# MetaFormer: Efficient Metadata-Guided Transformer for Patient-Aware 12-Lead Electrocardiogram Reconstruction

Yanchong Xie[1], Jiepeng Chen[1], Zefeng Li[1], Kai Zheng[2], Haoyang Huang[1], Minglong Zheng[3], Mingyue Cui[1(✉)], and Kai Huang[1]

[1] School of Computer Science and Engineering, Sun Yat-sen University, Guangzhou, China
cuimy@mail2.sysu.edu.cn
[2] School of Biomedical Engineering, Sun Yat-sen University, Guangzhou, China
[3] Sun Yat-sen Memorial Hospital, Sun Yat-sen University, Guangzhou, China

**Abstract.** Reconstructing 12-lead electrocardiogram (ECG) signals from single-lead wearable recordings bridges the gap between long-term monitoring for critical paroxysmal events and the clinical gold standard required for accurate diagnosis. However, existing methods usually treat this ill-posed inverse problem as a generic regression task, failing to resolve physiological ambiguities and consequently reconstructing toward population averages. To address these problems, we propose MetaFormer, an efficient reconstruction network that utilizes demographic characteristics to enforce patient-aware physical consistency. We introduce a meta tokenizer to parse heterogeneous clinical metadata into a unified semantic manifold, which establishes patient-specific physiological context to guide the reconstruction. To reduce the cross-lead spatial ambiguity, a metadata-guided Transformer aligns latent representations of the reconstructed signal with individual physiological baselines. Experimental results on two large-scale benchmarks demonstrate that MetaFormer outperforms other state-of-the-art methods while reducing computational cost by ∼73%, and closely approaches the performance of original 12-lead ECG signals on downstream arrhythmia classification. The code is available at: https://github.com/xych3/MetaFormer.

**Keywords:** Electrocardiogram · Signal reconstruction · Arrhythmia diagnosis · Wearable devices.

## 1 Introduction

Cardiovascular diseases (CVDs) are responsible for approximately one-third of global deaths annually [18], posing a substantial public health burden. Many life-threatening events, such as paroxysmal atrial fibrillation, often occur in out-of-hospital settings and are easily missed during a brief clinical 12-lead electrocardiogram (ECG) snapshot [3,21]. Benefiting from the increasing availability of wearable devices, continuous monitoring offers a practical solution to capture

these transient cardiac abnormalities [20,15]. However, wearable devices such as smartwatches typically adopt single-lead configurations [29] to maintain portability, which restricts the accurate diagnosis of localized pathologies [2,1] and fails to satisfy the clinical 12-lead gold standard [13].

Reconstructing the complete 12-lead ECG from a reduced lead set offers a promising solution to bridge this gap. Previous researches focuse on reconstructing from a subset of 12 leads, usually including three or six limb leads [5,14,25]. However, their reliance on multi-electrode setups restricts wearable applications. Recent studies attempt to synthesize the full 12-lead signals from a single reference lead. For instance, EKGAN [12] reconstructs 12-lead from Lead I with a conditional generative adversarial network (CGAN) [17], which aligns the latent representations of a 1D U-Net generator and an autoencoder with identical inputs to model the characteristics of the input lead. Following this, the parameter-efficient architecture mEcgNet [15] employs frequency-based separation to reconstruct the target waveforms, while the hybrid convolutional network ECGrecover [16] utilizes 1D temporal convolutions with 2D spatial filtering to reconstruct missing leads. Despite these advances, due to the inherent influence of individual anatomical and physiological characteristics on cardiac electrical activity [29], generating a complex spatial representation from a single observation remains challenging. In general, existing methods treat this cross-lead transformation as a generic regression, failing to resolve the physiological ambiguity where diverse spatial cardiac vectors project onto similar single-lead waveforms. Without patient-specific constraints, these methods inevitably regress toward population means and neglect the fine-grained morphologies required for precise clinical interpretation, leading to structural hallucinations and mode collapse.

