

TSRNET: SIMPLE FRAMEWORK FOR REAL-TIME ECG ANOMALY DETECTION WITH MULTIMODAL TIME AND SPECTROGRAM RESTORATION NETWORK

Nhat-Tan Bui^{1*}, Dinh-Hieu Hoang^{2,3*}, Thinh Phan¹, Minh-Triet Tran^{2,3},
Brijesh Patel⁴, Donald Adjero⁴, Ngan Le¹

¹AICV Lab, University of Arkansas, Arkansas, USA

²University of Science, Vietnam National University, Ho Chi Minh City, Vietnam

³John von Neumann Institute, Vietnam National University, Ho Chi Minh City, Vietnam

⁴West Virginia University, West Virginia, USA

ABSTRACT

The electrocardiogram (ECG) is a valuable signal used to assess various aspects of heart health, such as heart rate and rhythm. It plays a crucial role in identifying cardiac conditions and detecting anomalies in ECG data. However, distinguishing between normal and abnormal ECG signals can be a challenging task. In this paper, we propose an approach that leverages anomaly detection to identify unhealthy conditions using solely normal ECG data for training. Furthermore, to enhance the information available and build a robust system, we suggest considering both the time series and time-frequency domain aspects of the ECG signal. As a result, we introduce a specialized network called the Multimodal Time and Spectrogram Restoration Network (TSRNet) designed specifically for detecting anomalies in ECG signals. TSRNet falls into the category of restoration-based anomaly detection and draws inspiration from both the time series and spectrogram domains. By extracting representations from both domains, TSRNet effectively captures the comprehensive characteristics of the ECG signal. This approach enables the network to learn robust representations with superior discrimination abilities, allowing it to distinguish between normal and abnormal ECG patterns more effectively. Furthermore, we introduce a novel inference method, termed Peak-based Error, that specifically focuses on ECG peaks, a critical component in detecting abnormalities. The experimental result on the large-scale dataset PTB-XL has demonstrated the effectiveness of our approach in ECG anomaly detection, while also prioritizing efficiency by minimizing the number of trainable parameters. Our code is available at <https://github.com/UARK-AICV/TSRNet>.

Index Terms— ECG, Anomaly Detection, Spectrogram

1. INTRODUCTION

The electrocardiogram (ECG) is a crucial time series data that provides valuable insights into the detection and diag-

nosis of cardiac diseases. The morphology observed in each lead of ECG provides insight into the electrical activity occurring within specific segments of the heart. This information proves invaluable in identifying abnormal myocardial conditions [1, 2, 3, 4]. However, focusing solely on the classification of specific diseases through ECG analysis may result in the failure to detect various types of abnormal signals that were not encountered during training. Therefore, our attention shifts towards the detection of anomalies in ECG signals, a strategy with the potential to identify abnormal cardiac conditions in patients, regardless of the diversity and rarity of these conditions.

Unsupervised Deep Learning methods such as [5, 6, 7, 8, 9, 10, 11, 12] have demonstrated significant success in the realm of time series anomaly detection. Currently, BeatGAN [13], and Jiang et al. [14], which fall under the restoration-based (also called reconstruction-based) category, are primarily designed for ECG anomaly detection. Although these methods have attained notable advancements, they exclusively utilize the 1D ECG signals (i.e. time series domain), overlooking the potential of the 2D ECG spectrum (i.e. time-frequency domain). The ECG spectrogram, which represents meaningful time-frequency representation, has been successfully applied in ECG classification tasks [15, 2, 3]. Therefore, exploring the incorporation between time-series and spectrogram offers an alternative avenue for tackling the ECG anomaly detection challenge. As a result, this paper introduces Multimodal Time and Spectrogram Restoration Network (TSRNet), an unsupervised anomaly detection method designed for ECG anomaly detection. TSRNet combines both 1D ECG time series and 2D ECG spectrogram data using attention modules. The primary goals of our method are to extract and fuse meaningful features from both domains to robustly detect anomalies in ECG signals while maintaining simplicity and parameter efficiency. Our contributions include investigating the potential benefits of spectrograms in ECG anomaly detection, developing the TSRNet model, which incorporates a novel Peak-based Error, and evaluating

*Equal contribution

its performance through multiple experiments on the PTB-XL dataset [16].

