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 stages. The first stage generalizes the Dower transform via a linear least squares (LS) approach to estimate each lead from the others. The second stage employs a multilayer perceptron (MLP) to capture nonlinear residuals not captured by the LS stage. Finally, a recurrent neural network (RNN) is used to model temporal structure in the MLP residuals. These stages 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 signal-to-noise ratio (SNR) across all stages. The best reconstruction was achieved for the precordial leads, with V5 showing the highest average SNR of over 16 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 arrhythmias [1–4]. However, the complex interplay of shared versus lead-specific components of multi-lead ECGs remains underexplored. Although 12-lead ECGs provide complementary views of cardiac activity, their components are both redundant—due to shared physiological origins—and unique, reflecting each lead's anatomical orientation [5, 6]. Hypothetically, disentangling the *shared* information across leads from *unique*, lead-specific patterns, could enable more efficient and interpretable representations for 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-stage cascade pipeline that decomposes shared and lead-specific information while preserving interpretability. The first stage employs a linear least-squares (LS) model, extending the Dower transform [7], to capture dominant linear correlations across leads. The second stage is a multilayer perceptron (MLP) that estimates nonlinear residuals, modeling complex anatomical and physiological relationships not accounted for by the LS stage. Finally, the third stage 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 stages 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, the model demonstrates consistent improvements in reconstruction signal-to-noise ratio (SNR) in each stage. The results confirm the significance of the layered approach, which can be used as building blocks in explainable ECG foundation models for ECG interpretation pipelines.

Compared to prior work on lead imputation or generative ECG modeling [6], our focus lies not only in accurate reconstruction but in interpretability and decomposition. We treat reconstruction as a lens through which the structure of ECG signals can be disentangled.

2. Methods

An overview of the proposed architecture is shown in Figure 1. Accordingly, the model reconstructs a target ECG lead from the remaining leads using a three-stage residual architecture. Each stage is trained to progressively refine the estimate of the previous stage by modeling distinct types of inter-lead dependencies: linear spatial, nonlinear spatial, and spatiotemporal. 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).

