A Novel Temporal Generative Adversarial Network for Electrocardiography Anomaly Detection

Jing Qin[a], Fujie Gao[b], Zumin Wang[b,*], David C Wong[c], Zhibin Zhao[d], Samuel D Relton[e], Hui Fang[f]

[a]*College of Software Engineering, Dalian University, Dalian, China.*
[b]*College of Information Engineering, Dalian University, Dalian, China.*
[c]*Department of Computer Science and Centre for Health Informatics, University of Manchester, Manchester, U.K.*
[d]*School of Mechanical Engineering, Xi'an Jiaotong University, Xi'an, China.*
[e]*Leeds Institute of Health Sciences, University of Leeds, Leeds, U.K.*
[f]*Department of Computer Science, Loughborough University, Loughborough, U.K.*

Abstract

Cardiac abnormality detection from Electrocardiogram (ECG) signals is a common task for cardiologists. To facilitate efficient and objective detection, automated ECG classification by using deep learning based methods have been developed in recent years. Despite their impressive performance, these methods perform poorly when presented with cardiac abnormalities that are not well represented, or absent, in the training data. To this end, we propose a novel one-class classification based ECG anomaly detection generative adversarial network (GAN). Specifically, we embedded a Bi-directional Long-Short Term Memory (Bi-LSTM) layer into a GAN architecture and used a mini-batch discrimination training strategy in the

*Corresponding author

Email addresses: qinjing@dlu.edu.cn (Jing Qin), gaofujie@s.dlu.edu.cn (Fujie Gao), wangzumin@dlu.edu.cn (Zumin Wang), david.wong@manchester.ac.uk (David C Wong), zhaozhibin@stu.xjtu.edu.cn (Zhibin Zhao), s.d.relton@leeds.ac.uk (Samuel D Relton), h.fang@lboro.ac.uk (Hui Fang)

discriminator to synthesis ECG signals. Our method generates samples to match the data distribution from normal signals of healthy group so that a generalised anomaly detector can be built reliably. The experimental results demonstrate our method outperforms several state-of-the-art semi-supervised learning based ECG anomaly detection algorithms and robustly detects the unknown anomaly class in the MIT-BIH arrhythmia database. Experiments show that our method achieves the accuracy of 95.5% and AUC of 95.9% which outperforms the most competitive baseline by 0.7% and 1.7% respectively. Our method may prove to be a helpful diagnostic method for helping cardiologists identify arrhythmias.
Keywords: Electrocardiogram, Generative Adversarial Networks, Semi-supervised learning, One-class classification, MIT-BIH

---

1. Introduction

Electrocardiograms (ECG) are routinely used in clinical practice to identify cardiac abnormalities. They provide an inexpensive and non-invasive tool to facilitate accurate diagnosis of cardiovascular conditions, e.g., arrhythmia, coronary heart disease, heart attack and cardiomyopathy [1]. Correct interpretation of ECGs is time-consuming and requires expertise. Consequently, to identify cardiac abnormalities, researchers have worked to create automatic ECG signal classification algorithms and systems that require little human input [2, 3]. Deep neural network (DNN) models have supported state-of-the-art performances in these algorithms [4, 5, 6].

While DNN classifiers have primarily been trained in a supervised manner [6, 7, 8], one-class classification ECG anomaly detection is also an active research topic [9, 10, 11], which involves training a classifier with the presence of data from only one class [12]. It is typically used when there is high class imbalance. This often occurs in ECG datasets, in which there are many examples of normal ECG and far fewer examples of the myriad types of ECG abnormalities. With the development of advanced ECG anomaly detection, a more generalisable and reliable system could handle large variations in real clinical scenarios compared to those supervised methods [13].

An early attempt to detect ECG anomaly based used One-Class Support Vector Machines to determine the characteristics of a normal ECG [2]. To improve the performance when processing the high-dimensionality of the ECG data, more modern approaches using neural generative models, including variational autoencoder (VAE) and generative adversarial networks (GAN), have been proposed for anomaly detection [14, 15]. For example, Shin et al. [11] deployed AnoGAN [16] to synthesize ECG signals so that the augmented data could improve the model generalisation to enhance the anomaly detection. Although improved performance is evidenced in this work, the temporal constraint of ECG signals is not explored during the generative modelling training process.

In this paper, we propose a novel anomaly detection GAN (ECG-ADGAN) method for one-class ECG classification. Specifically, we embed a Bi-directional Long-Short Term Memory (Bi-LSTM) network into the GAN generator so that the temporal patterns in ECG signals are preserved. Further, we use a mini-batch strategy during GAN discriminator training to prevent the occurrence of mode collapse and improve the stability of the GAN training. The contributions of our paper are as follows:

- A novel one class classification GAN method is proposed to detect cardiac abnormalities. When training a cardiac anomaly detector in a semi-supervised learning manner, the model generalisation is enhanced by aligning synthesised data to real data distribution of ECG signals from a healthy control group.

- A Bi-LSTM layer is embedded into the generator of the neuro generative model to enforce the temporal constraint for better model training.

- Our experiments demonstrate that the mini-batch training guarantees the convergence stability of our model so that a reliable anomaly detector is achieved.

The remainder of the paper is structured as follows. We present related work in Section 2. The proposed method is explained in detail in Section 3. In Section 4, experimental results demonstrate the effectiveness of our proposed method when compared to several state-of-the-art semi-supervised learning algorithms. Finally, Section 5 presents conclusions and future work.

