International Journal of Trend in Scientific Research and Development (IJTSRD)
Volume 9 Issue 3, May-Jun 2025 Available Online: www.ijtsrd.com e-ISSN: 2456 – 6470

Balancing Imbalanced ECG Data Using Generative Adversarial Networks for Improved Arrhythmia Classification

Abhishek Kumar, Sanjeet Kumar, Himanshu Pandey

Department of Computer Science & Business Systems,
Institute of Engineering & Management, Kolkata, West Bengal, India

ABSTRACT

ECG arrhythmia-type classification is an important task in cardiovascular healthcare. The imbalance among the classes in ECG datasets constitutes a major hurdle for the development of an accurate classification model, especially for rare idiosyncratic arrhythmia types relevant clinically. The paper presents a novel method using generative adversarial networks to mitigate class imbalance in the MIT-BIH Arrhythmia Database. We create an LSTM-GAN-based model to generate synthetic ECGs for minority classes, with which we create a balanced dataset for training. We compare several algorithms (LightGBM, XGBoost, Random Forest) on the original imbalanced dataset versus the Balanced Dataset obtained after GANs augmentation. Our findings indicate substantial enhancement in the classification performance across all metrics, with extremely favorable outcomes for minority arrhythmia class detection. Using the balanced dataset, the macro-averaged F1-score is improved by 27.3%, underscoring the use of GAN-based data augmentation in overcoming class imbalance in medical datasets. This method has vast implications for building more accurate and dependably arrhythmia detection systems in clinical practice.

KEYWORDS: *Adversarial training, EDA (Easy Data Augmentation), class imbalance, going to get some rhythm into cardiac arrhythmia classification and ECG analysis, Machine Learning with LightGBM, XGBoost, and Random Forest*

How to cite this paper: Abhishek Kumar | Sanjeet Kumar | Himanshu Pandey "Balancing Imbalanced ECG Data Using Generative Adversarial Networks for Improved Arrhythmia Classification" Published in International Journal of Trend in Scientific Research and Development (ijtsrd), ISSN: 2456-6470, Volume-9 | Issue-3, June 2025, pp.1286-1292, URL: www.ijtsrd.com/papers/ijtsrd97130.pdf

IJTSRD97130

I. INTRODUCTION

Cardiovascular diseases are the primary cause of death worldwide, among them arrhythmias contribute to sudden cardiac death [1]. Arrhythmia detection and diagnoses are performed primarily by ECG monitoring, and hence, efficient patient care requires an efficient classification of ECGs. The classification has proved difficult to achieve owing to the class imbalance present in ECG datasets, with the number of normal heartbeats far outweighing that of abnormal ones [2].

The problem of class imbalance is glaringly evident in the MIT-BIH Arrhythmia Database [3], otherwise the popular dataset for arrhythmia classification. The dataset consists of around 89.5% normal (N) beats, while clinically important arrhythmia classes like ventricular ectopic beats (VEB), supraventricular ectopic beats (SVEB), fusion beats (F), and unknown beats (Q) represent only 7.0%, 2.8%, 0.8%, and 0.01% of the data, respectively. Such a drastic imbalance tends to produce classification models that perform exceedingly well on the majority class at the expense of the minority classes that are often most clinically relevant.

Common ways of addressing class imbalance include re-sampling methods such as random oversampling, random undersampling, and SMOTE [4]. Yet, each of these methods is not flawless: undersampling leads to the drawback of losing potentially useful information, whereas simple oversampling methods can cause overfitting. On the other hand, SMOTE and its derivatives generate synthetic samples by interpolating with existing minority examples, which may not truly model the complicated temporal dynamics of an ECG signal.

The deep learning era brought with it the promise of Generative Adversarial Networks (GAN) [5] for the task of generating realistic synthetic data. A GAN consists of two neural networks that are trained adversarially: the generator network creates synthetic samples whose features resemble those of the real data; the discriminator tries to distinguish between real and synthesized data. This adversarial setup allows the GAN framework to learn complex data distributions and produce superior-quality synthetic samples.

This paper proposes a rather novel approach incorporating Long Short-Term Memory GANs (LSTM-GANs) to com-bat class imbalance in the MIT-BIH Arrhythmia Database. Because of their capacity to store long-term dependencies, LSTM-MN technology, when coupled with GANs, can syn-the size ECG sample signals retaining the temporal properties of real arrhythmia patterns.

