



Multi-scale Cross-restoration Framework for Electrocardiogram Anomaly Detection

Aofan Jiang^{1,2}, Chaoqin Huang^{1,2,4}, Qing Cao³, Shuang Wu³, Zi Zeng³,
Kang Chen³, Ya Zhang^{1,2}, and Yanfeng Wang^{1,2}(✉)

¹ Shanghai Jiao Tong University, Shanghai, China

{stillunnamed, huangchaoqin, ya_zhang, wangyanfeng}@sjtu.edu.cn

² Shanghai AI Laboratory, Shanghai, China

³ Ruijin Hospital, Shanghai Jiao Tong University School of Medicine,
Shanghai, China

{cq30553, ck11208}@rjh.com.cn, {shuang-renata, zengzidoc}@sjtu.edu.cn

⁴ National University of Singapore, Singapore, Singapore

Abstract. Electrocardiogram (ECG) is a widely used diagnostic tool for detecting heart conditions. Rare cardiac diseases may be underdiagnosed using traditional ECG analysis, considering that no training dataset can exhaust all possible cardiac disorders. This paper proposes using anomaly detection to identify any unhealthy status, with normal ECGs solely for training. However, detecting anomalies in ECG can be challenging due to significant inter-individual differences and anomalies present in both global rhythm and local morphology. To address this challenge, this paper introduces a novel multi-scale cross-restoration framework for ECG anomaly detection and localization that considers both local and global ECG characteristics. The proposed framework employs a two-branch autoencoder to facilitate multi-scale feature learning through a masking and restoration process, with one branch focusing on global features from the entire ECG and the other on local features from heartbeat-level details, mimicking the diagnostic process of cardiologists. Anomalies are identified by their high restoration errors. To evaluate the performance on a large number of individuals, this paper introduces a new challenging benchmark with signal point-level ground truths annotated by experienced cardiologists. The proposed method demonstrates state-of-the-art performance on this benchmark and two other well-known ECG datasets. The benchmark dataset and source code are available at: <https://github.com/MediaBrain-SJTU/ECGAD>

Keywords: Anomaly Detection · Electrocardiogram

A. Jiang and C. Huang—Equal contribution.

Supplementary Information The online version contains supplementary material available at https://doi.org/10.1007/978-3-031-43907-0_9.

© The Author(s), under exclusive license to Springer Nature Switzerland AG 2023
H. Greenspan et al. (Eds.): MICCAI 2023, LNCS 14220, pp. 87–97, 2023.
https://doi.org/10.1007/978-3-031-43907-0_9

1 Introduction

The electrocardiogram (ECG) is a monitoring tool widely used to evaluate the heart status of patients and provide information on cardiac electrophysiology. Developing automated analysis systems capable of detecting and identifying abnormal signals is crucial in light of the importance of ECGs in medical diagnosis and the need to ease the workload of clinicians. However, training a classifier on labeled ECGs that focus on specific diseases may not recognize new abnormal statuses that were not encountered during training, given the diversity and rarity of cardiac diseases [8, 16, 23]. On the other hand, anomaly detection, which is trained only on normal healthy data, can identify any potential abnormal status and avoid the failure to detect rare cardiac diseases [10, 17, 21].

The current anomaly detection techniques, including one-class discriminative approaches [2, 14], reconstruction-based approaches [15, 30], and self-supervised learning-based approaches [3, 26], all operate under the assumption that models trained solely on normal data will struggle to process anomalous data and thus the substantial drop in performance presents an indication of anomalies. While anomaly detection has been widely used in the medical field to analyze medical images [12, 24] and time-series data [18, 29], detecting anomalies in ECG data is particularly challenging due to the substantial inter-individual differences and the presence of anomalies in both global rhythm and local morphology. So far, few studies have investigated anomaly detection in ECG [11, 29]. TSL [29] uses expert knowledge-guided amplitude- and frequency-based data transformations to simulate anomalies for different individuals. BeatGAN [11] employs a generative adversarial network to separately reconstruct normalized heartbeats instead of the entire raw ECG signal. While BeatGAN alleviates individual differences, it neglects the important global rhythm information of the ECG.