2. Related Work

2.1. Supervised learning based ECG arrhythmia classification

Supervised learning based automatic ECG signal classification has been developed for many years [17, 18]. Recently, DNN based methods have become dominant to achieve almost perfect performance when working on small-scale dataset, e.g., MIT BIH Arrhythmia Database [19]. For instances, Xu et al. [4] designed a 3-layer DNN algorithm to classify well-aligned heartbeat segments; Chu et al. [5] proposed to extract cross-lead features via a 2D CNN and use a Long Short-Term Memory (LSTM) network for the classification; Maweu et al. [20] propose a modular framework, CNN Explanations Framework for ECG Signals (CEFEs), for interpretable explanations; Mousavi et al. [21] introduced to extract feature representations via an autoencoder to train a bidirectional Recurrent Neural Network (RNN) for the arrhythmia detection.

Latest research focuses on improving model reliability and its generalisability when adapting a model to new datasets or a new working environment. Sellami et al. [7] used a residual network (Resnet) architecture and designed a batched weight loss function to deal with data imbalance problem. Zhang et al. [22] proposed a spatio-temporal attention-based convolutional recurrent neural network (STA-CRNN) to focus on representative features along both spatial and temporal axes. To provide a large-scale dataset that facilitates the build of reliable models, Alday et al. [23] organized the computing in cardiology challenge with large amount of training data which is categorized into 27 cardiac abnormalities. Working on the PhysioNet-2021 dataset, Shang et al. [24] proposed a multi-source adversarial feature learning to enhance model geralisability via extracting domain-invariant and discriminative representations.

Although models trained in a supervised manner achieved state-of-the-art performances in automated ECG signal classification tasks, they may be less reliable for identifying unknown cardiac abnormalities or those with few training examples. Thus, this is the main challenge to achieve a generalisable DNN model that can be applied in real applications [10].

2.2. GAN-based Anomaly Detection

Anomaly detection (AD) refers to a binary classification task that identifies data outliers which are significantly deviated from the distribution of normal data [25]. Unlike supervised learning based methods, AD typically use semi-supervised learning where outlier samples are rare or totally unavailable. Ruff et al. [15] summarise three main types of AD models: classification based methods, probabilistic methods and reconstruction methods. Each type has its advantages to deal with their specific tasks related to their assumptions. For instances, one class SVM [26] is one of the classification based methods, which assumes sufficient normal data points are provided so that the distance from the decision boundary to the origin is maximised; while probabilistic methods, e.g., Gaussian Mixture Model [27] and Kernel Density Estimation [28], assume the normal data obeys either a parametric or a non-parametric distribution so that the unseen anomalies located with low-density areas are detected. Due to the success of deep neural networks on applications with complex data topology, neural generative models, which are considered as hybrid probabilistic and reconstruction methods, have recently been emerging for the AD tasks.

One such deep learning method is GAN, which was first developed to generate realistic synthetic images [29, 30]. Due to its capability to synthesise samples that match a generative data distribution with high-dimensional and complex data, GAN has also been applied to various anomaly detection tasks [16, 25, 15]. In [16], the deep convolutional GAN architecture proposed in [30] was used to fit data distribution of optical coherence tomography (OCT) images of retina so that anomaly regions in these OCT images are identified.

In [25], Liu et al. also proposed a GAN based method for anomaly detection. They designed a multiple-generator structure to avoid mode collapsing problem so that the reference distribution of the entire dataset is fitted well by the GANs. As summarised in [15], GANs have advantage over many traditional anomaly detection methods since it combines both the strengths of probabilistic and reconstruction methods.

GAN networks have recently been witnessed in ECG anomaly detection tasks [14, 11]. For instances, in [14], BeatGAN was proposed to automatically detect anomalous beats so that clinicians could improve their diagnosis efficiency based on the arrhythmic ECG signals. In this work, they prove that GAN based methods achieve much better detection accuracy when compared to autoencoder based methods; in [11], AnoGAN was used to synthesise more data samples to improve the decision boundary for anomaly detection.

3. ECG-ADGAN model

In this section, we describe our proposed model in detail. We first illustrate and present the network architecture of our model. Following this, we explain our mini-batch discriminator training strategy, showing why it improves GAN convergence stability. Finally, we provide the model implementation detail for its reproducibility.

3.1. Network architecture and training stage

The ECG-ADGAN model architecture is depicted in Fig. 1. Similar to other GAN models, our model consists of a generator to synthesise the ECG signal and a discriminator to identify anomalous signals. The generator takes a 100-dimensional random vector as its input to feed into a Bi-LSTM layer that acts to extract temporal features in a latent space. Subsequently, this representation is processed by several 1-D convolutional and upsampling layers before it is connected to a dense layer to generate a synthesised ECG signal. To ensure the synthesised data fit well with the distribution of real samples, a discriminator is updated to distinguish synthesised signals from real signals. As illustrated in Fig. 1, the discriminator is composed of four 1-D convolutional layers followed by max pooling layer and one dense layer with 64 neurons.

