECG from reduced lead sets using an encoder–decoder convolutional neural network. Biomed Signal Process Control Jun. 2025;104:107486.

[7] Dower GE, Machado HB, Osborne JA. On deriving the electrocardiogram from vectorcardiographic leads. Clin Cardiol Feb. 1980;3(2):87–95.

[8] Wagner P, Strodthoff N, Bousseljot RD, Kreiseler D, Lunze FI, Samek W, Schaeffter T. PTB-XL, a large publicly available electrocardiography dataset. Sci Data May 2020;7:Art. no. 154.

Address for correspondence:

Mohammadsina Hassannia
101 Woodruff Circle, Atlanta, GA, 30322
sina.hassannia@dbmi.emory.edu