This paper proposes a novel multi-scale cross-restoration framework for ECG anomaly detection and localization. To our best knowledge, this is the first work to integrate both local and global characteristics for ECG anomaly detection. To take into account multi-scale data, the framework adopts a two-branch autoencoder architecture, with one branch focusing on global features from the entire ECG and the other on local features from heartbeat-level details. A multi-scale cross-attention module is introduced, which learns to combine the two feature types for making the final prediction. This module imitates the diagnostic process followed by experienced cardiologists who carefully examine both the entire ECG and individual heartbeats to detect abnormalities in both the overall rhythm and the specific local morphology of the signal [7]. Each of the branches employs a masking and restoration strategy, *i.e.*, the model learns how to perform temporal-dependent signal inpainting from the adjacent unmasked regions within a specific individual. Such context-aware restoration has the advantage of making the restoration less susceptible to individual differences. During testing, anomalies are identified as samples or regions with high restoration errors.

To comprehensively evaluate the performance of the proposed method on a large number of individuals, we adopt the public PTB-XL database [22] with only patient-level diagnoses, and ask experienced cardiologists to provide signal

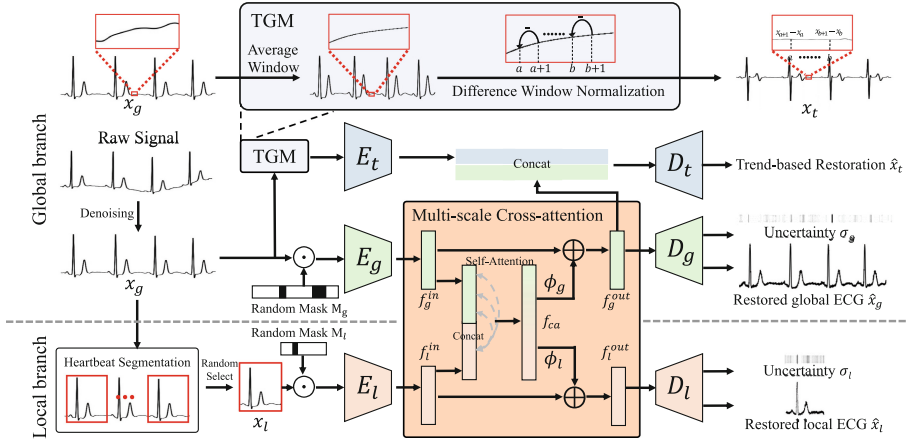


Fig. 1. The multi-scale cross-restoration framework for ECG anomaly detection.

point-level localization annotations. The resulting dataset is then introduced as a large-scale challenging benchmark for ECG anomaly detection and localization. The proposed method is evaluated on this challenging benchmark as well as on two traditional ECG anomaly detection benchmarks [6, 13]. The experimental results have shown that the proposed method outperforms several state-of-the-art methods for both anomaly detection and localization, highlighting its potential for real-world clinical diagnosis.

2 Method

In this paper, we focus on unsupervised anomaly detection and localization on ECGs, training based on only normal ECG data. Formally, given a set of N normal ECGs denoted as $\{x_i, i = 1, \dots, N\}$, where $x_i \in \mathbb{R}^D$ represents the vectorized representation of the i -th ECG consisting of D signal points, the objective is to train a computational model capable of identifying whether a new ECG is normal or anomalous, and localize the regions of anomalies in abnormal ECGs.

2.1 Multi-scale Cross-restoration

In Fig. 1, we present an overview of our two-branch framework for ECG anomaly detection. One branch is responsible for learning global ECG features, while the other focuses on local heartbeat details. Our framework comprises four main components: (i) masking and encoding, (ii) multi-scale cross-attention module, (iii) uncertainty-aware restoration, and (iv) trend generation module. We provide detailed explanations of each of these components in the following sections.