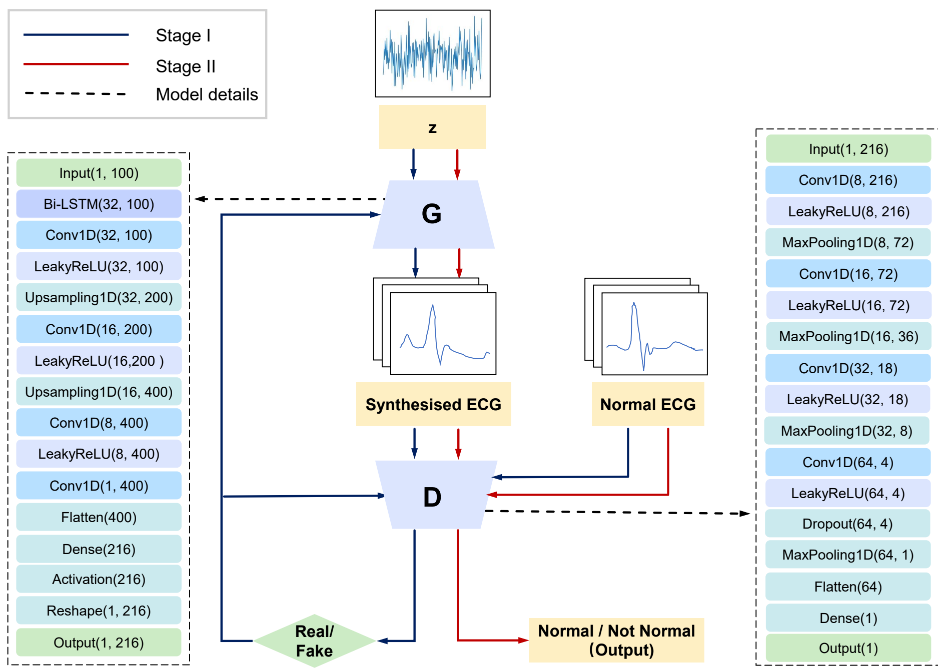

Figure 1: Architecture of the ECG-ADGAN network and its training stages. G and D represent a generator to and a discriminator respectively. In Stage I, the GAN is trained by iteratively updating both the generator and the discriminator which is used to distinguish between real/fake ECG sequences. In Stage II, the generator is frozen and the discriminator is further trained to detect anomaly ECG sequences from normal ECG signals.

The training of the anomaly detector is composed of two stages: In Stage I, we train the GAN by iteratively updating both the generator and the discriminator to ensure the generator can synthesize more realistic ECG signals. The loss function for training the generators and discriminators is the KL scatter (relative entropy):

$$T_{\text{function for training the}}(p\|q)=\sum_{i=1}^{n} p\left(x_i\right) \log \left(\frac{p\left(x_i\right)}{q\left(x_i\right)}\right) \tag{1}$$

where $p(x)$ represents the distribution of the true ECG samples and $q(x)$ the distribution of the generated ECG samples. The loss function convergence of both the generator and discriminator during the adversarial training in Stage I is illustrated in Fig. 2. It is observed that the generator does not learn the real data distribution at the beginning of the training since the discriminator could easily separate the generated data from the real data. During the 850 iterations, the generator gradually learns the real data distribution pattern to ensure that the adversarial learning between the generator and the discriminator reaches a Nash Equilibrium.

In Stage II, we freeze the generator and train the anomaly detection classifier with more iterations to ensure the detector becomes more distinctive to identify anomaly signals. These two-stage training strategy ensures the convergence of the GAN to align with the real ECG normal data distribution (illustrated in Fig. 2(a)) and achieves a more discriminant classifier for anomaly detection (illustrated in Fig. 2(b)).

Compared to other GAN based ECG anomaly detection models (e.g., [14, 11]), our model embeds a Bi-LSTM layer that creates strong temporal constraints on the signal synthesis; we hypothesize that this will lead to

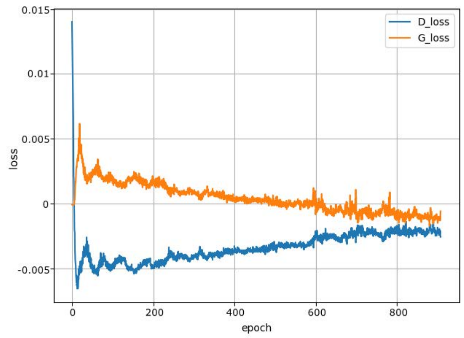

(a)

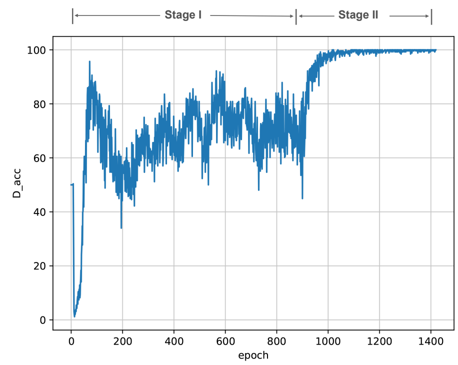

(b)

Figure 2: (a) Loss of the models in Stage I. (b) The discriminator accuracy with both Stage I and Stage II.

better overall performance.

The Bi-LSTM operations can be expressed in the following equations:

$$k_t = f\left(w_1 x_t + w_2 k_{t-1}\right) \tag{2}$$

$$k_t^{'} = f\left(w_3 x_t + w_5 k_{t+1}^{'}\right) \tag{3}$$

$$o_t = g\left(w_4 k_t + w_6 k_t^{'}\right) \tag{4}$$

where x_t denotes the input of current timestamp t, k_t denotes the output of forward layer at current timestamp, k_{t-1} denotes the output of forward layer at previous timestamp, k_t' denotes the output of backward layer at current timestamp, k_{t+1}' denotes the output of backward layer at next timestamp, and o_t is the final output by considering both the Bi-LSTM layers.