Notably, TSRNet is *trained exclusively on normal ECG samples* but is *tested on both normal and abnormal samples*. Our network follows a reconstruction-by-inpainting anomaly detection approach [17], where signals with a high restoration error in the inpainting task are identified as anomalies.

2. METHODOLOGY

In this section, we introduce our unsupervised method, which utilizes multi-lead ECGs (i.e. 12-lead) and their corresponding time-frequency spectrogram. Specifically, our method entails the design of a CNN-based model with cross-attention enhancements, exclusively trained on multimodal normal ECGs through a reconstruction task. Furthermore, we introduce a novel inference method that enhances the accuracy of distinguishing abnormal ECGs from normal ones in the inference phase, termed Peak-based Error. Our model is designed in an unsupervised manner, and we will provide detailed training and inference process as follows:

2.1. Training Process

Overall Architecture. Given a multi-lead ECG signal $x_{ecg} \in \mathbb{R}^{D \times N}$, where D is the signal length and N is the number of leads ($N = 12$ in our case). We employ Short Time Fourier Transform (STFT) [18] on x_{ecg} , resulting in the time-frequency spectral signal $x_{spec} \in \mathbb{R}^{D \times N \times H \times W}$. To introduce randomness as well as learn important signal patterns, both x_{ecg} and x_{spec} undergo a Masking-Out procedure (Section 2.1) to be randomly masked-out. Subsequently, they are separately fed into a 1D-CNN and a 2D-CNN encoders to extract meaningful features f_{ecg} and f_{spec} , respectively. Next, a shared weight two-layer cross-attention module is tailored to fuse the features f_{ecg} and f_{spec} from the two different domains to form a fused feature f_{fused} . Finally, the fused feature f_{fused} is processed by a 1D-CNN decoder to reconstruct the *original ECG signal* and predict the *uncertainty*, which implies the difficulty of reconstructing each data point. The overall architecture of the model is depicted in Figure 1.

Masking-Out In this work, we apply a simple yet efficient unsupervised learning approach called inpainting [19]. This approach involves partially masking the input data to encourage the model to learn crucial signal patterns from the training dataset. For 1D time series data, we apply masking across the sequence by uniformly randomly choosing 30% data points of each lead. For the 2D spectrogram, instead of masking pixels or square patches, we use the striping approach. That means, we only choose 20% time intervals, and masking out all values belonging to those intervals. It is noteworthy that the same mask is applied to all the leads. The masked times series and spectrogram are denoted by x_{ecg}^{masked} and x_{spec}^{masked} , respectively.

CNN-based Encoder We employ a 1D CNN encoder ϕ to distill features from the masked time series data x_{ecg}^{masked} and a 2D CNN encoder θ to extract frequency-related information from the masked spectrogram x_{spec}^{masked} .

$$f_{ecg} = \phi(x_{ecg}^{masked}) \text{ and } f_{spec} = \theta(x_{spec}^{masked}) \quad (1)$$

In particular, different leads are treated as different channels of the input. We do not apply padding in the 2D CNN encoder since padding introduces too much noisy information, making the model perform poorly, whereas padding is still used in the 1D CNN encoder. Each encoder consists of 5 convolutional blocks, including one convolution layer, leaky ReLU activation function, and batch normalization layer. At the end of θ , we add a 1D convolution layer to transform the 2D latent feature maps into 1D latent feature maps in such a way that the time dimension is preserved. Therefore, the latent features from both the encoders can be treated as sequences of features.