Masking and Encoding. Given a pair consisting of a global ECG signal $x_g \in \mathbb{R}^D$ and a randomly selected local heartbeat $x_l \in \mathbb{R}^d$ segmented from x_g for

training, as shown in Fig. 1, we apply two random masks, M_g and M_l , to mask x_g and x_l , respectively. To enable multi-scale feature learning, M_l is applied to a consecutive small region to facilitate detail restoration, while M_g is applied to several distinct regions distributed throughout the whole sequence to facilitate global rhythm restoration. The masked signals are processed separately by global and local encoders, E_g and E_l , resulting in global feature $f_g^{in} = E_g(x_g \odot M_g)$ and local feature $f_l^{in} = E_l(x_l \odot M_l)$, where \odot denotes the element-wise product.

Multi-scale Cross-attention. To capture the relationship between global and local features, we use the self-attention mechanism [20] on the concatenated feature of f_g^{in} and f_l^{in} . Specifically, the attention mechanism is expressed as $Attention(Q, K, V) = \text{softmax}(\frac{QK^T}{\sqrt{d_k}})V$, where Q, K, V are identical input terms, while $\sqrt{d_k}$ is the square root of the feature dimension used as a scaling factor. Self-attention is achieved by setting $Q = K = V = \text{concat}(f_g^{in}, f_l^{in})$. The cross-attention feature, f_{ca} , is obtained from the self-attention mechanism, which dynamically weighs the importance of each element in the combined feature. To obtain the final outputs of the global and local features, f_g^{out} and f_l^{out} , containing cross-scale information, we consider residual connections: $f_g^{out} = f_g^{in} + \phi_g(f_{ca})$, $f_l^{out} = f_l^{in} + \phi_l(f_{ca})$, where $\phi_g(\cdot)$ and $\phi_l(\cdot)$ are MLP architectures with two fully connected layers.

Uncertainty-Aware Restoration. Targeting signal restorations, features of f_g^{out} and f_l^{out} are decoded by two decoders, D_g and D_l , to obtain restored signals \hat{x}_g and \hat{x}_l , respectively, along with corresponding restoration uncertainty maps σ_g and σ_l measuring the difficulty of restoration for various signal points, where $\hat{x}_g, \sigma_g = D_g(f_g^{out})$, $\hat{x}_l, \sigma_l = D_l(f_l^{out})$. An uncertainty-aware restoration loss is used to incorporate restoration uncertainty into the loss functions,

$$\mathcal{L}_{global} = \sum_{k=1}^D \left\{ \frac{(x_g^k - \hat{x}_g^k)^2}{\sigma_g^k} + \log \sigma_g^k \right\}, \quad \mathcal{L}_{local} = \sum_{k=1}^d \left\{ \frac{(x_l^k - \hat{x}_l^k)^2}{\sigma_l^k} + \log \sigma_l^k \right\}, \quad (1)$$

where for each function, the first term is normalized by the corresponding uncertainty, and the second term prevents predicting a large uncertainty for all restoration pixels following [12]. The superscript k represents the position of the k -th element of the signal. It is worth noting that, unlike [12], the uncertainty-aware loss is used for restoration, but not for reconstruction.

Trend Generation Module. The trend generation module (TGM) illustrated in Fig. 1 generates a smooth time-series trend $x_t \in \mathbb{R}^D$ by removing signal details, which is represented as the smooth difference between adjacent time-series signal points. An autoencoder (E_t and D_t) encodes the trend information into $E_t(x_t)$, which are concatenated with the global feature f_g^{out} to restore the global ECG $\hat{x}_t = D_t(\text{concat}(E_t(x_t), f_g^{out}))$. The restoration loss is defined as the Euclidean distance between x_g and \hat{x}_t , $\mathcal{L}_{trend} = \sum_{k=1}^D (x_g^k - \hat{x}_t^k)^2$. This process guides global feature learning using time-series trend information, emphasizing rhythm characteristics while de-emphasizing morphological details.

Loss Function. The final loss function for optimizing our model during the training process can be written as

$$\mathcal{L} = \mathcal{L}_{global} + \alpha \mathcal{L}_{local} + \beta \mathcal{L}_{trend}, \quad (2)$$

where α and β are trade-off parameters weighting the loss function. For simplicity, we adopt $\alpha = \beta = 1.0$ as the default.