As illustrated in Fig. 3, the Bi-LSTM layer learns sets of weights to exploit the temporal correlations on both forward and backward directions of temporal signals. Thus, this layer could significantly contribute to the data distribution alignment for signals with strong temporal patterns. To justify this, we overlay 70 beats of ECG signals containing normal and premature ventricular contraction (PVC) rhythms in a plot by aligning them based on R-peak. Illustrated in the Fig. 4, it is observed that normal ECG signals have a clear temporal pattern with a small variation that could facilitate their identification and distinguish from those with PVC.

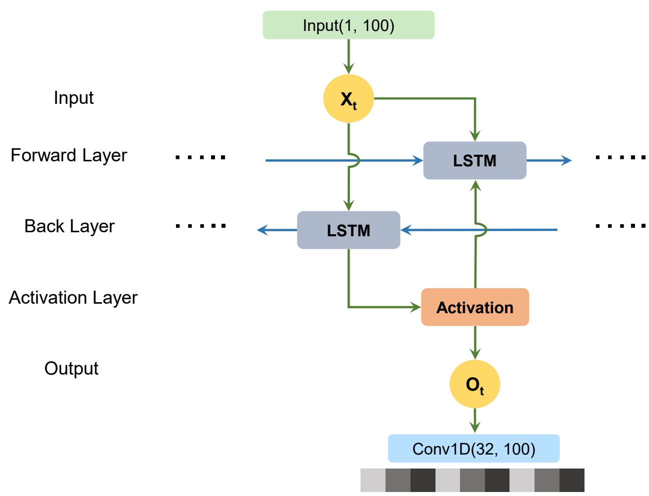

Figure 3: Architecture of Bi-LSTM in ECG-ADGAN

3.2. Mini-batch discriminator

The ultimate goal of training a GAN is to find the Nash Equilibrium of a zero-sum game. Simply put, a generator and a discriminator iteratively minimise their cost functions till convergence. However, in practice, it is hard to reach the Nash Equilibrium because the cost functions are non-convex and

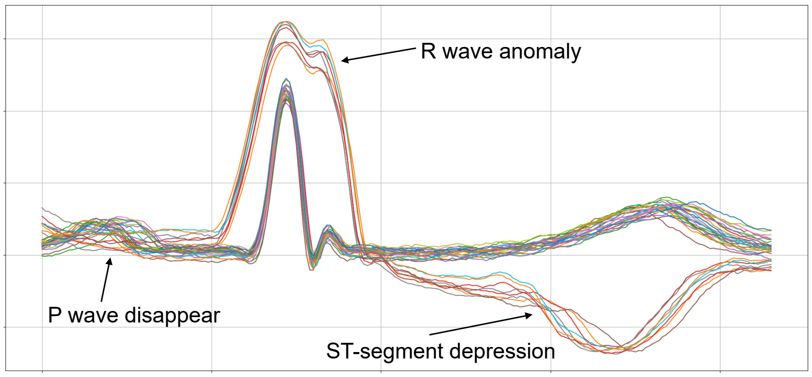

Figure 4: Heartbeat overlay drawing. We overlay 70 beats of ECG signals containing normal and premature ventricular contraction (PVC) rhythms in a plot by aligning them based on R-peak.

the parameter space is high-dimensional [31, 32]. Further, adversarial training of GANs penalises poorly generated samples, which makes the generator more likely to synthesise samples that have been validated, thereby losing data diversity and falling into a mode collapse. To avoid this here, we use mini-batch discrimination during Stage I training. Mini-batch discrimination was initially proposed in [31]. As expressed in Eq. (5) - (7), extra feature representations are created by calculating the similarity between batch samples. These are concatenated with original feature representations to reach a more reliable convergence of the discriminator.

$$o\left(\boldsymbol{x}_i\right)_b = \sum_{j=1}^{n} c_b\left(\boldsymbol{x}_i, \boldsymbol{x}_j\right) \in \mathbb{R} \tag{5}$$

$$o\left(\boldsymbol{x}_i\right) = \left[o\left(\boldsymbol{x}_i\right)_1, o\left(\boldsymbol{x}_i\right)_2, \ldots, o\left(\boldsymbol{x}_i\right)_B\right] \in \mathbb{R}^B \tag{6}$$

$$o(\mathbf{X}) \in \mathbb{R}^{n \times B} \tag{7}$$

where c denotes the $L1$ distance between each sample features x_i, n denotes the feature dimension, and o denotes the sum of c, b denotes the batch index and B is the batch size. The combination of $o(x)_b$ with original sample features improves convergence during GAN training as well as diversifies synthesised samples.

3.3. Implementation of the model

The pseudocode of our model is presented in Algorithm 1, and the code is available at `https://github.com/gaofujie1997/ECG-ADGAN`. In addition, we visualize the synthesised samples of the generator at several key training stages that are correspondent to significant value changes of loss functions. Fig. 5 illustrates an example synthesized (a) at the beginning of the training, the output of the generator is a random vector. (b) after 100 iterations, the output reveals some of the characteristics of the ECG. (c) after 300 iterations, the generator learns the characteristics of the ECG well, but the waveform is still not smooth. (d) after 850 iterations, generator can output realistic synthetic ECGs. There are a clear convergence pattern since the generator gradually aligns well with normal ECG signals.

4. Experiment

4.1. Experimental settings

The algorithms were deployed on a tower workstation with a Core i9-10920X processor (3.5GHz), Quadro RTX6000 graphics card (24GB of video

Algorithm 1 ECG-ADGAN

Input: Normal ECG Signal

Output: Normal or Not Normal Flags

Require: GAN

1: Initialize the parameters of G and D.