In this paper, we propose MetaFormer, an efficient reconstruction network that explicitly incorporates patient-specific physiological attributes as semantic priors to constrain the solution space. Specifically, our method consists of a meta tokenizer to embed heterogeneous metadata into a patient-aware physiological context, and a metadata-guided Transformer that aligns latent representations with these individual baselines to resolve cross-lead spatial ambiguity. We compare our MetaFormer with state-of-the-art reconstruction methods [12,15,16] on two large-scale ECG benchmarks PTB-XL [24] and CSN [28]. The results show that MetaFormer achieves superior reconstruction accuracy and efficiency. Moreover, the reconstructed signals closely approach the performance of the original 12-lead ECGs. The main contributions of this paper are as follows: (1) We propose MetaFormer, a novel metadata-guided architecture that reconstructs the full 12-lead ECG signal from a single reference lead in a patient-aware generation process. (2) A meta tokenizer is introduced by constructing a structured semantic manifold from diverse demographic attributes, effectively leveraging a learnable null-space formulation for robust feature parsing. (3) To solve the physiological ambiguity in the ill-posed inverse reconstruction, we introduce a metadata-guided Transformer that aligns the latent representations with individualized constraints to reconstruct precise diagnostic waveforms.

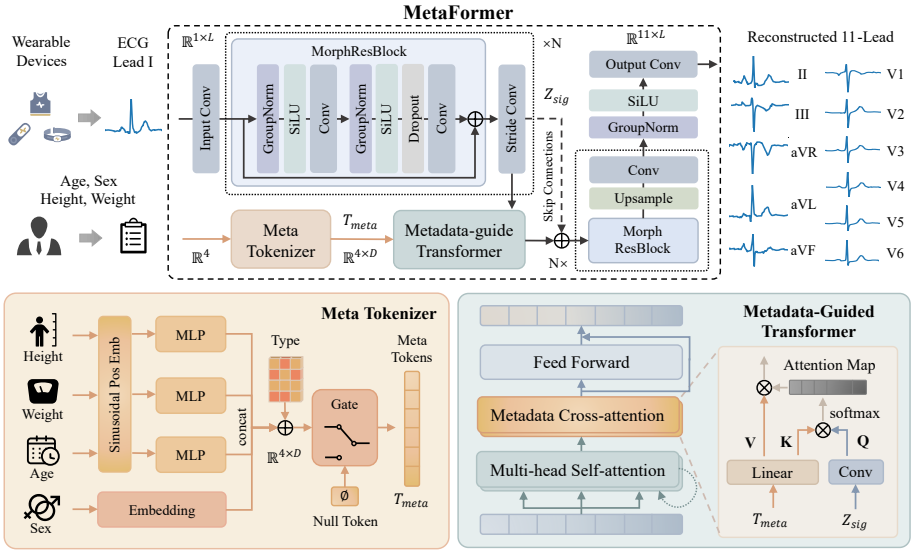


Fig. 1: The framework of MetaFormer. The architecture includes a meta tokenizer that encodes metadata into semantic tokens, and a metadata-guided Transformer backbone that integrates these patient-specific priors via cross-attention to reconstruct the complete 12-lead ECG from a single reference lead.

## 2 Methods

As shown in Figure 1, we propose MetaFormer, an efficient end-to-end framework that integrates patient-specific physiological priors to reconstruct from single lead. Specifically, the reference lead $\mathbf{X}_{ref}$ is first processed by a series of MorphResBlocks, which extract hierarchical morphological features and compress the signal into a latent representation $\mathbf{Z}_{sig}$. Concurrently, heterogeneous attributes are processed by the meta tokenizer and mapped into a structured semantic space to a sequence of high-dimensional meta tokens $\mathbf{T}_{meta}$. Then, both latent features and meta tokens are fed into the metadata-guided Transformer, where a cross-attention mechanism dynamically recalibrates the signal representations using the demographic priors to resolve cross-lead spatial ambiguity. Finally, these physiologically modulated features are upsampled and decoded through a symmetrical residual architecture to restore the full 12-lead ECG.

**Meta Tokenizer** Cardiac electrical activity is constrained by the anatomical geometry and physiological characteristics of the subject, including the position of the heart relative to the electrodes and the conduction velocity. To leverage these critical heterogeneity priors, we propose the meta tokenizer that explicitly models patient attributes as a structured semantic context. This module maps heterogeneous clinical metadata, denoted as $\mathcal{M} = \{m_1, m_2, \ldots, m_N\}$ (e.g., age,

sex, height, weight), into a sequence of high-dimensional tokens $\mathbf{T}_{meta} \in \mathbb{R}^{N \times d}$, where $N$ is the number of attributes and $d$ is the embedding dimension. To handle different data modalities, we apply distinct mapping functions. For a continuous attribute $m_i \in \mathbb{R}$, we employ a sinusoidal encoding strategy [23] to preserve numerical continuity, followed by a multi-layer perceptron (MLP) projection. The embedding $\mathbf{e}_i$ is obtained as:

$$\mathbf{e}_i = \text{MLP}\left(\left[\sin(m_i/\Omega^{2j/d}), \cos(m_i/\Omega^{2j/d})\right]_{j=1}^{d/2}\right), \tag{1}$$

where $\Omega$ is a frequency scaling factor and $j$ denotes the feature dimension index. For a discrete attribute $m_i$, we map the categorical value via a learnable lookup table to obtain $\mathbf{e}_i$. To differentiate attribute modalities within the sequence, we add a learnable type embedding $\mathbf{p}_i \in \mathbb{R}^d$ to each encoded feature. In clinical scenarios, metadata is frequently incomplete. We formalize a robust parsing mechanism by introducing learnable null tokens $\mathbf{e}_{\emptyset,i}$ for each attribute. Let $c_i \in \{0, 1\}$ be a binary indicator representing the availability of attribute $m_i$. The specific token representation is formulated as:

$$\mathbf{t}_i = c_i(\mathbf{e}_i + \mathbf{p}_i) + (1 - c_i)\mathbf{e}_{\emptyset,i}. \tag{2}$$

The final contextual representation is constructed as $\mathbf{T}_{meta} = [\mathbf{t}_1, \mathbf{t}_2, \ldots, \mathbf{t}_N]^\top$. During training, we randomly set $c_i = 0$ with a probability $p_{drop}$ to simulate missing data, enabling the network to construct a robust semantic manifold and learn an adaptive reliance on available physiological priors.

**Metadata-Guided Transformer** Reconstructing the complete 12-lead ECG from a single lead is an ill-posed inverse problem, as multiple high-dimensional cardiac vector trajectories project onto similar single-lead waveforms. To overcome this, we design the metadata-guided Transformer, which utilizes patient-specific tokens in the signal reconstruction process to align the latent representations with physical constraints. The single-lead input $\mathbf{X}_{ref} \in \mathbb{R}^{1 \times L}$ is first encoded into a latent feature map $\mathbf{Z}_{sig} \in \mathbb{R}^{L' \times d}$ via a series of MorphResBlocks. Each block consists of 1D convolution layers, group normalization [27], SiLU activation [7], and a dropout layer, with a residual connection to preserve morphological details. Downsampling is performed via strided convolutions to capture semantic features. To align the latent features with physiological constraints, we introduce a metadata-guided attention mechanism at the bottleneck. Specifically, a multi-head self-attention mechanism models the global temporal dynamics and long-range dependencies across the cardiac cycle. Subsequently, a metadata-guided cross-modal attention mechanism explicitly calibrates these temporal features using patient-specific priors to align the reconstructed spatial topology with personalized physiological constraints. Specifically, the ECG features $\mathbf{Z}_{sig}$ serve as the Query ($\mathbf{Q}_{\mathbf{sig}}$), and the meta tokens $\mathbf{T}_{meta}$ serve as both

Key ($\mathbf{K_{meta}}$) and Value ($\mathbf{V_{meta}}$):

$$\begin{aligned} \mathbf{Q_{sig}} = \mathbf{Z}_{sig}\mathbf{W}_Q, \quad \mathbf{K_{meta}} = \mathbf{T}_{meta}\mathbf{W}_K, \quad \mathbf{V_{meta}} = \mathbf{T}_{meta}\mathbf{W}_V, \\ \text{Attention}(\mathbf{Q_{sig}}, \mathbf{K_{meta}}, \mathbf{V_{meta}}) = \text{Softmax}\left(\frac{\mathbf{Q_{sig}}\mathbf{K_{meta}}^\top}{\sqrt{d_k}}\right)\mathbf{V_{meta}}, \end{aligned} \tag{3}$$

where $\mathbf{W}_Q, \mathbf{W}_K, \mathbf{W}_V$ are learnable matrices, and $\sqrt{d_k}$ is the scaling factor based on the key dimension. The modulated features are then processed by a position-wise feed-forward network (FFN) composed of two linear transformations with a GeLU activation [9]. Linear interpolation restores temporal resolution, and skip connections integrate high-resolution features to ensure the preservation of fine-grained morphological details. Finally, the modulated representations are decoded to reconstruct the remaining 11 target leads $\hat{\mathbf{X}}_{target} \in \mathbb{R}^{11\times L}$.