2.2 Anomaly Score Measurement

For each test sample x , local ECGs from the segmented heartbeat set $\{x_{l,m}, m = 1, \dots, M\}$ are paired with the global ECG x_g one at a time as inputs. The anomaly score $\mathcal{A}(x)$ is calculated to estimate the abnormality,

$$\mathcal{A}(x) = \sum_{k=1}^D \frac{(x_g^k - \hat{x}_g^k)^2}{\sigma_g^k} + \sum_{m=1}^M \sum_{k=1}^d \frac{(x_{l,m}^k - \hat{x}_{l,m}^k)^2}{\sigma_{l,m}^k} + \sum_{k=1}^D (x_g^k - \hat{x}_t^k)^2, \quad (3)$$

where the three terms correspond to global restoration, local restoration, and trend restoration, respectively. For localization, an anomaly score map is generated in the same way as Eq. (3), but without summing over the signal points. The anomalies are indicated by relatively large anomaly scores, and vice versa.

3 Experiments

Datasets. Three publicly available ECG datasets are used to evaluate the proposed method, including PTB-XL [22], MIT-BIH [13], and Keogh ECG [6].

- **PTB-XL** database includes clinical 12-lead ECGs that are 10 s in length for each patient, with only patient-level annotations. To build a new challenging anomaly detection and localization benchmark, 8167 normal ECGs are used for training, while 912 normal and 1248 abnormal ECGs are used for testing. We provide signal point-level annotations of 400 ECGs, including 22 different abnormal types, that were annotated by two experienced cardiologists. To our best knowledge, we are the first to explore ECG anomaly detection and localization across various patients on such a complex and large-scale database.
- **MIT-BIH** arrhythmia dataset divides the ECGs from 44 patients into independent heartbeats based on the annotated R-peak position, following [11]. 62436 normal heartbeats are used for training, while 17343 normal and 9764 abnormal beats are used for testing, with heartbeat-level annotations.
- **Keogh ECG** dataset includes 7 ECGs from independent patients, evaluating anomaly localization with signal point-level annotations. For each ECG, there is an anomaly subsequence that corresponds to a pre-ventricular contraction, while the remaining sequence is used as normal data to train the model. The ECGs are partitioned into fixed-length sequences of 320 by a sliding window with a stride of 40 during training and 160 during testing.

Table 1. Anomaly detection and anomaly localization results on PTB-XL database. Results are shown in the patient-level AUC for anomaly detection and the signal point-level AUC for anomaly localization, respectively. The best-performing method is in **bold**, and the second-best is underlined.

Method	Year	detection	localization
DAGMM [30]	2018	0.782	0.688
MADGAN [9]	2019	0.775	0.708
USAD [1]	2020	0.785	0.683
TranAD [18]	2022	0.788	0.685
AnoTran [25]	2022	0.762	0.641
TSL [29]	2022	0.757	0.509
BeatGAN [11]	2022	<u>0.799</u>	<u>0.715</u>
Ours	2023	0.860	0.747

Table 2. Anomaly detection results on MIT-BIH dataset, comparing with state-of-the-arts. Results are shown in terms of the AUC and F1 score for heartbeat-level classification. The best-performing method is in **bold**, and the second-best is underlined.

Method	Year	F1	AUC
DAGMM [30]	2018	0.677	0.700
MSCRED [27]	2019	0.778	0.627
USAD [1]	2020	0.384	0.352
TranAD [18]	2022	0.621	0.742
AnoTran [25]	2022	0.650	0.770
TSL [29]	2022	0.750	0.894
BeatGAN [11]	2022	<u>0.816</u>	<u>0.945</u>
Ours	2023	0.883	0.969

Table 3. Anomaly localization results on Keogh ECG [6] dataset, comparing with several state-of-the-arts. Results are shown in the signal point-level AUC. The best-performing method is in **bold**, and the second-best is underlined.