2: Fix G, randomly select n real ECG samples and n synthesised samples from G using the defined noise z, which are used to train D to distinguish between real and fake as much as possible.

3: Update D k times followed by 1 times update of G.

4: When the change in the loss of D and G is less than 5%, stop training G but continue training D until the loss of D no longer improves, saving the model of D as M.

5: Using model M for anomaly detection.

6: **return** Normal or Not Normal

memory) running on 32GB RAM, and a deep learning environment using Window Server 2019 + Python 3.8 + Keras 2.2. In this paper, the MIT-BIH arrhythmia database [19], an open-source dataset collected by the Massachusetts Institute of Technology, was selected as the experimental dataset. The MIT-BIH arrhythmia database contains 48 dual channel ECGs of 30 minutes in length, with a sampling rate of 360 Hz. The location of R-peak was annotated by using a simple slope-sensitive QRS detector and the type of beats was annotated by two cardiologists [19].

To preprocess the data, we used the method in [33] to find R-peak and segment each recording into epochs containing a single cardiac cycle and the length of each cycle was sampled to 216 dimensions. Following this, as shown

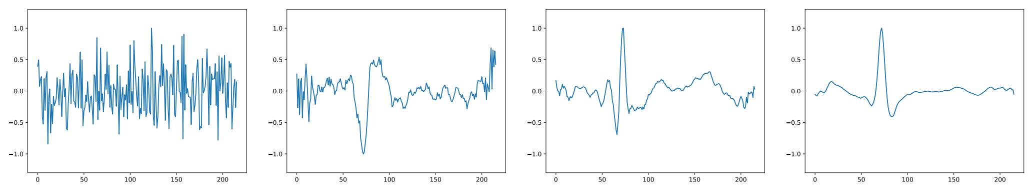

(a) (b) (c) (d)

Figure 5: ECG generated at different stages. (a) at the beginning of the training, (b) after 100 iterations, (c) after 300 iterations and (d) after 850 iterations. It shows our generator converges to model data distribution of normal ECG sequences which can be further used for ECG anomaly detection.

in Fig 6, we applied wavelet threshold denoising [34]. Specifically, the Eq. (8) and (9) defined the adaptive threshold in the noise removing process for obtaining the input vector of our model.

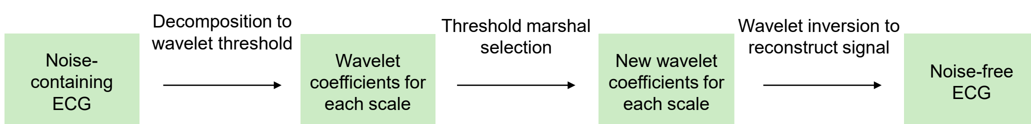

Figure 6: Wavelet threshold denoising method. There are three steps to preprocess original ECG sequences to noise-free sequences before feeding into our proposed model.

$$t = \sigma \sqrt{2 \log (M)} \tag{8}$$

$$\omega_t = \begin{cases} [\operatorname{sgn}(\omega)](|\omega| - t), & |\omega| \geq t \\ 0, & |\omega| < t \end{cases} \tag{9}$$

where M is the length of the high-frequency coefficients for each scale of the wavelet, scale is 9, $\sigma = 0.674$ is a constant value, t is the threshold, and

ω are the high-frequency wavelet coefficients.

We follow the principle defined in [35] to divide the data into a training set and a test set for fair comparison. Particularly, any heartbeat records with participant IDs in the test set are not included in the training dataset no matter if they are normal or abnormal heartbeats. This resolves the potential bias if similar patterns from same IDs have been learned at the training stage. The division of the data sets on MIT-BIH is presented in Table 1. For the evaluation comparison, we use the same test data partitioned in [11].

Table 1: Division of the training and test sets on MIT-BIH

Data sets	Records	Type	Number of heartbeats
Train	101, 103, 112, 113, 115, 117, 121, 122, 123, 230	Normal heartbeats	85717
Test	100, 102, 104, 105, 106, 107, 108, 109, 111, 114, 116, 118, 119, 124, 200, 201, 202, 203, 205, 207, 208, 209, 210, 212, 213, 214, 215, 217, 219, 220, 221, 222, 223, 228, 231, 232, 233, 234	Normal and Abnormal heartbeats	2005

In Table 2, we provide the abnormality categories in MIT-BIH. The AAMI standards are adopted in our experiments and results have been compared to those of [11] and [14]. We treat samples in Supraventricular (S), Ventricular (V), Fusion (F) and Unknown (Q) as anomalous samples and all normal heartbeats (N) as normal samples to make a binary classification task. For the method evaluation, accuracy, precision, recall, F1 and AUC are used as metrics to compare our model with benchmark methods.

Table 2: Composition of Train Data and Test Data

AAMI category	MIT-BIH category	Train	Test
Normal(N)	N, L, R	85717	1000
Supraventricular(S)	A, a, J, S, j, e	-	330
Ventricular(V)	V, E	-	330
Fusion(F)	F	-	330
Unknown(Q)	/, f, Q	-	15

4.2. Comparison with state-of-the-art semi-supervised learning methods

In our work, we compared the performance of our model to OC-SVM [2], AnoGAN [16], GANomaly [36] and BeatGAN [14]. For fair comparison, we re-implemented these benchmark methods on the same dataset. As shown in Table 3, compared to OC-SVM method, all the synthesised based methods achieved much better performance. Our method achieved the best performance, with an accuracy of 95.5%, precision of 96.9%, recall of 91.8%, F1 of 94.3% and AUC of 95.9%. Overall, the model's better performance provides evidence that the temporal constraints with Bi-LSTM and mini-batch training are useful.