**Learning** We optimize MetaFormer using a joint objective function that combines mean squared error (MSE) loss and Pearson correlation coefficient (PCC) losses to ensure both signal fidelity and morphological precision. The MSE loss $\mathcal{L}_{mse}$ minimizes point-wise voltage errors for accurate amplitude recovery:

$$\mathcal{L}_{mse} = \frac{1}{B \cdot C \cdot L}\sum_{i=1}^{B} \|\hat{\mathbf{X}}^{(i)} - \mathbf{X}_{gt}^{(i)}\|_2^2. \tag{4}$$

where $B$ is the batch size, $C = 11$ denotes the number of output channels, $L$ is the signal length, $\hat{\mathbf{X}}^{(i)}$ is the reconstructed signal, and $\mathbf{X}_{gt}^{(i)}$ is the ground truth. The Pearson loss $\mathcal{L}_{pcc}$ captures the global structure of the ECG:

$$\mathcal{L}_{pcc} = 1 - \frac{1}{B \cdot C}\sum_{i=1}^{B}\sum_{c=1}^{C} \frac{\text{Cov}(\hat{\mathbf{x}}_c^{(i)}, \mathbf{x}_{c,gt}^{(i)})}{\sigma_{\hat{\mathbf{x}}_c^{(i)}} \cdot \sigma_{\mathbf{x}_{c,gt}^{(i)}}}, \tag{5}$$

where $\text{Cov}(\cdot)$ denotes the covariance, and $\sigma$ represents the standard deviation along the temporal dimension. The total loss is formulated as $\mathcal{L}_{total} = \mathcal{L}_{mse} + \lambda\mathcal{L}_{pcc}$, where $\lambda$ balances the relative importance between absolute fidelity and morphological consistency, and is set to 0.1 according to grid search.

## 3 Experiments

**Datasets and Implementation** We evaluate our proposed method on two large-scale public benchmarks: PTB-XL [24] and Chapman–Shaoxing–Ningbo (CSN) [28]. The PTB-XL dataset comprises 21,799 recordings from 18,869 patients across 71 arrhythmia statements that conform to the SCP-ECG standard [22]. The CSN dataset contains 45,152 recordings with SNOMED-CT [4] annotations by up to two cardiologists. Following the preprocessing strategy in [12,15,16], we normalize all signals to the range of [-1, 1], apply a band-pass filter with a cutoff frequency of [0.05, 150] Hz, and resample d to a fixed length of

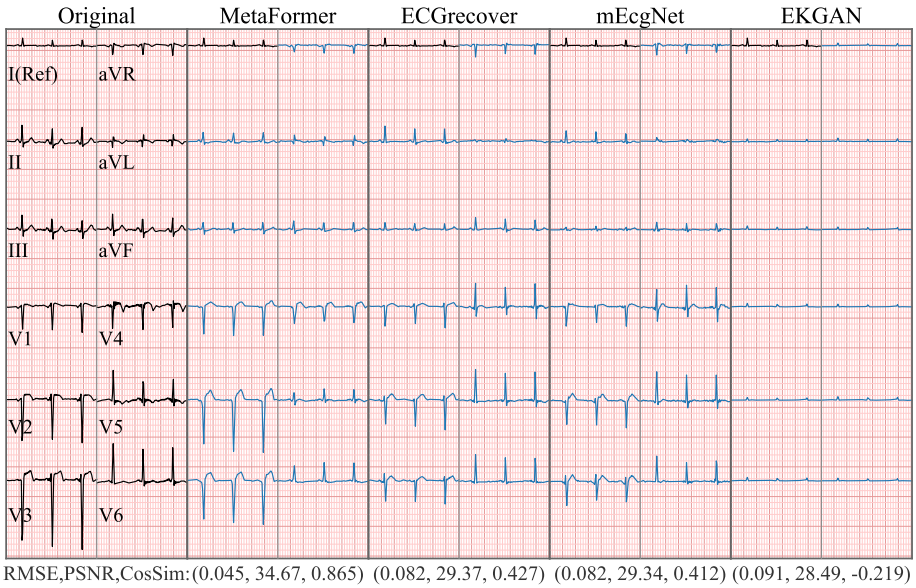


Fig. 2: Visualization of reconstruction performance. The black traces represent the ground-truth ECG, while the blue traces denote the 12-lead ECG reconstructed by different methods. Ref, Reference.

512. The MetaFormer operates across four stages with the channels progressively expanded as $\{32, 64, 96, 128\}$. For the meta tokenizer, attributes are projected into a 256-dimensional space. The Metadata-Guided Transformer utilizes 8-head self-attention ($d_h = 64$) and 8-head cross-attention ($d_h = 32$), with a dropout rate of 0.3 to prevent overfitting. We implement the model in PyTorch and train on a single NVIDIA RTX A100 GPU using the AdamW optimizer with a weight decay of 0.1. The learning rate follows a cosine annealing schedule, warming up from 0 to $1 \times 10^{-4}$ over the first 5 epochs before decaying to $1 \times 10^{-6}$.