Methods	Year	A	B	C	D	E	F	G	Avg
DAGMM [30]	2018	0.672	0.612	0.805	0.713	0.457	0.662	0.676	0.657
MSCRED [27]	2019	0.667	0.633	0.798	0.714	0.461	0.746	0.659	0.668
MADGAN [9]	2019	0.688	0.702	0.833	0.664	0.463	0.692	0.678	0.674
USAD [1]	2020	0.667	0.616	0.795	0.715	0.462	0.649	0.680	0.655
GDN [4]	2021	0.695	0.611	0.790	0.674	0.458	0.648	0.671	0.650
CAE-M [28]	2021	0.657	0.618	0.802	0.715	0.457	0.708	0.671	0.661
TranAD [18]	2022	0.647	0.623	<u>0.820</u>	0.720	0.446	0.780	0.680	0.674
AnoTran [25]	2022	0.739	0.502	0.792	<u>0.799</u>	0.498	0.748	0.711	0.684
BeatGAN [11]	2022	<u>0.803</u>	0.623	0.783	0.747	<u>0.506</u>	0.757	0.852	<u>0.724</u>
Ours	2023	0.832	<u>0.641</u>	0.819	0.815	0.543	<u>0.760</u>	<u>0.833</u>	0.749

Evaluation Protocols. The performance of anomaly detection and localization is quantified using the area under the Receiver Operating Characteristic curve (AUC), with a higher AUC value indicating a better method. To ensure comparability across different annotation levels, we used patient-level, heartbeat-level, and signal point-level AUC for each respective setting. For heartbeat-level classification, the F1 score is also reported following [11].

Implementation Details. The ECG is pre-processed by a Butterworth filter and Notch filter [19] to remove high-frequency noise and eliminate ECG baseline wander. The R-peaks are detected with an adaptive threshold following [5], which

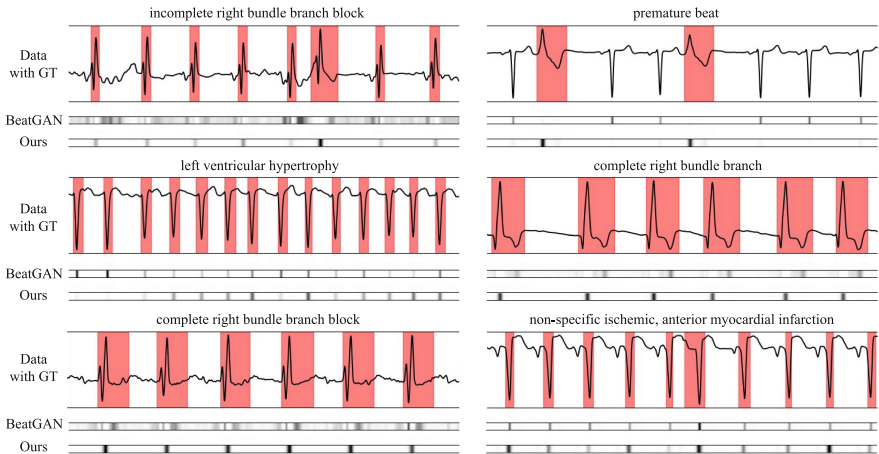


Fig. 2. Anomaly localization visualization on PTB-XL with different abnormal types. Ground truths are highlighted in red boxes on the ECG data, and anomaly localization results for each case, compared with the state-of-the-art method, are attached below. (Color figure online)

does not require any learnable parameters. The positions of the detected R-peaks are then used to segment the ECG sequence into a set of heartbeats.

We use a convolutional-based autoencoder, following the architecture proposed in [11]. The model is trained using the AdamW optimizer with an initial learning rate of $1e-4$ and a weight decay coefficient of $1e-5$ for 50 epochs on a single NVIDIA GTX 3090 GPU, with a single cycle of cosine learning rate used for decay scheduling. The batch size is set to 32. During testing, the model requires 2365M GPU memory and achieves an inference speed of 4.2 fps.

3.1 Comparisons with State-of-the-Arts

We compare our method with several time-series anomaly detection methods, including heartbeat-level detection method BeatGAN [11], patient-level detection method TSL [29], and several signal point-level anomaly localization methods [1, 4, 9, 18, 25, 27, 28, 30]. For a fair comparison, we re-trained all the methods under the same experimental setup. For those methods originally designed for signal point-level tasks only [1, 9, 18, 25, 30], we use the mean value of anomaly localization results as their heartbeat-level or patient-level anomaly scores.