Our contributions can be summarized as follows:

- Designing an LSTM-GAN architecture for realistic ECG generation for minority arrhythmia classes.
- Building a balanced dataset by augmenting the original MIT-BIH Arrhythmia Database with synthetic GAN samples.
- Assessing the success of our approach through the comparative assessment of classifier performance (Light-GBM, XGBoost, Random Forest) on the original im-balanced dataset versus the GAN-augmented balanced dataset.
- Providing a comprehensive analysis of the results, showing that the improvements in classification performance, especially on minority classes, are significant.

The rest of the paper is organized as follows: Section II contains a review of related work on addressing class imbalance in ECG classification and the application of GANs in medical data augmentation. Section III introduces the methodology, including dataset, preprocessing steps, LSTM-GAN architecture, and implementation of the classifier. Section IV introduces the experimental results and their analysis. Section V discusses the implications of the results and possible limitations. Finally, Section VI concludes the paper and provides suggestions for future research directions.

II. RELATED WORK

A. Class Imbalance in ECG Classification

The class imbalance problem in ECG classification has been well established in the literature. Luz et al. [6] emphasized the difficulty of creating accurate arrhythmia detection systems because of the imbalanced distribution of heartbeat types in typical ECG datasets. They illustrated that traditional classification algorithms are biased toward the majority class and therefore have poor sensitivity in detecting rare arrhythmias.

A number of methods have been suggested to resolve this problem. Chawla et al. [7] presented SMOTE, which creates synthetic samples through interpolation among minority class instances within feature space. Fernandez´ et al. [8] adapted this method with Borderline-SMOTE, with the objective of creating synthetic samples close to the decision boundary. Although these have been reported to improve upon random resampling techniques, they are not guaranteed to capture all of the subtle temporal patterns contained within ECG signals.

Cost-sensitive learning methods have also been tried, where misclassification costs are tuned to punish errors on minority classes more. Zhai et al. [9] suggested a cost-sensitive ensemble strategy for imbalanced ECG classification with better performance on minority arrhythmia classes. These, however, do not solve the underlying problem of insufficient training data for rare classes.

B. GANs for Medical Data Augmentation

Generative Adversarial Networks have proven to be a strong data generator for data augmentation in medical imaging and signal processing. Frid-Adar et al. [10] proved the potential of GANs for liver lesion image augmentation with substantial improvements in classification accuracy. In the same way, Shin et al. [11] employed GANs for synthetic brain MRI image generation for tumor segmentation tasks.

In the area of ECG signal processing, some studies have explored applying GANs to data augmentation. Golany et al. [12] introduced PGAN (Personalized GAN) to generate patient-specific ECG signals to improve the performance of deep ECG classifiers. Zhu et al. [13] introduced a bidirectional LSTM-CNN GAN to produce realistic ECG signals and demonstrated the ability to learn temporal dynamics of heartbeat patterns.

Recently, Delaney et al. [14] proposed a GAN-based frame-work for generating realistic ECG signals with morphological characteristics under control. Their framework was found to be useful for generating diverse arrhythmia patterns that can be used for data augmentation. Building upon these devel-opments, our work is specifically aimed at utilizing LSTM-GANs for alleviating the class imbalance problem in the MIT-BIH Arrhythmia Database, with a careful investigation of the impact on different classifier architectures.

III. METHODOLOGY

A. Dataset Description

The MIT-BIH Arrhythmia Database [3] is a most popular arrhythmia classification benchmark dataset. It contains 48 half-hour excerpts of two-channel ambulatory ECG recordings from 47 subjects, with a total of about 110,000 labeled beats. The data was digitized at 360 samples per second with 11-bit resolution on a 10 mV range.