4.3. Experiments for unknown cardiac abnormality detection

We further conducted a cross-validation experiment to demonstrate our model generalisation over supervised learning method when facing unknown cardiac abnormalities. Specifically, as shown in Table 4, we train a model by excluding one abnormality category each time and evaluate the model performance by using this category in the test stage. For the balanced data

Table 3: Performance of each method

Methods	Accuracy	Precision	Recall	F1	AUC
OC-SVM	0.804	0.765	0.805	0.785	0.794
AnoGAN	0.927	0.873	0.910	0.891	0.894
GANomaly	0.934	0.914	0.895	0.904	0.937
BeatGAN	0.948	0.953	0.907	0.929	0.943
ECG-ADGAN	**0.955**	**0.969**	**0.918**	**0.943**	**0.959**

test, we randomly sampled 330 normal ECG segments to conduct each experiment. The CNN+LSTM supervised learning model proposed in [37] was selected for the comparison. In Table 4, it was observed that the ECG-ADGAN had better generalisability to detect unknown cardiac abnormalities although CNN+LSTM model achieved high accuracy on detect known cardiac abnormalities. Further, the confusion matrix of the four experiments are illustrated in Fig. 7 to demonstrate how our model detected each unknown cardiac abnormalities. When the four classes were treated as unknown abnormalities respectively, the misclassification rates were much lower when using the semi-supervised learning method. Thus, the matrix further confirmed that the semi-supervised learning method achieved more generalised model for the unknown cardiac abnormality detection.

4.4. Ablation experiment

To prove the effectiveness of Bi-LSTM and mini-batch discrimination training, we have conducted an ablation study by taking them off from our proposed model respectively. The results in Table 5 show that, although

Table 4: Performance of CNN+LSTM and ECG-ADGAN

		CNN+LSTM								ECG-ADGAN			
Train Label		Test with train labels				Test with new abnormal label				Acc	Precision	Recall	F1
Normal	Abnormal	Acc	Precision	Recall	F1	Acc	Precision	Recall	F1				
N	S,V,Q	0.970	0.973	0.967	0.970	0.753	0.973	0.676	0.798	0.939	0.967	0.917	0.941
N	S,F,Q	0.976	0.979	0.973	0.976	0.924	0.979	0.883	0.928	0.965	0.967	0.964	0.965
N	V,F,Q	0.976	0.982	0.970	0.976	0.697	0.982	0.626	0.764	0.903	0.967	0.858	0.909
N	S,V,F	0.967	1.000	0.938	0.968	0.800	1.000	0.714	0.833	0.966	0.933	1.000	0.966

both elements have contributed to the performance improvement, Bi-LSTM plays a relatively more important role in our proposed method compared to the mini-batch discrimination since the temporal pattern preservation makes the most useful constraint to model ECG sequences. With the contributions from both, we achieved the best performance in terms of all the evaluation metrics.

Table 5: Ablation experiment

Bi-LSTM	Mini-batch Discrimination	Accrucy	Precision	Recall	F1	AUC
×	×	0.929	0.887	0.895	0.891	0.897
×	✓	0.936	0.923	0.903	0.913	0.931
✓	×	0.941	0.943	0.905	0.924	0.939
✓	✓	0.955	0.969	0.918	0.943	0.959

5. Discussions and Conclusions

In this paper, we have proposed a novel ECG anomaly detection GAN method for the one-class classification analysis task since it is more suitable to model the high-dimensional ECG signal data distribution. Compared to previous GAN based ECG anomaly detection work, our method has better

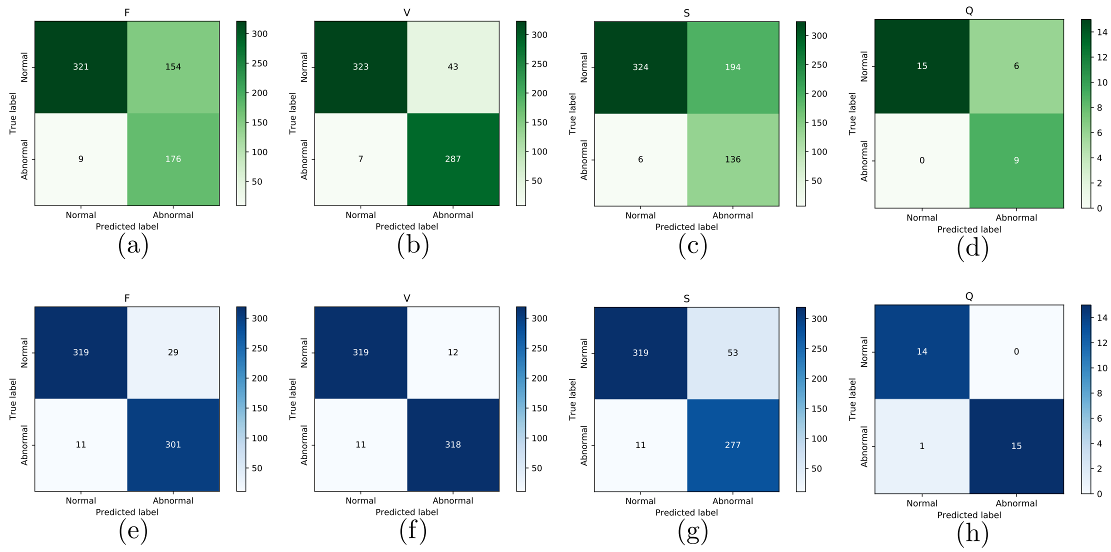

Figure 7: Confusion matrix of Table 4. (a)-(d) are CNN+LSTM: (a)F, (b)V, (c)S, (d)Q. (e)-(f) are ECG-ADGAN: (e)F, (f)V, (g)S, (h)Q. It shows that our proposed method achieved more generalised model when compared to a supervised CNN+LSTM model.