Anomaly Detection. The anomaly detection performance on PTB-XL is summarized in Table 1. The proposed method achieves 86.0% AUC in patient-level anomaly detection and outperforms all baselines by a large margin (10.3%). Table 2 displays the comparison results on MIT-BIH, where the proposed method achieves a heartbeat-level AUC of 96.9%, showing an improvement of 2.4% over the state-of-the-art BeatGAN (94.5%). Furthermore, the F1-score of the proposed method is 88.3%, which is 6.7% higher than BeatGAN (81.6%).

Table 4. Ablation studies on PTB-XL dataset. Factors under analysis are: the masking and restoring (MR), the multi-scale cross-attention (MC), the uncertainty loss function (UL), and the trend generation module (TGM). Results are shown in the patient-level AUC in % of five runs. The best-performing method is in **bold**.

MR	MC	UL	TGM	AUC
				70.4 \pm 0.3
✓				80.4 \pm 0.7
	✓			80.3 \pm 0.3
		✓		72.8 \pm 2.0
			✓	71.2 \pm 0.5
✓	✓			84.8 \pm 0.8
✓	✓	✓		85.2 \pm 0.4
✓	✓	✓	✓	86.0\pm0.1

Table 5. Sensitivity analysis w.r.t. mask ratio on PTB-XL dataset. Results are shown in the patient-level AUC of five runs. The best-performing method is in **bold**, and the second-best is underlined.

Mask Ratio	AUC
0%	80.2 \pm 0.2
10%	85.2 \pm 0.2
20%	<u>85.5\pm0.3</u>
30%	86.0\pm0.1
40%	84.9 \pm 0.3
50%	83.8 \pm 0.1
60%	82.9 \pm 0.1
70%	75.8 \pm 1.0

Anomaly Localization. Table 1 presents the results of anomaly localization on our proposed benchmark for multiple individuals. The proposed method achieves a signal point-level AUC of 74.7%, outperforming all baselines (3.2% higher than BeatGAN). It is worth noting that TSL, which is not designed for localization, shows poor performance in this task. Table 3 shows the signal point-level anomaly localization results for each independent individual on Keogh ECG. Overall, the proposed method achieves the best or second-best performance compared to other methods on six subsets and the highest mean AUC among all subsets (74.9%, 2.5% higher than BeatGAN), indicating its effectiveness. The proposed method shows a lower standard deviation (± 10.5) across the seven subsets compared to TranAD (± 11.3) and BeatGAN (± 11.0), which indicates good generalizability of the proposed method across different subsets.

Anomaly Localization Visualization. We present visualization results of anomaly localization on several samples from our proposed benchmark in Fig. 2, with ground truths annotated by experienced cardiologists. Regions with higher anomaly scores are indicated by darker colors. Our proposed method outperforms BeatGAN in accurately localizing various types of ECG anomalies, including both periodic and episodic anomalies, such as incomplete right bundle branch block and premature beats. Our method though provides narrower localization results than ground truths, as it is highly sensitive to abrupt unusual changes in signal values, but still represents the important areas for anomaly identification, a fact confirmed by experienced cardiologists.

3.2 Ablation Study and Sensitivity Analysis

Ablation studies were conducted on PTB-XL to confirm the effectiveness of individual components of the proposed method. Table 4 shows that each module

contributes positively to the overall performance of the framework. When none of the modules were employed, the method becomes a ECG reconstruction approach with a naive L2 loss and lacks cross-attention in multi-scale data. When individually adding the MR, MC, UL, and TGM modules to the baseline model without any of them, the AUC values improve from 70.4% to 80.4%, 80.3%, 72.8%, and 71.2%, respectively, demonstrating the effectiveness of each module. Moreover, as the modules are added in sequence, the performance improves step by step from 70.4% to 86.0% in AUC, highlighting the combined impact of all modules on the proposed framework.