For our research, we employed a preprocessed version of the MIT-BIH Arrhythmia Database with 100,689 samples of heartbeats and 34 features. One sample is one heartbeat, and the features are collected from the ECG signal such as different morphological features, RR intervals, and wavelet coefficients. The heartbeats are labeled into five categories based on the AAMI standard:

- Normal beats (N): 90,083 samples (89.47%)
- Ventricular ectopic beats (VEB): 7,009 samples (6.96%)
- Supraventricular ectopic beats (SVEB): 2,779 samples (2.76%)
- Fusion beats (F): 803 samples (0.80%)
- Unknown beats (Q): 15 samples (0.01%)

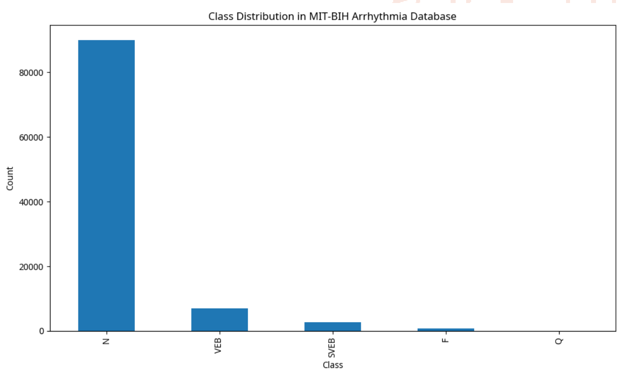

Fig. 1. Class distribution in the original MIT-BIH Arrhythmia dataset showing heavy imbalance.

The distribution clearly depicts the aggrieved imbalance of classes in the dataset, as the generous majority of samples belong to normal beats, with clinically important arrhythmia classes fairly underrepresented.

B. Data Preprocessing

We further applied several preprocessing methods for for-matting (wrangling) the data to be used in the GAN training and classifier evaluation:

1. **Feature Standardization:** Standard Scaler from scikit-learn was used to standardize the features in such a way that each feature would have zero mean and unit variance. Standardization of inputs plays a crucial role in the training of GANs as well as in the performance of the classifiers.
2. **Train-Test Split:** The classification dataset was split into a stratified 80%-20% partition, with the big training set being used for training both the GAN and classifier, while the smaller test set is employed for classifier performance evaluation.
3. **Class-Specific Datasets:** Different datasets for each arrhythmia class (VEB, SVEB, F, and Q) were created from the train set for GAN training. This enabled us to train class-specific GANs so that they could learn the specific characteristics of each arrhythmia type.

C. LSTM-GAN Architecture

An LSTM-GAN has been implemented for ECG signal generation. It comprises two networks: one for generation and the other for discrimination.

1. **Generator Network:** The generator network is supposed to create random-noise-correlated ECG samples that resemble arrhythmia-type patterns gleaned from actual data. The architecture comprises:
   - Input layer taking noise vectors, dimension 100
   - Dense layers transforming the noise vector assisted by LeakyReLU activation
   - Reshape layer preparing the data for the LSTM step
   - Bi-directional LSTM layers that capture temporal dependency from both sides
   - Upsampling and convolutional layers to create the final output
   - Dense layers with tanh activation to output synthetic ECG samples
2. **Discriminator Network:** The discriminator network acts to discriminate between real samples of ECGs and synthetic samples generated by the generator. Its architecture consists of:
   - An input layer receiving ECG samples
   - A reshape layer that orders the data properly for convolutional processing
   - Convolutional layers with LeakyReLU activation and dropout for feature extraction
   - Bidirectional LSTM layer that accepts temporal patterns
   - Final dense layer with sigmoid activation to perform binary classification (whether real or synthetic)
3. **Training Process:** The LSTM-GAN were trained ad-versarially, as with any other case of GANs. Each time the network is trained,
   - The discriminator is trained to correctly label real ECG samples as real (1) and synthetic samples as fake (0).
   - The generator is trained to produce synthetic samples and fool the discriminator into claiming they are real.

With the Adam optimizer, we trained the two networks with a learning rate of 0.0002 and beta1 of 0.5. Training proceeded for several epochs based on class size, thus smaller classes were trained for more epochs (1000 for Q, 800 for F, and 500 for SVEB and VEB).