**Comparison with Baseline Methods** We compare MetaFormer with three state-of-the-art baselines: EKGAN [12], mEcgNet [15], and ECGrecover [16]. Evaluation metrics include root mean square error (RMSE) [26], peak signal-to-noise ratio (PSNR) [26], percent root mean square difference (PRD) [11], and cosine similarity (CosSim) [19]. As shown in Table 1, MetaFormer consistently outperforms all baselines on PTB-XL and CSN, as expected. ECGrecover achieves a suboptimal reconstruction performance by employing a hybrid 1D/2D U-Net architecture to model both inter-lead and intra-lead relationships. However, it treats the reconstruction as a generic mapping task and leads to physiological ambiguity among subjects. In contrast, MetaFormer overcomes this limitation with the proposed metadata-guided Transformer. By explicitly utilizing metadata to constrain the solution space, it achieves the lowest RMSE of 0.063 and

Table 1: Performance and complexity comparison of different methods.

| Method | Input | PTB-XL | | | CSN | | | Params (M) | FLOPs (G) |
|---|---|---|---|---|---|---|---|---|---|
| | | RMSE↓ | PSNR↑ | CosSim↑ | RMSE↓ | PSNR↑ | CosSim↑ | | |
| EKGAN | Lead I | 0.111 | 31.15 | 0.231 | 0.109 | 33.41 | 0.303 | 12.52 | <u>1.29</u> |
| mEcgNet | Lead I | 0.072 | 35.20 | 0.760 | 0.078 | 36.68 | 0.713 | <u>2.09</u> | 1.52 |
| ECGrecover | Lead I | <u>0.067</u> | <u>35.77</u> | <u>0.795</u> | <u>0.072</u> | <u>37.28</u> | <u>0.758</u> | 5.93 | 4.59 |
| **Ours** | Lead I+Meta | **0.063** | **36.30** | **0.817** | **0.068** | **37.78** | **0.783** | **1.61** | **0.34** |

Table 2: Ablation study of MetaFormer on PTB-XL dataset.

(a) Prior Injection

| Method | RMSE↓ | CosSim↑ | AUC↑ | F1↑ |
|---|---|---|---|---|
| w/o Metadata | 0.075 | 0.770 | 0.94 | 0.69 |
| Concat | 0.069 | 0.792 | 0.95 | 0.71 |
| AdaIN | 0.066 | 0.805 | 0.95 | 0.73 |
| **GuidedAttn** | **0.063** | **0.817** | **0.96** | **0.75** |

(b) Metadata Embedding

| Configuration | RMSE↓ | CosSim↑ | AUC↑ | F1↑ |
|---|---|---|---|---|
| Linear Projection | 0.067 | 0.802 | 0.95 | 0.72 |
| **Meta Tokenizer** | **0.063** | **0.817** | **0.96** | **0.75** |
| w/o Null Token | 0.066 | 0.806 | 0.95 | 0.73 |
| w/o Type Embed | 0.065 | 0.811 | 0.96 | 0.74 |

0.067 and the highest CosSim of 0.817 and 0.783 on PTB-XL and CSN, respectively. Furthermore, by effectively modeling the patient attributes as semantic priors to recalibrate the latent features, MetaFormer achieves superior performance with only 1.61M parameters and 0.34 GFLOPs, reducing the parameter count by 72.8% and FLOPs by 92.6% compared to the competitive ECGrecover.

**Qualitative Visualization** As shown in Figure 2, we also visualize the reconstruction performance of baselines. We can observer that MetaFormer accurately recovers the fine-grained morphological details, including the sharp QRS complexes and pathological ST-T wave morphologies in the precordial leads. In contrast, the GAN-based method (EKGAN) exhibits mode collapse with phase misalignment, while ECGrecover and mEcgNet neglect distinct diagnostic features due to their tendency to regress towards a population mean representation.