We conduct a sensitivity analysis on the mask ratio, as shown in Table 5. Restoration with a 0% masking ratio can be regarded as reconstruction, which takes an entire sample as input and its target is to output the input sample. Results indicate that the model's performance first improves and then declines as the mask ratio increases from 0% to 70%. This trend is due to the fact that a low mask ratio can limit the model's feature learning ability during restoration, while a high ratio can make it increasingly difficult to restore the masked regions. Therefore, there is a trade-off between maximizing the model's potential and ensuring a reasonable restoration difficulty. The optimal mask ratio is 30%, which achieves the highest anomaly detection result (86.0% in AUC).

4 Conclusion

This paper proposes a novel framework for ECG anomaly detection, where features of the entire ECG and local heartbeats are combined with a masking-restoration process to detect anomalies, simulating the diagnostic process of cardiologists. A challenging benchmark, with signal point-level annotations provided by experienced cardiologists, is proposed, facilitating future research in ECG anomaly localization. The proposed method outperforms state-of-the-art methods, highlighting its potential in real-world clinical diagnosis.

Acknowledgement. This work is supported by the National Key R&D Program of China (No. 2022ZD0160702), STCSM (No. 22511106101, No. 18DZ2270700, No. 21DZ1100100), 111 plan (No. BP0719010), the Youth Science Fund of National Natural Science Foundation of China (No. 7210040772) and National Facility for Translational Medicine (Shanghai) (No. TMSK-2021-501), and State Key Laboratory of UHD Video and Audio Production and Presentation.

References

1. Audibert, J., Michiardi, P., Guyard, F., Marti, S., Zuluaga, M.A.: Usad: unsupervised anomaly detection on multivariate time series. In: Proceedings of the 26th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining, pp. 3395–3404 (2020)
2. Chalapathy, R., Menon, A.K., Chawla, S.: Robust, deep and inductive anomaly detection. In: Ceci, M., Hollmén, J., Todorovski, L., Vens, C., Džeroski, S. (eds.) ECML PKDD 2017. LNCS (LNAI), vol. 10534, pp. 36–51. Springer, Cham (2017). https://doi.org/10.1007/978-3-319-71249-9_3

3. Chen, L., Bentley, P., Mori, K., Misawa, K., Fujiwara, M., Rueckert, D.: Self-supervised learning for medical image analysis using image context restoration. *Med. Image Anal.* **58**, 101539 (2019)
4. Deng, A., Hooi, B.: Graph neural network-based anomaly detection in multivariate time series. In: *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 35, pp. 4027–4035 (2021)
5. van Gent, P., Farah, H., van Nes, N., van Arem, B.: Analysing noisy driver physiology real-time using off-the-shelf sensors: Heart rate analysis software from the taking the fast lane project. *J. Open Res. Softw.* **7**(1) (2019)
6. Keogh, E., Lin, J., Fu, A.: Hot sax: finding the most unusual time series subsequence: Algorithms and applications. In: *Proceedings of the IEEE International Conference on Data Mining*, pp. 440–449. Citeseer (2004)
7. Khan, M.G.: Step-by-step method for accurate electrocardiogram interpretation. In: *Rapid ECG Interpretation*, pp. 25–80. Humana Press, Totowa, NJ (2008)
8. Kiranyaz, S., Ince, T., Gabbouj, M.: Real-time patient-specific ECG classification by 1-d convolutional neural networks. *IEEE Trans. Biomed. Eng.* **63**(3), 664–675 (2015)
9. Li, D., Chen, D., Jin, B., Shi, L., Goh, J., Ng, S.-K.: MAD-GAN: multivariate anomaly detection for time series data with generative adversarial networks. In: Tetko, I.V., Kůrková, V., Karpov, P., Theis, F. (eds.) *ICANN 2019*. LNCS, vol. 11730, pp. 703–716. Springer, Cham (2019). https://doi.org/10.1007/978-3-030-30490-4_56
10. Li, H., Boulanger, P.: A survey of heart anomaly detection using ambulatory electrocardiogram (ECG). *Sensors* **20**(5), 1461 (2020)
11. Liu, S., et al.: Time series anomaly detection with adversarial reconstruction networks. *IEEE Trans. Knowl. Data Eng.* (2022)
12. Mao, Y., Xue, F.-F., Wang, R., Zhang, J., Zheng, W.-S., Liu, H.: Abnormality detection in chest X-Ray images using uncertainty prediction autoencoders. In: Martel, A.L., et al. (eds.) *MICCAI 2020*. LNCS, vol. 12266, pp. 529–538. Springer, Cham (2020). https://doi.org/10.1007/978-3-030-59725-2_51
13. Moody, G.B., Mark, R.G.: The impact of the mit-bih arrhythmia database. *IEEE Eng. Med. Biol. Mag.* **20**(3), 45–50 (2001)
14. Ruff, L., et al.: Deep one-class classification. In: *International Conference on Machine Learning*, pp. 4393–4402. PMLR (2018)
15. Schlegl, T., Seeböck, P., Waldstein, S.M., Langs, G., Schmidt-Erfurth, U.: f-anogan: fast unsupervised anomaly detection with generative adversarial networks. *Med. Image Anal.* **54**, 30–44 (2019)
16. Shaker, A.M., Tantawi, M., Shedeed, H.A., Tolba, M.F.: Generalization of convolutional neural networks for ECG classification using generative adversarial networks. *IEEE Access* **8**, 35592–35605 (2020)
17. Shen, L., Yu, Z., Ma, Q., Kwok, J.T.: Time series anomaly detection with multi-resolution ensemble decoding. In: *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 35, pp. 9567–9575 (2021)
18. Tuli, S., Casale, G., Jennings, N.R.: Tranad: deep transformer networks for anomaly detection in multivariate time series data. In: *International Conference on Very Large Databases 15*(6), pp. 1201–1214 (2022)
19. Van Gent, P., Farah, H., Van Nes, N., Van Arem, B.: Heartpy: a novel heart rate algorithm for the analysis of noisy signals. *Transport. Res. F: Traffic Psychol. Behav.* **66**, 368–378 (2019)
20. Vaswani, A., et al.: Attention is all you need. In: *Advances in Neural Information Processing Systems 30* (2017)