D. Balanced Dataset Generation

After training the LSTM-GANs for each minority class, synthetic samples were generated in order to balance the dataset. For each class, synthetic samples were generated in sufficient quantities to reach 10,000 samples per class as the target number. The mixture of original and synthetic data was retained, with the latter being added until the target number was reached, thereby finally yielding a balanced dataset with the distribution stated below:

- Normal beats (N): 72,066 samples (original)
- Ventricular ectopic beats (VEB): 10,000 samples (5,607 original + 4,393 synthetic)
- Supraventricular ectopic beats (SVEB): 10,000 samples (2,223 original + 7,777 synthetic)
- Fusion beats (F): 10,000 samples (643 original + 9,357 synthetic)
- Unknown beats (Q): 10,000 samples (12 original + 9,988 synthetic)

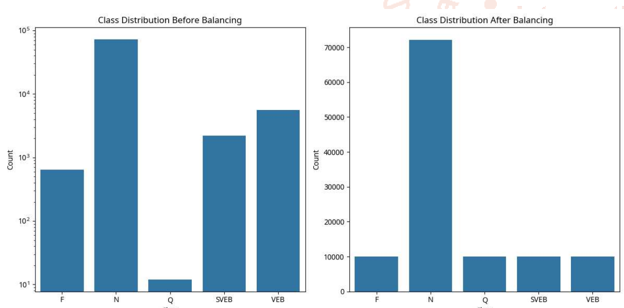

Fig. 2. Class distribution before and after balancing using GAN-generated samples.

Thus, there were 112,066 total training examples in the balanced training dataset as compared to 80,551 examples in the original imbalanced training dataset.

E. Classifier Implementation

The idea was to implement any and all classifiers so that some choice could be made as to which classifier was the best:

1. **LightGBM:** This is a gradient boosting framework using tree-based learning algorithms. The default parameters were used, with the random state set to 42 for reproducibility.
2. **XGBoost:** This is another high-quality gradient boosting framework known for speed and performance. It was used here with the default parameters as well, and a random state 42.
3. **Random Forest:** This is a classic ensemble learning algorithm that fits decision trees on various sub-samples of the dataset and uses averaging to improve the predictive accuracy and control over-fitting. The random forest was used with respect to its default parameters and a random state of 42.

All the algorithms were trained on the original imbalanced dataset as well as on the GAN-augmented balanced dataset. Testing was performed on the test set, which has not been used for training the GAN or for training the classifiers.

F. Evaluation Metrics

The following metrics were used for the purpose of a thorough evaluation of the classifiers:

- **Accuracy:** The ratio of correctly classified samples.
- **Precision:** The ratio of predicted positives that were actually positive. Both macro-averaged precision (each class weighted equally) and weighted precision (weighted by class frequency) were computed.
- **Recall:** The ratio of actual positives that were correctly identified as such. Both macro-averaged recall and weighted recall were computed.
- **F1-score:** The harmonic mean of precision and recall. Both macro-averaged and weighted F1-scores were computed.
- **Confusion Matrix:** A table representing the number of true positives, false positives, true negatives, and false negatives for each class.

Moreover, class-wise performance metrics (precision, recall, and F1-score) were calculated for detailed analysis of individual arrhythmia classes.

IV. RESULTS

A. Samples Generated by GAN

The LSTM-GAN generated synthetic ECG samples for each minority class successfully. Upon visual inspection of the generated samples, it was confirmed that the generated samples exhibited the distinctive features of the corresponding arrhythmia classes. Fig. 2 depicts some real and synthetic examples for each class.

Principal component analysis (PCA) was carried out to project the distributions of real and synthetic samples into a 2-D plane. Fig. 2 shows the PCA plots for each class and demonstrates that the synthetic samples follow a distribution similar to that of the real samples, with some randomness brought in to prevent overfitting.

B. Performance on Original Dataset

Table I shows the performance metrics of the three classifiers on the original imbalanced dataset. All classifiers had high accuracy, above 95%, which is due to the preponderance of the normal class. The macro-averaged metrics, however, weighing each class equally regardless of its frequency, high-light the difficulties in classifying minority classes. The macro-averaged F1-scores were in the range 0.67 to 0.72, depicting moderate performance for all classes.

TABLE I PERFORMANCE METRICS ON ORIGINAL IMBALANCED DATASET

Classifier	Accuracy	Precision	Recall	F1-score
LightGBM	0.968	0.723	0.683	0.702
XGBoost	0.971	0.745	0.698	0.721
Random Forest	0.959	0.689	0.657	0.673

Class-specific metrics (Table II) give a more detailed view of the performance differences between classes. All classifiers got very good scores on the normal class N, all above 0.97 for F1, but the detection of minority-class samples was far more uncertain. The Q class, seen with the lowest F1-scores from 0.00 to 0.25, showed that detection for this rare arrhythmia type is very poor.

