Deep Convolutional Neural Network-Based ECG Analyzer

Author
Mane Hakobyan
mane_hakobyan@edu.aua.am

Supervisor
Arthur Ghulyan
arthurghulian@gmail.com

May 3, 2023

Abstract

This article introduces a deep-learning-based instrument designed to identify cardiac arrhythmias. Electrocardiography (ECG) serves as a prevalent, non-invasive diagnostic technique employed to evaluate the electrical function of the heart. Conducting an ECG test entails placing electrodes on a patient's chest, arms, and legs, thereby capturing the electrical impulses produced by the heart These impulses are subsequently amplified, filtered, and depicted as a graphical waveform, illustrating the heart's electrical patterns.

Our primary goal is to incorporate this tool within a medical application as an assisting tool that will help physicians in Patient's continues care for chronic diseases' prevention and management. Utilizing such advanced instruments can facilitate the identification of heart irregularities even in the absence of cardiologists.

The methodology implemented in this research project involves an existing approach operating on unfiltered ECG signals. The chosen approach demonstrates exceptional performance compared to other existing methods. This study thereby contributes significantly to the fields of heart disease diagnosis and management, ECG signal analysis, and deep learning applications in detecting arrhythmias and other cardiovascular diseases.
Keywords: *heart diseases, ECG, ECG signal analysis, cardiovascular diseases, deep learning, arrhythmia detection.*

1 Introduction

Cardiovascular diseases is the leading cause of mortality worldwide. In 2019 alone, approximately 17.9 million individuals died from such diseases, accounting for 32% of all global fatalities (Centers for Disease Control and Prevention (2020)). Early detection and diagnosis of cardiac pathologies can significantly reduce the death rate associated with cardiovascular disorders. Consequently, the emphasis is on making the detection of cardiovascular anomalies as widely accessible as possible. Electrocardiography (ECG) is a non-invasive diagnostic technique devoid of risks, pain, and harm to the human body, making it an ideal choice for this purpose.

In countries such as Armenia, the Ministry of Health mandates the presence of an ECG machine in every medical clinic and hospital. Furthermore, most ambulances are equipped with ECG machines. Given these circumstances, it can be reasonably argued that analyzing ECG signals is the most accessible approach for detecting cardiovascular diseases, particularly arrhythmias. This accessibility underscores the importance of developing advanced tools and methods to efficiently analyze and interpret ECG data to improve early diagnosis and treatment of heart-related conditions.

While our tool can be used to identify arrhythmias through the utilization of resting (standard) ECGs, our primary objective is contributing to the development of continuous care, which is done within the patient's residence. As our goal is the monitoring heart rate and rhythm, there are two main ECG types that we find best to use for our tool. One of them is using a 24-hour two or three lead Holter monitor (uses II, V1 and/or V5 leads). This device, based on the Galvanometer's principle, records electrocardiographic signals from an individual engaged in daily activities.

Despite the Holter monitor's efficiancy in arrhythmia detection, an alternative ECG type offers greater accessibility to the general population, ease of wear, minimal discomfort, and superior outcomes compared to both standard ECG and Holter monitors. The adhesive patch monitor, a single lead ECG variant (commonly using lead II), is compact, portable, and records for a continuous 14-day period, increasing the likelihood of identifying infrequent or otherwise undetectable rhythm alterations. A clinical research Barrett et al. (2014) study revealed that, in comparison to the 61 arrhythmia events detected by the Holter monitor in the same study group, the adhesive patch monitor identified 96 arrhythmia events.

Heart has characteristic electrical activity patterns: The initial peak, known as the P wave, represents the dispersion of the electrical impulse (excitation) across the heart's two atria. As the atria contract, blood is pumped into the ventricles before the atria subsequently relax. The electrical impulse then reaches the ventricles, which is observed in the Q, R, and S waves of the ECG, collectively referred to as the QRS complex. The ventricles contract, and then the T wave indicates the cessation of the electrical impulse's propagation, followed by the relaxation of the ventricles.

Various heart diseases and irregular heartbeats can be identified through ECG analysis. Examining the appearance and development of these irregularities aids in determining their underlying causes. Remarkably, the electrical activity of the heart can be measured on the skin's surface, even at considerable distances from the heart. The standard "12-lead ECG" employs a total of ten electrodes: six on the chest and one each on the lower arms and calves. These electrodes are connected via cables to an ECG machine, which translates the received signals into an ECG graph and stores it for further examination and interpretation.