**Ablation Study** Table 2 presents the ablation study of the proposed MetaFormer on the PTB-XL dataset. From the table, we can observe that models with physiological priors consistently outperform the baseline, and our metadata-guided attention superior to direct concatenation and AdaIN [10], as expected. This mainly because the guided attention mechanism utilizes metadata to recalibrate latent features, which preserve patient-specific structural details compared to directly aggregation and distribution modulation. We further verify the attribute embedding strategy of our proposed meta tokenizer. It can be noted that replacing the proposed tokenizer with a simple linear projection or removing the learnable null tokens and type embeddings leads to performance degradation. A

Table 3: Arrhythmia classification performance of different baselines. SR, Sinus Rhythm; AFIB, Atrial Fibrillation; AF, Atrial Flutter; CRBBB, Complete Right Bundle Branch Block; CLBBB, Complete Left Bundle Branch Block; SVT, Supraventricular Tachycardia; PVC, Premature Ventricular Contractions.

| Dataset | Method | SR | | AFIB | | CRBBB | | CLBBB | | PVC | | Average | |
|---|---|---|---|---|---|---|---|---|---|---|---|---|---|
| | | AUC | F1 | AUC | F1 | AUC | F1 | AUC | F1 | AUC | F1 | AUC | F1 |
| PTB-XL | Original | 0.86 | 0.91 | 0.95 | 0.63 | 0.99 | 0.83 | 0.99 | 0.82 | 0.97 | 0.64 | 0.95 | 0.77 |
| | EKGAN | 0.67 | 0.70 | 0.80 | 0.00 | 0.81 | 0.00 | 0.89 | 0.00 | 0.79 | 0.06 | 0.79 | 0.15 |
| | mEcgNet | 0.86 | 0.90 | 0.94 | 0.55 | 0.96 | 0.37 | **0.99** | 0.72 | 0.95 | 0.40 | 0.94 | 0.59 |
| | ECGrecover | **0.87** | **0.91** | **0.95** | 0.60 | **0.99** | 0.73 | **0.99** | **0.83** | 0.96 | 0.62 | 0.95 | 0.74 |
| | **Ours** | **0.87** | **0.91** | **0.95** | **0.62** | **0.99** | **0.74** | **0.99** | 0.81 | **0.97** | **0.63** | **0.96** | **0.75** |

| Dataset | Method | SR | | AF | | SVT | | CRBBB | | PVC | | Average | |
|---|---|---|---|---|---|---|---|---|---|---|---|---|---|
| | | AUC | F1 | AUC | F1 | AUC | F1 | AUC | F1 | AUC | F1 | AUC | F1 |
| CSN | Original | 0.98 | 0.88 | 0.95 | 0.78 | 0.99 | 0.71 | 0.99 | 0.66 | 0.98 | 0.64 | 0.98 | 0.73 |
| | EKGAN | 0.96 | 0.81 | 0.91 | 0.52 | 0.98 | 0.47 | 0.65 | 0.00 | 0.88 | 0.00 | 0.88 | 0.36 |
| | mEcgNet | 0.97 | 0.85 | 0.92 | 0.57 | **0.99** | 0.60 | 0.95 | 0.18 | 0.92 | 0.31 | 0.95 | 0.50 |
| | ECGrecover | **0.98** | **0.88** | 0.95 | 0.76 | **0.99** | 0.63 | 0.98 | 0.68 | 0.95 | 0.54 | 0.97 | 0.70 |
| | **Ours** | **0.98** | **0.88** | **0.96** | **0.77** | **0.99** | **0.69** | **0.99** | **0.70** | **0.96** | **0.56** | **0.98** | **0.72** |

potential reason is that the simple linear projection fails to capture the high-dimensional ordinal relationships of physiological variables, while the absence of specific tokens degrades the model's ability to differentiate between attribute modalities, resulting in suboptimal latent parsing. These results validate the effectiveness of leveraging patient-specific priors in ECG reconstruction.

**Arrhythmia Diagnosis** To further verify the clinical utility of the reconstructed 12-lead ECG signals in arrhythmia diagnosis, we independently train the baseline classifier proposed by [8] using original signals from PTB-XL and CSN datasets. Diagnosis performance is quantified by the area under the receiver operating characteristic curve (AUC) and F1-score (F1) on 12-lead signals reconstructed from Lead I by different methods. As shown in Table 3, MetaFormer successfully reconstructs the pathological features and preserves diagnostic integrity, achieving the averaging AUCs of 0.96 and 0.98 with corresponding F1-scores of 0.75 and 0.72, which exhibits a minimum performance discrepancy of less than 3% compared to the ground truth. While mEcgNet and ECGrecover show competitive results, they tend to reconstruct generic waveform templates and struggle in modeling complex pathological variations, which limits the accurate reconstruction of the abnormal morphology and rapid, irregular ventricular rates in severe arrhythmias like supraventricular tachycardia (SVT) [6].