Cross-Attention Fusion To combine information from time series and spectrum domains extracted by the encoders ϕ and θ , we introduce an innovative cross-attention fusion module to learn robust and informative features as shown in Figure 2. First, we concatenate the two features f_{ecg} and f_{spec} , resulting in feature f_{concat} . Then we apply two shared weight self-attention layers along with skip connections. The motivation to duplicate the attention layers is to enhance the capacity of the model but not too much so that the model is able to learn the necessary long-range relationship but not too strong to attain generalization ability. Finally, the processed sequence is channel-wise refined into a shorter sequence by a linear layer $\mathcal{F}$. This can be considered as a regularization technique, preventing the model from reconstructing even noises.

$$\begin{aligned} f_{concat} &= \text{concat}(f_{ecg}, f_{spec}) \\ f_{att}^1 &= f_{concat} + \text{att}(f_{concat}) \\ f_{att}^2 &= f_{att}^1 + \text{att}(f_{att}^1) \\ f_{fused} &= \mathcal{F}(f_{att}^2) \end{aligned} \quad (2)$$

In this context, $\mathcal{F}$ comprises normalization and a ReLU activation function. The resulting sequence f_{fused} possesses the same length as that of f_{ecg} .

CNN-based Decoder and Objective Function In the final stage, the 1D CNN decoder ω takes f_{fused} as the input and returns the reconstructed ECG signal y , accompanied by an uncertainty σ as a measurement of difficulty of reconstructing each data point, i.e. $(y, \sigma) = \omega(f_{fused})$.

In the training procedure, we optimize the model with the uncertainty-aware restoration loss [14] as follows:

$$\mathcal{L} = \frac{\sum_{i=1}^D \exp(-\sigma^i) \times (y^i - x_{ecg}^i)^2 + \sigma^i}{D} \quad (3)$$

Here, D is the number of data points in time-series. Intuitively, we can consider the first term $\exp(-\sigma^i) \times (y^i - x_{ecg}^i)^2$




Fig. 1. The overall of our proposed TSRNet framework: Starting with a 12-lead ECG time-series signal x_{ecg} , we generate a time-frequency spectral signal x_{spec} using Short Time Fourier Transform (STFT). Both x_{ecg} and x_{spec} undergo a Masking-Out procedure before being processed by a 1D-CNN encoder ϕ and a 2D-CNN encoder θ , resulting in f_{ecg} and f_{spec} , respectively. A shared-weight two-layer cross-attention module is proposed to fuse these features, creating a fused feature f_{fused} . Finally, a 1D-CNN decoder ω processes the fused feature f_{fused} to generate the reconstructed ECG signal y and the uncertainty σ .




Fig. 2. Illustration of Cross-Attention Fusion. Two features, f_{ecg} and f_{spec} , are concatenated to form f_{concat} . Subsequently, this concatenated feature undergoes processing by two shared-weight self-attention layers. $\oplus$ denotes the element-wise addition.




Fig. 3. Visualization of Peak-based Error in the ECG signal. We only consider the R-peaks when detecting the anomaly signal, which is marked by the pink segments.

as a constraint on the relationship between the uncertainty σ and the squared error $(y - x_{ecg})^2$: if the error is big, then the uncertainty must be also big to mitigate the error. In addition, optimizing the second term σ encourages the model to reconstruct the original signal more accurately.

2.2. Inference with Peak-based Error

In our training phase, the model exclusively learns from normal ECG signals. Consequently, during the inference phase, any signals associated with anomalies (i.e., unknown classes) will exhibit higher reconstruction errors compared to those from normal signals. As a result, we can directly use the aforementioned uncertainty-aware restoration loss as an indicator of signal anomalies. In other words, the larger the restoration loss $\mathcal{L}$ is, the more likely the signal has an anomaly. However, recognizing that anomalies are often detected around the R-peaks, we introduce a Peak-based Error, defined as follows:

$$\mathcal{E} = \frac{\sum_{i \in \text{mask}} \exp(-\sigma^i) \times (y^i - x_{ecg}^i)^2 + \sigma^i}{||\text{mask}||} \quad (4)$$

Here, the variable *mask* is defined as a window around R-peaks as shown in Figure 3. The efficacy of this Peak-based

Error strategy is proved by the ablation study shown in the Section 3. The classification of an input ECG signal as normal or abnormal is contingent upon $\mathcal{E}$.