performance because of two distinctive designs, which include: (i) we embed Bi-LSTM layer in our architecture so that the temporal relation is explored in the learning process; and (ii) we use a mini-batch training strategy to stabilise the GAN convergence. The experimental results demonstrate our method outperforms several state-of-the-art semi-supervised learning based ECG anomaly detection algorithms and reached 95.5% accuracy and 95.9% AUC on the MIT-BIH arrhythmia database. Further, we conduct a cross-validation experiment to leave one abnormality class as unknown abnormality each time to train a supervised learning model. Based on the experimental results, it proves that the generalisability of our model is more superior to identify unknown abnormalities when compared to supervised learning methods.

Although our proposed method offered an unsupervised way to detect ECG anomalies and demonstrated an improved performance compared to the state-of-the-art anomaly detection methods, the method cannot distinguish specific types of anomalies without human annotation. In future work, we will investigate to identify these specific anomaly types with few human annotations in a weakly-supervised manner. In addition, the training method of ECG-ADGAN can be extended to anomaly detection on other time-series data, such as log anomaly detection, audio anomaly detection and video anomaly detection.

Data Availability

We have disclosed data and source code in our work to facilitate subsequent research and make contributions to the community.

Declaration of Competing Interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Acknowledgement

This work was supported by the Youth Fund Project of the National Nature Fund of China under Grant 62002038.

References

[1] A. Y. Hannun, P. Rajpurkar, M. Haghpanahi, G. H. Tison, C. Bourn, M. P. Turakhia, A. Y. Ng, Cardiologist-level arrhythmia detection and classification in ambulatory electrocardiograms using a deep neural network, Nature medicine 25 (1) (2019) 65–69.

[2] S.-M. Woo, H.-J. Lee, B.-J. Kang, S.-W. Ban, Ecg signal monitoring using one-class support vector machine, in: Proceedings of the 9th WSEAS International Conference on Applications of Electrical Engineering, Penang, Malaysia, March, 2010, pp. 23–25.

[3] S. Kiranyaz, T. Ince, M. Gabbouj, Real-time patient-specific ecg classification by 1-d convolutional neural networks, IEEE Transactions on Biomedical Engineering 63 (3) (2016) 664–675.

[4] S. S. Xu, M.-W. Mak, C.-C. Cheung, Towards end-to-end ecg classification with raw signal extraction and deep neural networks, IEEE journal of biomedical and health informatics 23 (4) (2018) 1574–1584.

[5] J. Chu, H. Wang, W. Lu, A novel two-lead arrhythmia classification system based on cnn and lstm, Journal of Mechanics in Medicine and Biology 19 (03) (2019) 1950004.

[6] Z. Zhao, H. Fang, S. D. Relton, R. Yan, Y. Liu, Z. Li, J. Qin, D. C. Wong, Adaptive lead weighted resnet trained with different duration signals for classifying 12-lead ecgs, in: 2020 Computing in Cardiology, IEEE, 2020, pp. 1–4.

[7] A. Sellami, H. Hwang, A robust deep convolutional neural network with batch-weighted loss for heartbeat classification, Expert Systems with Applications 122 (2019) 75–84.

[8] G. A. Tadesse, H. Javed, K. Weldemariam, Y. Liu, J. Liu, J. Chen, T. Zhu, Deepmi: Deep multi-lead ecg fusion for identifying myocardial infarction and its occurrence-time, Artificial Intelligence in Medicine 121 (2021) 102192.

[9] T. Golany, K. Radinsky, Pgans: Personalized generative adversarial networks for ecg synthesis to improve patient-specific deep ecg classification, in: Proceedings of the AAAI Conference on Artificial Intelligence, Vol. 33, 2019, pp. 557–564.

[10] M. Braei, S. Wagner, Anomaly detection in univariate time-series: A survey on the state-of-the-art, arXiv preprint arXiv:2004.00433 (2020).

[11] D.-H. Shin, R. C. Park, K. Chung, Decision boundary-based anomaly detection model using improved anogan from ecg data, IEEE Access 8 (2020) 108664–108674.

[12] P. Perera, P. Oza, V. M. Patel, One-class classification: A survey, arXiv preprint arXiv:2101.03064 (2021).

[13] H. Li, P. Boulanger, A survey of heart anomaly detection using ambulatory electrocardiogram (ecg), Sensors 20 (5) (2020) 1461.

[14] B. Zhou, S. Liu, B. Hooi, X. Cheng, J. Ye, Beatgan: Anomalous rhythm detection using adversarially generated time series., in: IJCAI, 2019, pp. 4433–4439.

[15] L. Ruff, J. R. Kauffmann, R. A. Vandermeulen, G. Montavon, W. Samek, M. Kloft, T. G. Dietterich, K.-R. Müller, A unifying review of deep and shallow anomaly detection, Proceedings of the IEEE 109 (5) (2021) 756–795.