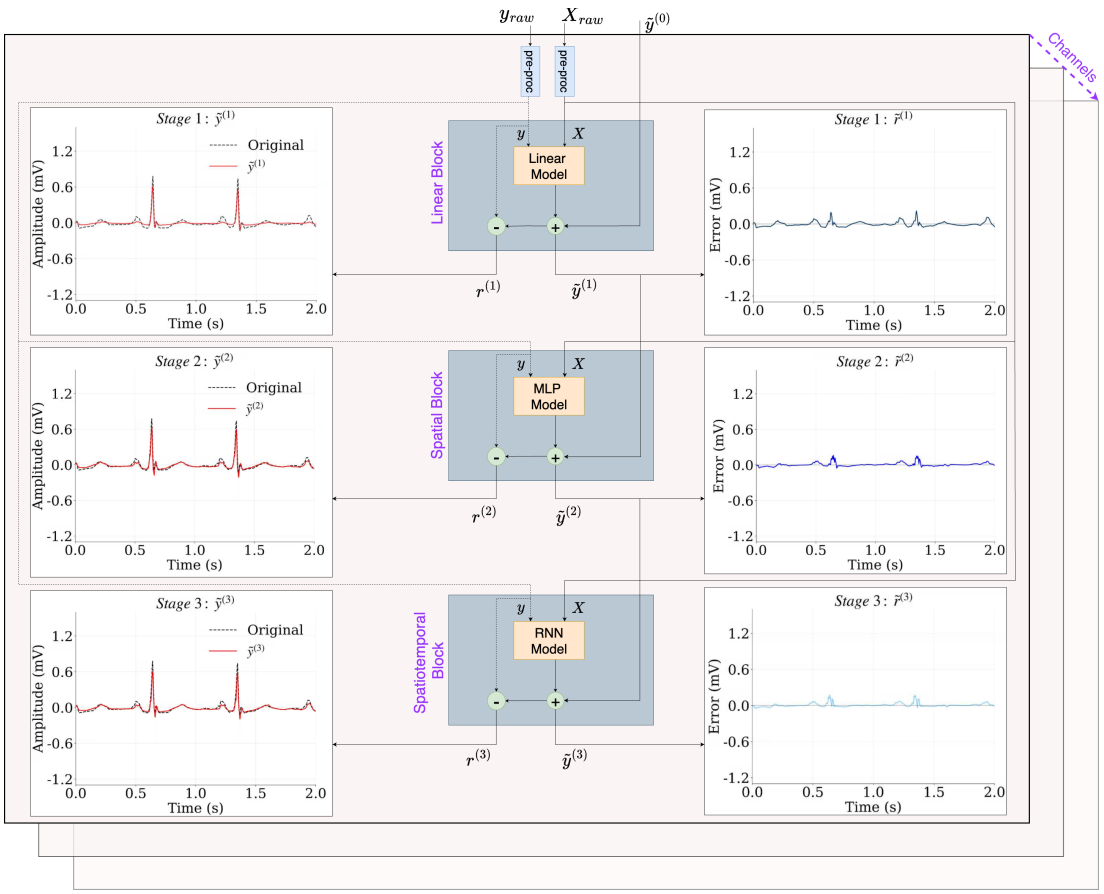

Figure 1. The proposed residual-based iterative model reconstructs a target ECG lead ($\mathbf{y}$) from other leads ($\mathbf{X}$) in three refinement stages. Stage 1 applies a linear least squares, yielding $\tilde{y}^{(1)}$ with a 4.68 dB average SNR (across all leads). Stage 2 uses an MLP spatial block to produce $\tilde{y}^{(2)}$ with 7.20 dB average SNR. Stage 3 employs an RNN spatio-temporal block, giving $\tilde{y}^{(3)}$ with 7.29 dB average SNR, demonstrating progressive enhancement across stages.

2.1. Linear Least Squares

The first stage uses a linear least squares (LS) regression to estimate the target lead as a weighted combination of the remaining leads. Formally, given the matrix $\mathbf{X} \in \mathbb{R}^{(L-1)\times T}$ representing $(L-1)$ ECG leads over T time steps, we solve for $\mathbf{w}_{\text{LS}} \in \mathbb{R}^{1\times(L-1)}$ that minimizes

$$\|\mathbf{y} - \mathbf{w}_{\text{LS}}\mathbf{X}\|_2^2, \tag{1}$$

where $\mathbf{y} \in \mathbb{R}^{1\times T}$ is the target ECG lead and $\|\cdot\|_2$ denotes the Euclidean (ℓ_2) norm. This stage generalizes the classical Dower Transform [7], which maps vectorcardiographic Frank leads to the 8- or 12-lead system. This stage captures dominant linear correlations, serving as a first-order linear approximation for the shared information between leads.

2.2. Nonlinear Least Squares with Multilayer Perceptron

While the LS projection in Stage 1 captures coarse linear inter-lead relationship, it cannot model nonlinear relationships. To address this, the second stage of the proposed architecture is a multilayer perceptron (MLP). For each time point t, the MLP receives $\mathbf{x}_t \in \mathbb{R}^{L-1}$ (i.e., the values of the other leads at the same time step) and predicts a nonlinear correction to the residual $r_t^{(1)} = y_t - \hat{y}_t^{(1)}$ from the first stage. Formally, the MLP learns a function $f_{\text{MLP}} : \mathbb{R}^{L-1} \to \mathbb{R}$ that minimizes the following objective:

$$\min_{\theta_{\text{MLP}}} \sum_{t=1}^{T} \left(y_t - \hat{y}_t^{(1)} - f_{\text{MLP}}(\mathbf{x}_t; \theta_{\text{MLP}}) \right)^2, \tag{2}$$

where θ_{MLP} are the MLP parameters. This formulation is structurally analogous to the linear least squares step in Stage 1, but with the added capacity to model nonlinear

spatial dependencies between the leads. The output of this stage is an updated prediction $\hat{y}^{(2)} = \hat{y}^{(1)} + f_{\text{MLP}}(\mathbf{X})$ that incorporates nonlinear refinements to the initial approximation.

2.3. Temporal Modeling using Recurrent Neural Network (RNN)

The third stage models temporal dependencies in the residual signal that remain after the spatial modeling in Stage 2. Let $r^{(2)} = y - \hat{y}^{(2)}$ denote the residual after the linear (LS) and nonlinear (MLP) steps. 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 s input signal $\mathbf{x}_{t-s:t+s} \in \mathbb{R}^{(L-1)\times(2s+1)}$ is extracted and passed to the RNN. The network uses this bidirectional approach to predict 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 stages were instantaneous and agnostic to the sampling frequency, the RNN stage is sampling rate dependent.