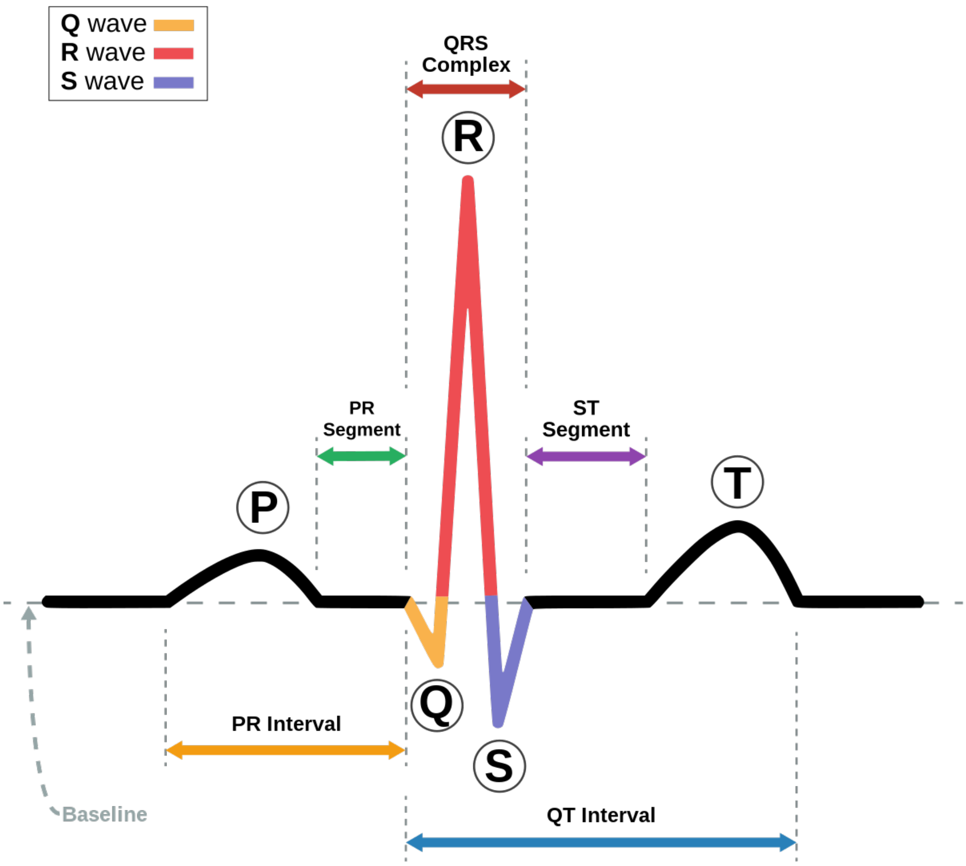

Figure 1: ECG Pattern.

The precise identification of the QRS complex is of paramount importance for the diagnosis and surveillance of numerous cardiac disorders. Electrocardiogram (ECG) signals, however, are frequently contaminated by noise, which can compromise the reliability of QRS detection. Various approaches have been devised to mitigate noise interference in ECG signals, encompassing digital filtering, wavelet transformation, and adaptive thresholding methodologies. A prominent example of such techniques is the Pan-Tompkins algorithm, which is extensively employed for QRS detection. This algorithm integrates both filtering and thresholding methods to identify the QRS complex accurately. It initially applies a band-pass filter to eliminate undesired frequency components, followed by a differentiation filter that accentuates the QRS complex. Subsequently, the signal undergoes squaring to amplify the QRS amplitude, and a moving average filter is employed to refine the signal's smoothness. The final stage incorporates a dynamic thresholding technique to detect the QRS complex effectively.

Once the QRS complexes have been accurately identified, a range of classification algorithms can be utilized to categorize distinct arrhythmias. These algorithms leverage features derived from the ECG signal, including heart rate variability, QRS complex amplitude and duration, and the existence of atypical waveforms. Widely-used classification