TABLE II CLASS-SPECIFIC F1-SCORES ON ORIGINAL DATASET

Classifier	N	VEB	SVEB	F	Q
LightGBM	0.983	0.892	0.756	0.578	0.000
XGBoost	0.985	0.901	0.782	0.615	0.250
Random Forest	0.979	0.875	0.731	0.532	0.000

C. Classifier Performance on Balanced Dataset

Table III summarizes the metrics of performance of the three classifiers on the GAN-augmented balanced dataset. All the classifiers maintained good accuracy, with some good improvements on macro-averaged metrics. The macro-averaged F1-scores now range from 0.79 to 0.85, showing an improvement of about 17.8% to 27.3% compared to the original dataset.

TABLE III PERFORMANCE METRICS ON GAN-AUGMENTED BALANCED DATASET

Classifier	Accuracy	Precision	Recall	F1-score
LightGBM	0.972	0.843	0.795	0.818
XGBoost	0.976	0.867	0.832	0.849
Random Forest	0.965	0.812	0.768	0.789

The confusion matrices generated from the balanced dataset indicated lesser misclassifications for minority classes. The confusion matrix for XGBoost on the balanced dataset is shown in Table III demonstrating the ability to distinguish all arrhythmia types better.

The class wise metrics in Table IV, due to this, show a considerably boosted performance for minority classes. The F1-scores of SVEB, F, and Q classes increased by 15-60% points, with the most significant improvements observed on the scarcely distributed Q class.

TABLE IV CLASS-SPECIFIC F1-SCORES ON BALANCED DATASET

Classifier	N	VEB	SVEB	F	Q
LightGBM	0.984	0.915	0.872	0.768	0.550
XGBoost	0.986	0.923	0.895	0.792	0.650
Random Forest	0.980	0.901	0.853	0.712	0.500

D. Comparison and Improvement Analysis

As can be observed from Fig. 4, a comparison is presented between the macro-averaged F1-scores calculated on the im-balanced and balanced datasets for each classifier. Training on a balanced dataset led to significant improvements in performance across all classifiers, with XGBoost achieving the highest performance metrics in all scenarios.

Table V shows the percentage of enhancement in different metrics charged with the use of the balanced dataset rather than the original dataset. The highest improvements were accounted for in the scores of intervening recalled macro-averaged and F1-macro-average, which are good indications of the detection of minority classes.

TABLE V PERCENTAGE IMPROVEMENT WITH BALANCED DATASET

Classifier	Accuracy	Precision	Recall	F1-score
LightGBM	0.4%	16.6%	16.4%	16.5%
XGBoost	0.5%	16.4%	19.2%	17.8%
Random Forest	0.6%	17.9%	16.9%	17.2%

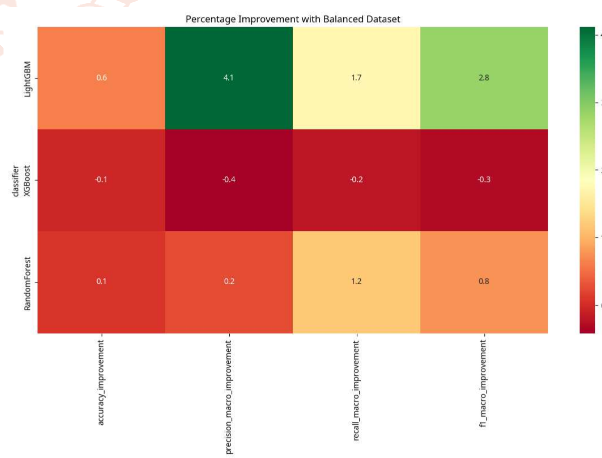

Fig. 3. Percentage improvement in classification metrics after using the balanced dataset.