3. EXPERIMENTS

A. Dataset and Evaluation Metrics.

Dataset: We conduct the experiments on PTB-XL dataset [16], including 12-lead ECGs that are 10s in length for each patient. For a fair comparison, we follow the data splitting from the previous work of Jiang et al. [14], which utilizes 8,167 normal ECGs for training, and 912 normal with 1,248 abnormal ECGs for testing.

Metrics: We evaluate the network's accuracy using the area under the Receiver Operating Characteristic curve (AUC), while the network's complexity is measured by the number of trainable parameters (params), and inference time (second).

C. Performance Comparison.

We compare our TSRNet with several SOTA methods in time series anomaly detection, e.g. DAGMM [5], MAD-GAN [6], USAD [7], TranAD [12], Anomaly Transformer [8], Zheng et al. [20], BeatGAN [13], and Jiang et al. [14]. All the performance results are excerpted from Jiang et al. [14]. Those SOTA methods belong to the unsupervised time series anomaly detection approach with the same setting as our method. As a result, we regard them as the most suitable options for comparison in order to highlight the effectiveness of our proposed method.

Analyzing the results presented in Table 1, it is evident that our TSRNet outperforms other SOTA methods while keeping its model size compact. It is worth highlighting that our model achieves the same peak performance as Jiang et al.'s [14]. However, our model is not only designed as a simpler framework but also requires fewer learned parameters. Consequently, our model boasts significantly faster inference speeds (more than $\times 6$), as demonstrated in Table 2. It's worth mentioning that the inference time results are averaged over three runs for each method.

4. CONCLUSION

In this study, we introduced TSRNet, a simple and real-time framework designed for ECG anomaly detection. TSRNet

Table 1. ECG anomaly detection comparison on PTB-XL dataset.

Methods	AUC $\uparrow$
DAGMM (ICLR'18) [5]	0.782
MAD-GAN (ICANN'19) [6]	0.775
USAD (SIGKDD'20) [7]	0.785
TranAD (VLDB'22) [12]	0.788
Anomaly Transformer (ICLR'22) [8]	0.762
Zheng et al. (MICCAI'22) [20]	0.757
BeatGAN (TKDE'22) [13]	0.799
Jiang et al. (MICCAI'23) [14]	0.860
TSRNet (Ours)	0.860

Table 2. Detailed comparison between two closely-matched performance methods, i.e. our TSRNet and Jiang et al. [14] on number of trainable parameters (M) and inference time (s).

Methods	AUC $\uparrow$	Params(M) $\downarrow$	Inference Time(s) $\downarrow$
Jiang et al. [14]	0.860	7.09	0.19
TSRNet (Ours)	0.860	4.39	0.03

is rooted in a restoration-based approach and trained in an unsupervised manner using inpainting mechanism. Notably, TSRNet introduces an innovative perspective by emphasizing the substantial impact of both time-frequency and time-series domains on ECG anomaly detection. Additionally, we introduced an effective cross-attention mechanism to merge information from both modalities, enabling the model to leverage the valuable characteristics embedded in each. For efficient inference, we proposed a Peak-based Error strategy, which prioritizes R-peaks to classify an input ECG signal as either normal or abnormal. The experimental results demonstrate that TSRNet achieves SOTA performance (AUC = 0.860) while maintaining real-time inference capabilities (33.3 fps) and a compact model size (4.39M params).

Acknowledgement: Nhat-Tan Bui, Thinh Phan, and Ngan Le are supported by the National Science Foundation (NSF) under Award No OIA-1946391 RII Track-1, NSF 1920920 RII Track 2 FEC, NSF 2223793 EFRI BRAID, NSF 2119691 AI SUSTAIN, NSF 2236302. Minh-Triet Tran is sponsored by VNU-HCM (DS2020-42-01). Dinh-Hieu Hoang is funded by VINIF, VINBIGDATA (code VINIF.2022.ThS.JVN.04).