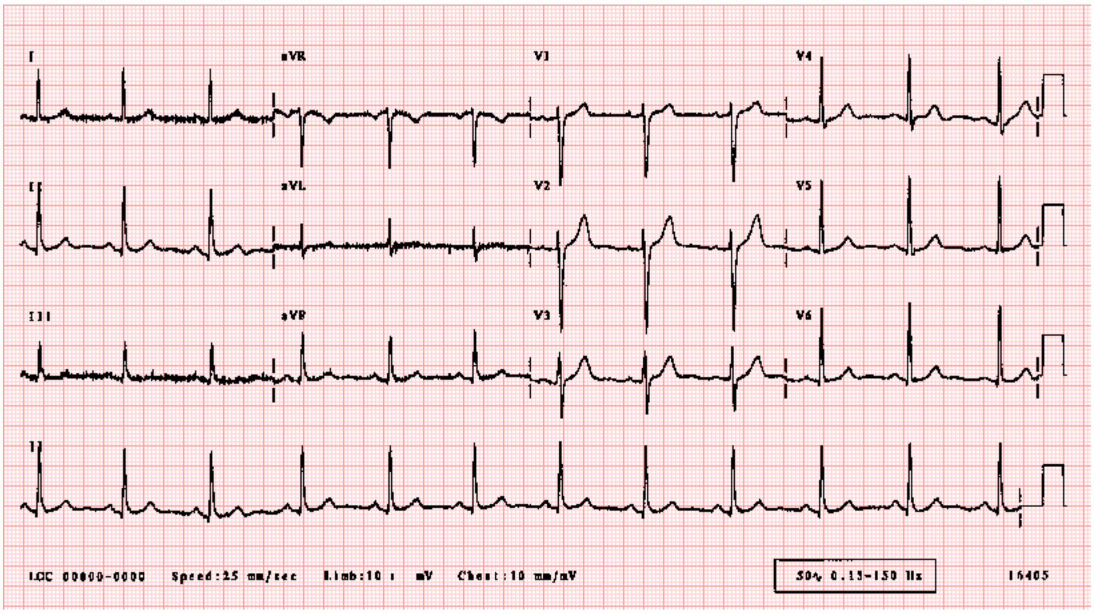

Figure 2: Normal Adult 12-lead ECG.

algorithms encompass support vector machines, artificial neural networks, decision trees, and logistic regression models.

Upon conducting a thorough review of numerous scholarly articles, we opted to base our Arrhythmia detection system on the research paper titled ”Arrhythmia Detection Using Deep Convolutional Neural Network with Long Duration ECG Signals.Yildirim, Pawiak, Tan, and Acharya (2018)” In order to achieve improved results, we implemented several modifications to the original model and tested on a different dataset for training purposes. Our proposed approach relies on a convolutional neural network (CNN) architecture specifically designed for the classification of brief ECG signal segments. The deep network structure, comprised of 16 layers, incorporates conventional CNN layers.

The input for this advanced network configuration consists of 3600 raw ECG signal samples, each with an extended duration. Notably, the classifier network’s output delivers predictions for signal classes without necessitating QRS detection and segmentation, which distinguishes it from traditional methodologies. To thoroughly assess the network’s performance, we utilized a database containing 1000 ECG fragments and conducted experimental studies involving 13-, 15-, and 17-class cases. In the article, the authors describe conducted experimental studies involving 13-, 15-, and 17-class cases. We conduct our experiment on 12 distinct classes.

2 Methods

2.1 Data

The MIT-BIH Arrhythmia Database George B. Moody (2010) comprises 48 electrocardiograms (ECG) records, each with a duration of slightly over 30 minutes. These records have been extracted from a larger set of more than 4000 long-term Holter recordings obtained by the Beth Israel Hospital Arrhythmia Laboratory between 1975 and 1979. The database is devided into two distinct groups. The first group (23 records) serves as a representative sample of waveforms and artifacts, while the second group (25 records) is specifically curated to encompass rare yet clinically significant phenomena.

The subjects involved in the study include 25 males with an age range of 32 to 89 years and 22 females with an age range of 23 to 89 years. In most cases, the ECG lead configuration consists of a modified limb lead II (MLII) as the upper signal and a modified lead V1 as the lower signal, with both electrode placements on the chest. However, certain records employed alternative lead configurations due to surgical dressings or other specific circumstances (Mainly they are V2 or V5 leads).

The ECG recordings in the MIT-BIH Arrhythmia Database were digitized at a rate of 360 samples per second per channel, featuring 11-bit resolution over a ten mV range. To ensure accuracy and reliability, a minimum of two independent cardiologists annotated each record. In cases of disagreement, the annotators reached a consensus to establish the computer-readable reference annotations for each beat. In total, the database contains approximately 110,000 annotationsE. (2000).

We have 12 different annotations Ribeiro et al. (2020): normal beat, atrial premature beat, atrial flutter, atrial fibrillation, pre-extitation, ventricular bigeminy, ventricular trigeminy, idioventricular rhythm, ventricular flutter, left bundle branch block beat, right bundle branch block beat and pacemaker rhythm.

2.2 Preprocessing and Model

The ECG signals from the MIT-BIH Arrhythmia Database were preprocessed by segmenting them into 10-second windows, detrending, normalizing amplitudes, and applying wavelet-based denoising. To capture time and frequency domain information, Empirical Mode Decomposition (EMD) is utilized for feature extraction, resulting in a set of intrinsic mode functions (IMFs). In order to achieve optimal performance in the classification task, various hyperparameters were selected and tuned for the 16-layer deep convolutional network. Key hyperparameters considered in this study include the learning rate, number of epochs, and batch size.

The learning rate was set to 0.0003, which determines the step size at each iteration during the optimization process. An appropriate learning rate ensures that the model converges to a solution in a reasonable amount of time without overshooting or getting stuck in local minima. The number of epochs was set to 100, which represents the number of complete passes through the entire training dataset.