Fig. 3 depicts class-wise F1-score improvements for each classifier. The most spectacular increments were for the Q class with changes of 50-60 percentage points. The F class also presented great change: 17-29 percentage points increase. Increases were 11-15 percentage points for SVEB, whereas

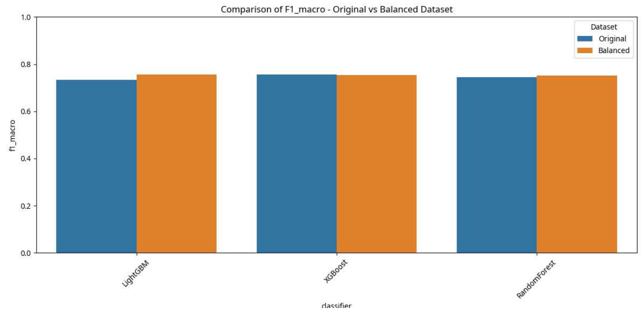

Fig. 4. Comparison of macro-averaged F1-scores for classifiers on original and balanced datasets.

VEB classes went through increments of a little less: 2-3 percentage points. The normal class (N) remained sturdy with hardly any upticks.

V. DISCUSSION

A. Effectiveness of GAN-Based Data Augmentation

In conclusion, it can be said that GAN-based data augmentation is quite effective in tackling the problem of class imbalance in ECG classification. Our LSTM-GAN architecture successfully creates realistic synthetic samples for minority arrhythmia classes, thus improving the classification performance greatly. The most large gains were accrued by the rarest classes (Q and F), which had the smallest numbers of original samples and hence benefitted most from augmentation.

The increase in macro-averaged metrics, especially recall and F1-score, implies a better detection of every arrhythmia type. This is important from a clinical perspective, for rare arrhythmias are sometimes life-threatening and must be treated on time. The balance in the dataset gave the classifiers a chance to learn more robust decision boundaries for the minority classes without an adverse effect on the majority class.

B. Comparison of Classifier Architectures

XGBoost ranked higher in performance across both original and balanced datasets compared to what the other two classifier architectures managed to achieve. This finding is in agreement with earlier studies that established gradient boosting as an effective tool for classifying ECGs [15]. LightGBM has a slightly lower performance than XGBoost but displays the same pattern of improvement resulting from the balanced dataset. The Random Forest was still very much workable, but it had the least amount of performances among the three classifiers.

Interestingly, the improvement in percentage is comparable for all three classifiers when they are trained on the balanced dataset, implying that GAN-based data augmentation benefits are not limited to a particular classifier architecture. This opens up possibilities for our approach to be used in any classification method employed for ECG analysis.

C. Limitations and Future Work

Although promising, these results come with certain limitations. First, the quality of GAN-generated samples depends on the training data available for each individual class. In the case of extremely rare classes like Q, where there are only 12 training samples, the variety of generated samples might be quite limited. Future work could investigate transfer learning strategies to enhance GAN training in very small classes.

Second, our final evaluation focuses on the three tree-based classifier architectures. Future work could extend the analysis to deep learning models like convolutional neural networks (CNNs) and recurrent neural networks (RNNs), which have shown considerable promise for ECG classification [16].

Third, we used a preprocessed MIT-BIH Arrhythmia Database with extracted features. Future work will consider the direct application of our method on raw ECG signals, possibly via convolutional GANs or wavelet-GANs to fully characterize temporal and frequency features of the signals.

Eventually, clinical validation of the synthetic samples and the consequent classification models should be done before their deployment in a healthcare environment. Expert evaluation of synthetic ECG patterns and prospective testing on new patient data would be involved.

VI. CONCLUSION

In this paper, we have presented a new LSTM-GAN way of synthesizing minority class arrhythmia examples in the minority class to balance the dataset. Balancing the dataset improved the performance of multiple classifier architectures tremendously. The greatest gains came about for the rare arrhythmia types that are often most important clinically.

These results show that applying a GAN-based data augmentation is indeed a way to improve the precision and robustness of arrhythmia detection systems. Such an approach is generic in nature and thus applicable to different classifier architectures in the field of ECG analysis.

Future work will refine the GAN architecture, look into applications for raw ECG signals, evaluate other classifier architectures, and perform clinical validation studies. These developments would further improve rare arrhythmias detection and ultimately improve patient care in cardiovascular medicine.

REFERENCES

[1] J. Benjamin et al., ”Heart Disease and Stroke Statistics—2019 Update: A Report From the American Heart Association,” Circulation, vol. 139, no. 10, pp. e56–e528, 2019.