Formally, the RNN learns a function f_{RNN} with parameters θ_{RNN} that minimizes the temporal residual loss:

$$\min_{\theta_{\text{RNN}}} \sum_{t=1}^{T} \left(y_t - \hat{y}_t^{(2)} - f_{\text{RNN}}(\mathbf{x}_{t-s:t+s}; \theta_{\text{RNN}}) \right)^2. \quad (3)$$

After applying the RNN correction, the final model output is given by:

$$\hat{y}_t^{(3)} = \hat{y}_t^{(2)} + f_{\text{RNN}}(\mathbf{x}_{t-s:t+s}).$$

2.4. Final Reconstruction

The final estimate of the target lead is computed as the cumulative sum of the outputs from all three stages:

$$\hat{y}_t^{(3)} = \hat{y}_t^{(1)} + f_{\text{MLP}}(\mathbf{x}_t) + f_{\text{RNN}}(\mathbf{x}_{t-s:t+s}), \quad (4)$$

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

3. Results

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

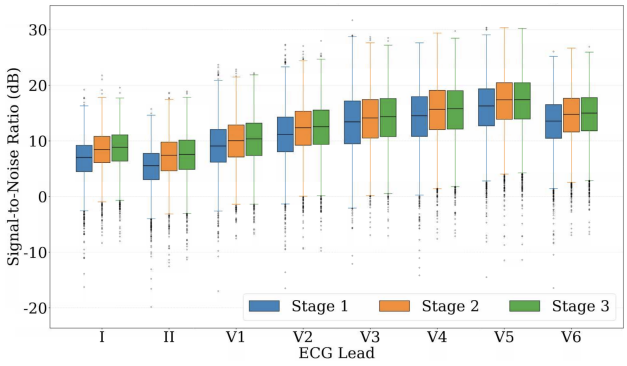

Figure 2. Signal-to-Noise Ratio (SNR) for lead reconstruction across the eight directly measured ECG leads (I, II, V1–V6) at three stages of the pipeline: Stage 1 (LS), Stage 2 (LS + MLP), and Stage 3 (LS + MLP + RNN). Boxplots summarize the SNR variability across the left-out validation set.

Table 1 and Figure 2 summarize the results. Accordingly, the three-stage architecture yields consistent SNR improvements across all ECG leads. The LS stage 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 target lead: Stage 1 captures the overall waveform but distorts high-frequency components, Stage 2 improves spatial alignment and improves the SNR to 7.20 dB, and Stage 3 further refines temporal details, yielding a final SNR of 7.29 dB. Overall, the model achieves a 7.29 dB in SNR from Stage 1 to Stage 3, highlighting how successive stages refine both morphology and temporal fidelity.

Reconstruction accuracy varies across leads, reflecting differences in spatial redundancy. Limb leads I and II are generally more difficult to recover, indicating greater lead-specific information that is harder to infer from others. Among the precordial leads, V4 and V5 consistently achieve the highest SNRs, with V5 reaching 16.16 dB after Stage 3. This confirms our hypothesis that lead-specific components—irrecoverable from other views—correspond to lower reconstruction SNRs and carry unique diagnostic value.

4. Discussion

The results demonstrate that ECG leads can be effectively reconstructed from one another using a structured, multi-stage model that progressively captures linear, non-

linear, and temporal dependencies. The consistent improvement in SNR across stages confirms the significance of the modular design: each stage contributes unique explanatory power by modeling what the previous stage fails to capture. In particular, the final RNN stage—though relatively simple—adds temporal refinement that proves useful for capturing temporal dynamics.

From a machine and deep learning standpoint, the key significance of the proposed model is the explicit decomposition of ECG leads into shared (redundant) and specific (non-redundant) components, which enables its usage as a signal decomposition block in more sophisticated architectures such as ECG foundation models used for algorithmic diagnosis. The per-lead reconstruction analysis confirms that certain leads (e.g., V3, V4, V5) are more redundant and predictable from others, while others (e.g., the limb leads I, II) carry more unique information. These differences align with anatomical expectations: mid-precordial leads lie closer to the heart and often share more "near-field" signal structure, whereas limb leads capture orthogonal "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 specific components may respond differently to pathological changes—providing new waveform features for disease detection and interpretation. Finally, by quantifying inter-lead predictability, the model can support more efficient ECG compression strategies where only the most informative leads are stored or transmitted. 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 valuable building block for pretraining or diagnostics, extending beyond lead reconstruction.

5. Conclusion

We proposed a modular, three-stage 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.

Table 1. Mean SNR ($\pm$ std) over 5 folds for different measured ECG leads (in dB). Note the progressive improvement from Stage 1 to 3, and varied performance across leads.

ECG Lead	Stage 1 (LS)	Stage 2 (MLP)	Stage 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

Address for correspondence:

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

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.