## 4 Conclusion

This study introduces MetaFormer, an efficient framework designed to reconstruct standard 12-lead ECGs from single-lead wearable recordings. By incorporating metadata as physiological priors, MetaFormer achieves patient-aware reconstruction across both normal and pathological rhythms, preserving the subtle diagnostic details for reliable arrhythmia classification while maintaining computational efficiency. While MetaFormer demonstrates superior performance in reconstruction, constrained by the inherent information loss in single-lead observation, recovering the morphological details of complex cardiac electrical rhythms remains differences. Nevertheless, our method expands the clinical utility of wearable devices, providing a new perspective for long-term cardiac monitoring.

**Disclosure of Interests.** The authors have no competing interests to declare that are relevant to the content of this article.

## References

1. Al-Zaiti, S.S., Martin-Gill, C., Zègre-Hemsey, J.K., Bouzid, Z., Faramand, Z., Alrawashdeh, M.O., Gregg, R.E., Helman, S., Riek, N.T., Kraevsky-Phillips, K., et al.: Machine learning for ecg diagnosis and risk stratification of occlusion myocardial infarction. Nature Medicine **29**(7), 1804–1813 (2023)
2. Attia, Z.I., Harmon, D.M., Dugan, J., Manka, L., Lopez-Jimenez, F., Lerman, A., Siontis, K.C., Noseworthy, P.A., Yao, X., Klavetter, E.W., et al.: Prospective evaluation of smartwatch-enabled detection of left ventricular dysfunction. Nature medicine **28**(12), 2497–2503 (2022)
3. Attia, Z.I., Noseworthy, P.A., Lopez-Jimenez, F., Asirvatham, S.J., Deshmukh, A.J., Gersh, B.J., Carter, R.E., Yao, X., Rabinstein, A.A., Erickson, B.J., et al.: An artificial intelligence-enabled ecg algorithm for the identification of patients with atrial fibrillation during sinus rhythm: a retrospective analysis of outcome prediction. The Lancet **394**(10201), 861–867 (2019)
4. Benson, T., Grieve, G.: Snomed ct. In: Principles of Health Interoperability: FHIR, HL7 and SNOMED CT, pp. 293–324. Springer (2020)
5. Cho, Y., Kwon, J.m., Kim, K.H., Medina-Inojosa, J.R., Jeon, K.H., Cho, S., Lee, S.Y., Park, J., Oh, B.H.: Artificial intelligence algorithm for detecting myocardial infarction using six-lead electrocardiography. Scientific reports **10**(1), 20495 (2020)
6. Delacrétaz, E.: Supraventricular tachycardia. New England Journal of Medicine **354**(10), 1039–1051 (2006)
7. Elfwing, S., Uchibe, E., Doya, K.: Sigmoid-weighted linear units for neural network function approximation in reinforcement learning. Neural networks **107**, 3–11 (2018)
8. Hannun, A.Y., Rajpurkar, P., Haghpanahi, M., Tison, G.H., Bourn, C., Turakhia, M.P., Ng, A.Y.: Cardiologist-level arrhythmia detection and classification in ambulatory electrocardiograms using a deep neural network. Nature medicine **25**(1), 65–69 (2019)
9. Hendrycks, D., Gimpel, K.: Gaussian error linear units (gelus). arXiv preprint arXiv:1606.08415 (2016)