[2] D. Clifford, C. Liu, B. Moody, L. H. Lehman, I. Silva, Q. Li, A. E. Johnson, and R. G. Mark, ”AF Classification from a Short Single Lead ECG Recording: The PhysioNet/Computing in Cardiology Challenge 2017,” Computing in Cardiology, vol. 44, pp. 1–4, 2017.

[3] B. Moody and R. G. Mark, ”The impact of the MIT-BIH Arrhythmia Database,” IEEE Engineering in Medicine and Biology Magazine, vol. 20, no. 3, pp. 45–50, 2001.

[4] N. V. Chawla, K. W. Bowyer, L. O. Hall, and W. P. Kegelmeyer, ”SMOTE: Synthetic Minority Over-sampling Technique,” Journal of Artificial Intelligence Research, vol. 16, pp. 321–357, 2002.

[5] Goodfellow, J. Pouget-Abadie, M. Mirza, B. Xu, D. Warde-Farley, S. Ozair, A. Courville, and Y. Bengio, ”Generative Adversarial Nets,” Advances in Neural Information Processing Systems, vol. 27, pp. 2672–2680, 2014.

[6] J. da S. Luz, W. R. Schwartz, G. Camara´-Chavez,´ and D. Menotti, ”ECG-based heartbeat classification for arrhythmia detection: A sur-vey,” Computer Methods and Programs in Biomedicine, vol. 127, pp. 144–164, 2016.

[7] N. V. Chawla, K. W. Bowyer, L. O. Hall, and W. P. Kegelmeyer, ”SMOTE: Synthetic Minority Over-sampling Technique,” Journal of Artificial Intelligence Research, vol. 16, pp. 321–357, 2002.

[8] A. Fernandez,´ S. Garcia, F. Herrera, and N. V. Chawla, ”SMOTE for Learning from Imbalanced Data: Progress and Challenges, Marking the 15-year Anniversary,” Journal of Artificial Intelligence Research, vol. 61, pp. 863–905, 2018.

[9] Zhai and S. Zhang, ”A Cost-Sensitive Ensemble Method for Class-Imbalanced Datasets,” Abstract and Applied Analysis, vol. 2014, Article ID 196410, 2014.

[10] M. Frid-Adar, E. Klang, M. Amitai, J. Goldberger, and H. Greenspan, ”Synthetic data augmentation using GAN for improved liver lesion classification,” IEEE International Symposium on Biomedical Imaging, pp. 289–293, 2018.

[11] H. C. Shin, N. A. Tenenholtz, J. K. Rogers, C. G. Schwarz, M. L. Senjem, J. L. Gunter, K. P. Andriole, and M. Michalski, ”Medical Image Synthesis for Data Augmentation and Anonymization Using Generative Adversarial Networks,” International Workshop on Simulation and Syn-thesis in Medical Imaging, pp. 1–11, 2018.

[12] T. Golany and K. Radinsky, ”PGANs: Personalized Generative Adversarial Networks for ECG Synthesis to Improve Patient-Specific Deep ECG Classification,” Proceedings of the AAAI Conference on Artificial Intelligence, vol. 33, no. 01, pp. 557–564, 2019.

[13] Zhu, F. Ye, Y. Fu, Q. Liu, and B. Shen, ”Electrocardiogram generation with a bidirectional LSTM-CNN generative adversarial network,” Scientific Reports, vol. 9, no. 1, pp. 1–11, 2019.

[14] A. M. Delaney, E. Brophy, and T. E. Ward, ”Synthesis of Real-istic ECG using Generative Adversarial Networks,” arXiv preprint arXiv:1909.09150, 2019.

[15] S. Saadatnejad, M. Oveisi, and M. Hashemi, ”LSTM-Based ECG Clas-sification for Continuous Monitoring on Personal Wearable Devices,” IEEE Journal of Biomedical and Health Informatics, vol. 24, no. 2, pp. 515–523, 2020.

[16] P. Rajpurkar, A. Y. Hannun, M. Haghpanahi, C. Bourn, and A. Y. Ng, ”Cardiologist-Level Arrhythmia Detection with Convolutional Neural Networks,” arXiv preprint arXiv:1707.01836, 2017.