5. REFERENCES

- [1] Alan Davies and Alwyn Scott, *Starting to Read ECGs*, Springer, 2014. **1**
- [2] Thinh Phan, Duc Le, et al., “Multimodality Multi-Lead ECG Arrhythmia Classification using Self-Supervised Learning,” in *IEEE EMBS BHI*, 2022. **1**
- [3] Minh Duc Le and Vidhiwar others Rathour, “Multi-module Recurrent Convolutional Neural Network with Transformer Encoder for ECG Arrhythmia Classification,” in *IEEE EMBS BHI*, 2021. **1**
- [4] Duc Le, Sang Truong, Patel Brijesh, Donald Adjeroh, and Ngan Le, “sCL-ST: Supervised Contrastive Learning with Semantic Transformations for Multiple Lead ECG Arrhythmia Classification,” *IEEE Journal of Biomedical and Health Informatics*, 2023. **1**
- [5] Bo Zong, Qi Song, et al., “Deep Autoencoding Gaussian Mixture Model for Unsupervised Anomaly Detection,” in *ICLR*, 2018. **1, 3, 4**
- [6] Dan Li, Dacheng Chen, et al., “MAD-GAN: Multivariate anomaly detection for time series data with generative adversarial networks,” in *ICANN*, 2019. **1, 3, 4**
- [7] Julien Audibert, Pietro Michiardi, et al., “USAD: UnSupervised Anomaly Detection on Multivariate Time Series,” in *SIGKDD*, 2020. **1, 3, 4**
- [8] Jiehui Xu, Haixu Wu, Jianmin Wang, and Mingsheng Long, “Anomaly Transformer: Time Series Anomaly Detection with Association Discrepancy,” in *ICLR*, 2022. **1, 3, 4**
- [9] Chuxu Zhang, Dongjin Song, et al., “A Deep Neural Network for Unsupervised Anomaly Detection and Diagnosis in Multivariate Time Series Data,” in *AAAI*, 2019. **1**
- [10] Ailin Deng and Bryan Hooi, “Graph Neural Network-Based Anomaly Detection in Multivariate Time Series,” in *AAAI*, 2021. **1**
- [11] Yuxin Zhang, Yiqiang Chen, and othres, “Unsupervised Deep Anomaly Detection for Multi-Sensor Time-Series Signals,” *IEEE TKDE*, 2023. **1**
- [12] Shreshth Tuli, Giuliano Casale, and Nicholas R Jennings, “TranAD: Deep Transformer Networks for Anomaly Detection in Multivariate Time Series Data,” *arXiv preprint arXiv:2201.07284*, 2022. **1, 3, 4**
- [13] Shenghua Liu, Bin Zhou, Quan Ding, Bryan Hooi, Zhengbo Zhang, Huawei Shen, and Xueqi Cheng, “Time Series Anomaly Detection With Adversarial Reconstruction Networks,” *IEEE Transactions on Knowledge and Data Engineering*, 2023. **1, 3, 4**
- [14] Aofan Jiang, Chaoqin Huang, and othres, “Multi-scale Cross-restoration Framework for Electrocardiogram Anomaly Detection,” in *MICCAI*, 2023. **1, 2, 3, 4**
- [15] Xiaoming Sun, Pengfei Liu, et al., “Automatic classification of electrocardiogram signals based on transfer learning and continuous wavelet transform,” *Ecological Informatics*, 2022. **1**
- [16] Patrick Wagner, Nils Strodthoff, et al., “PTB-XL, a large publicly available electrocardiography dataset,” *Scientific Data*, 2020. **2, 3**
- [17] Vitjan Zavrtanik, Matej Kristan, et al., “Reconstruction by inpainting for visual anomaly detection,” *PR*, 2021. **2**
- [18] Daniel Griffin and Jae Lim, “Signal estimation from modified short-time Fourier transform,” *IEEE Transactions on Acoustics, Speech, and Signal Processing*, 1984. **2**
- [19] Chaoqin Huang, Qinwei Xu, et al., “Self-Supervised Masking for Unsupervised Anomaly Detection and Localization,” *IEEE TM*, 2022. **2**
- [20] Yaojia Zheng, Zhouwu Liu, et al., “Task-oriented Self-supervised Learning for Anomaly Detection in Electroencephalography,” in *MICCAI*, 2022. **3, 4**