[16] T. Schlegl, P. Seeböck, S. M. Waldstein, U. Schmidt-Erfurth, G. Langs, Unsupervised anomaly detection with generative adversarial networks to guide marker discovery, in: International conference on information processing in medical imaging, Springer, 2017, pp. 146–157.

[17] M. I. Owis, A. H. Abou-Zied, A.-B. Youssef, Y. M. Kadah, Study of features based on nonlinear dynamical modeling in ecg arrhythmia detection and classification, IEEE transactions on Biomedical Engineering 49 (7) (2002) 733–736.

[18] T. Mar, S. Zaunseder, J. P. Martínez, M. Llamedo, R. Poll, Optimization of ecg classification by means of feature selection, IEEE transactions on Biomedical Engineering 58 (8) (2011) 2168–2177.

[19] G. B. Moody, R. G. Mark, The impact of the mit-bih arrhythmia database, IEEE Engineering in Medicine and Biology Magazine 20 (3) (2001) 45–50.

[20] B. M. Maweu, S. Dakshit, R. Shamsuddin, B. Prabhakaran, Cefes: A cnn explainable framework for ecg signals, Artificial Intelligence in Medicine 115 (2021) 102059.

[21] S. Mousavi, F. Afghah, Inter-and intra-patient ecg heartbeat classification for arrhythmia detection: a sequence to sequence deep learning

approach, in: ICASSP 2019-2019 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP), IEEE, 2019, pp. 1308–1312.

[22] J. Zhang, A. Liu, M. Gao, X. Chen, X. Zhang, X. Chen, Ecg-based multi-class arrhythmia detection using spatio-temporal attention-based convolutional recurrent neural network, Artificial Intelligence in Medicine 106 (2020) 101856.

[23] E. A. P. Alday, A. Gu, A. J. Shah, C. Robichaux, A.-K. I. Wong, C. Liu, F. Liu, A. B. Rad, A. Elola, S. Seyedi, et al., Classification of 12-lead ecgs: the physionet/computing in cardiology challenge 2020, Physiological measurement 41 (12) (2020) 124003.

[24] Z. Shang, Z. Zhao, H. Fang, S. Relton, D. Murphy, Z. Hancox, R. Yan, D. Wong, Deep discriminative domain generalization with adversarial feature learning for classifying ecg signals, in: 2021 Computing in Cardiology (CinC), Vol. 48, IEEE, 2021, pp. 1–4.

[25] Y. Liu, Z. Li, C. Zhou, Y. Jiang, J. Sun, M. Wang, X. He, Generative adversarial active learning for unsupervised outlier detection, IEEE Transactions on Knowledge and Data Engineering 32 (8) (2019) 1517–1528.

[26] Y. Chen, X. S. Zhou, T. S. Huang, One-class svm for learning in image retrieval, in: Proceedings 2001 International Conference on Image Processing (Cat. No. 01CH37205), Vol. 1, IEEE, 2001, pp. 34–37.

[27] R. Mohammadi-Ghazi, Y. M. Marzouk, O. Büyüköztürk, Conditional classifiers and boosted conditional gaussian mixture model for novelty detection, Pattern recognition 81 (2018) 601–614.

[28] Y.-C. Chen, A tutorial on kernel density estimation and recent advances, Biostatistics & Epidemiology 1 (1) (2017) 161–187.

[29] I. Goodfellow, J. Pouget-Abadie, M. Mirza, B. Xu, D. Warde-Farley, S. Ozair, A. Courville, Y. Bengio, Generative adversarial nets, in: Z. Ghahramani, M. Welling, C. Cortes, N. Lawrence, K. Q. Weinberger (Eds.), Advances in Neural Information Processing Systems, Vol. 27, Curran Associates, Inc., 2014.

[30] A. Radford, L. Metz, S. Chintala, Unsupervised representation learning with deep convolutional generative adversarial networks, arXiv preprint arXiv:1511.06434 (2015).

[31] T. Salimans, I. Goodfellow, W. Zaremba, V. Cheung, A. Radford, X. Chen, X. Chen, Improved techniques for training gans, in: D. Lee, M. Sugiyama, U. Luxburg, I. Guyon, R. Garnett (Eds.), Advances in Neural Information Processing Systems, Vol. 29, Curran Associates, Inc., 2016.

[32] I. J. Goodfellow, On distinguishability criteria for estimating generative models, arXiv preprint arXiv:1412.6515 (2014).

[33] A. Lourenço, H. Silva, C. Carreiras, et al., Outlier detection in non-intrusive ecg biometric system, in: International Conference Image Analysis and Recognition, Springer, 2013, pp. 43–52.

[34] L.-K. Shark, C. Yu, Denoising by optimal fuzzy thresholding in wavelet domain, Electronics Letters 36 (6) (2000) 581–582.

[35] G. Garcia, G. Moreira, D. Menotti, E. Luz, Inter-patient ecg heartbeat classification with temporal vcg optimized by pso, Scientific reports 7 (1) (2017) 1–11.

[36] S. Akcay, A. Atapour-Abarghouei, T. P. Breckon, Ganomaly: Semi-supervised anomaly detection via adversarial training, in: Asian conference on computer vision, Springer, 2018, pp. 622–637.

[37] C. Chen, Z. Hua, R. Zhang, G. Liu, W. Wen, Automated arrhythmia classification based on a combination network of cnn and lstm, Biomedical Signal Processing and Control 57 (2020) 101819.