21. Venkatesan, C., Karthigaikumar, P., Paul, A., Satheeskumaran, S., Kumar, R.: ECG signal preprocessing and SVM classifier-based abnormality detection in remote healthcare applications. *IEEE Access* **6**, 9767–9773 (2018)
22. Wagner, P., et al.: Ptb-xl, a large publicly available electrocardiography dataset. *Sci. Data* **7**(1), 1–15 (2020)
23. Wang, J., et al.: Automated ECG classification using a non-local convolutional block attention module. *Comput. Methods Programs Biomed.* **203**, 106006 (2021)
24. Wolleb, J., Bieder, F., Sandkühler, R., Cattin, P.C.: Diffusion models for medical anomaly detection. In: *Medical Image Computing and Computer Assisted Intervention-MICCAI 2022*, pp. 35–45. Springer, Cham (2022). https://doi.org/10.1007/978-3-031-16452-1_4
25. Xu, J., Wu, H., Wang, J., Long, M.: Anomaly transformer: time series anomaly detection with association discrepancy. In: *International Conference on Learning Representations* (2022)
26. Ye, F., Huang, C., Cao, J., Li, M., Zhang, Y., Lu, C.: Attribute restoration framework for anomaly detection. *IEEE Trans. Multimedia* **24**, 116–127 (2022)
27. Zhang, C., et al.: A deep neural network for unsupervised anomaly detection and diagnosis in multivariate time series data. In: *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 33, pp. 1409–1416 (2019)
28. Zhang, Y., Chen, Y., Wang, J., Pan, Z.: Unsupervised deep anomaly detection for multi-sensor time-series signals. *IEEE Trans. Knowl. Data Eng.* **35**(2), 2118–2132 (2021)
29. Zheng, Y., Liu, Z., Mo, R., Chen, Z., Zheng, W.s., Wang, R.: Task-oriented self-supervised learning for anomaly detection in electroencephalography. In: *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pp. 193–203. Springer, Cham (2022). https://doi.org/10.1007/978-3-031-16452-1_19
30. Zong, B., et al.: Deep autoencoding gaussian mixture model for unsupervised anomaly detection. In: *International Conference on Learning Representations* (2018)