10. Huang, X., Belongie, S.: Arbitrary style transfer in real-time with adaptive instance normalization. In: Proceedings of the IEEE international conference on computer vision. pp. 1501–1510 (2017)
11. Jalaleddine, S.M., Hutchens, C.G., Strattan, R.D., Coberly, W.A.: Ecg data compression techniques-a unified approach. IEEE transactions on Biomedical Engineering **37**(4), 329–343 (1990)
12. Joo, J., Joo, G., Kim, Y., Jin, M.N., Park, J., Im, H.: Twelve-lead ecg reconstruction from single-lead signals using generative adversarial networks. In: International Conference on Medical Image Computing and Computer-Assisted Intervention. pp. 184–194. Springer (2023)
13. Kligfield, P., Gettes, L.S., Bailey, J.J., Childers, R., Deal, B.J., Hancock, E.W., Van Herpen, G., Kors, J.A., Macfarlane, P., Mirvis, D.M., et al.: Recommendations for the standardization and interpretation of the electrocardiogram: part i: the electrocardiogram and its technology: a scientific statement from the american heart association electrocardiography and arrhythmias committee, council on clinical cardiology; the american college of cardiology foundation; and the heart rhythm society endorsed by the international society for computerized electrocardiology. Circulation **115**(10), 1306–1324 (2007)
14. Lee, J., Kim, M., Kim, J.: Reconstruction of precordial lead electrocardiogram from limb leads using the state-space model. IEEE Journal of Biomedical and Health Informatics **20**(3), 818–828 (2015)
15. Lee, J., Yoo, Y., Kim, J., Lim, D., Yang, G., Yoo, C.: Parameter-efficient 12-lead ecg reconstruction from a single lead. In: International Conference on Medical Image Computing and Computer-Assisted Intervention. pp. 431–441. Springer (2025)
16. Lence, A., Granese, F., Fall, A., Hanczar, B., Salem, J.E., Zucker, J.D., Prifti, E.: Ecgrecover: a deep learning approach for electrocardiogram signal completion. In: Proceedings of the 31st ACM SIGKDD Conference on Knowledge Discovery and Data Mining V. 1. pp. 2359–2370 (2025)
17. Mirza, M., Osindero, S.: Conditional generative adversarial nets. arXiv preprint arXiv:1411.1784 (2014)
18. Naghavi, M., Ong, K.L., Aali, A., Ababneh, H.S., Abate, Y.H., Abbafati, C., Abbasgholizadeh, R., Abbasian, M., Abbasi-Kangevari, M., Abbastabar, H., et al.: Global burden of 288 causes of death and life expectancy decomposition in 204 countries and territories and 811 subnational locations, 1990–2021: a systematic analysis for the global burden of disease study 2021. The Lancet **403**(10440), 2100–2132 (2024)
19. Salton, G., Wong, A., Yang, C.S.: A vector space model for automatic indexing. Communications of the ACM **18**(11), 613–620 (1975)
20. Steinhubl, S.R., Waalen, J., Edwards, A.M., Ariniello, L.M., Mehta, R.R., Ebner, G.S., Carter, C., Baca-Motes, K., Felicione, E., Sarich, T., et al.: Effect of a home-based wearable continuous ecg monitoring patch on detection of undiagnosed atrial fibrillation: the mstops randomized clinical trial. Jama **320**(2), 146–155 (2018)
21. Takase, B., Ikeda, T., Shimizu, W., Abe, H., Aiba, T., Chinushi, M., Koba, S., Kusano, K., Niwano, S., Takahashi, N., et al.: Jcs/jhrs 2022 guideline on diagnosis and risk assessment of arrhythmia. Circulation Journal **88**(9), 1509–1595 (2024)
22. Trigo, J.D., Chiarugi, F., Alesanco, Á., Martínez-Espronceda, M., Chronaki, C.E., Escayola, J., Martínez, I., García, J.: Standard-compliant real-time transmission of ecgs: Harmonization of iso/ieee 11073-phd and scp-ecg. In: 2009 Annual International Conference of the IEEE Engineering in Medicine and Biology Society. pp. 4635–4638. IEEE (2009)

23. Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A.N., Kaiser, Ł., Polosukhin, I.: Attention is all you need. Advances in neural information processing systems **30** (2017)
24. Wagner, P., Strodthoff, N., Bousseljot, R.D., Kreiseler, D., Lunze, F.I., Samek, W., Schaeffter, T.: Ptb-xl, a large publicly available electrocardiography dataset. Scientific data **7**(1), 1–15 (2020)
25. Wang, L.d., Zhou, W., Xing, Y., Liu, N., Movahedipour, M., Zhou, X.g.: A novel method based on convolutional neural networks for deriving standard 12-lead ecg from serial 3-lead ecg. Frontiers of Information Technology & Electronic Engineering **20**(3), 405–413 (2019)
26. Wang, Z., Bovik, A.C.: Mean squared error: Love it or leave it? a new look at signal fidelity measures. IEEE signal processing magazine **26**(1), 98–117 (2009)
27. Wu, Y., He, K.: Group normalization. In: Proceedings of the European conference on computer vision (ECCV). pp. 3–19 (2018)
28. Zheng, J., Zhang, J., Danioko, S., Yao, H., Guo, H., Rakovski, C.: A 12-lead electrocardiogram database for arrhythmia research covering more than 10,000 patients. Scientific data **7**(1), 48 (2020)
29. Zimetbaum, P., Goldman, A.: Ambulatory arrhythmia monitoring: choosing the right device. Circulation **122**(16), 1629–1636 (2010)