Other hyperparameters include in_channels_, num_segments_in _record, segment_len, num_records, num_classes, and allow_label_leakage. These hyperparameters were chosen based on the characteristics of the dataset and the problem at hand, such as the number of input channels, the number of segments in a record, the length of each segment, the total number of records, and the number of output classes.

First, the dataset was loaded and stored as a Pandas DataFrame. The dataset was then filtered by the target label for each arrhythmia class, and the data were divided into training, validation, and testing subsets based on the desired ratios. For example, the normal beat ECG data was divided into 198 training samples, 44 validation samples, and 41 test samples. This process was repeated for each arrhythmia class, and the corresponding subsets are concatenated to form the final training, validation, and testing sets.

Once the dataset has been split, the custom class is used to create PyTorch Dataset objects for each subset. These Dataset objects are then passed to the DataLoader class, which is responsible for loading the data in batches during the training, validation, and testing phases. The DataLoader class takes the following parameters: dataset, batch_size, and shuffle. The batch_size was set according to the hyperparameter discussed earlier, and the shuffle parameter was set to True, ensuring that the data was randomly shuffled before each epoch.

A 16-layer deep convolutional network was developed for the classification of ECG signals according to cardiac arrhythmia. This deep network model facilitates automatic classification of input fragments through an end-to-end structure, eliminating the need for hand-crafted feature extraction or selection steps. The model's architecture comprises classical CNN layers, with a predominant 1D-CNN structure. 1D convolution layers process feature maps, which represent ECG fragments, using various weight sizes. The model's first layer applies 1D convolution with 128 weight vectors on the input ECG signals, followed by batch normalization for each batch. The 1D max pooling layer generates new feature maps by taking maximum values in a specified region on the feature maps obtained from previous layers, thus reducing feature map sizes and computational cost. Alternative methods, such as average values, may also be employed.

Subsequent layers involve repeated convolution processes using different weight sizes, followed by batch normalization and pooling. The deep network features a flattened layer at the 14th layer, transforming multidimensional input feature vectors into one-dimensional output data. The features obtained from the flattened layer are then fed to a dense-connected neural network layer with 512 units. The network's final layer consists of a softmax layer with units equal to the number of output classes, enabling the prediction of the input data's class membership. Dropout parameters are incorporated in some layers to prevent overfitting during the learning phase.

The model's layer numbers, types, and parameters were fine-tuned using a brute force technique to optimize performance on validation sets. The developed 16-layer model demonstrated the highest classification results for long-duration ECG signals.

2.3 Visualization

To visualize the waveforms, we use WFDB, which is a native Python waveform-database package created by the PhysioNet team. It allows one to read, write, and process the signals and annotations. You can see an example of visualizations in Figure 3.

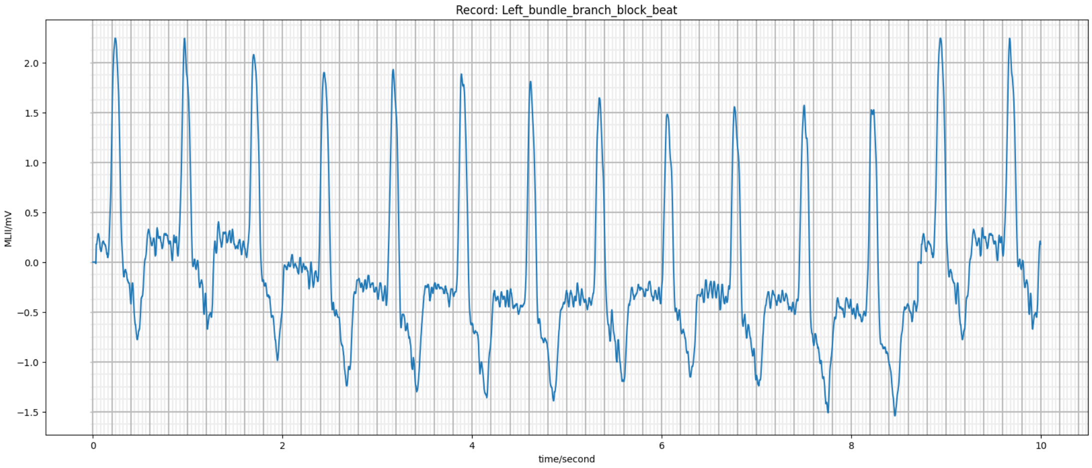

Figure 3: Left Bundle Block Beat

2.4 Tool

For the user interface of our tool, we have employed Streamlit, an open-source application framework designed for Python. Streamlit has been specifically developed to facilitate the creation of interactive applications that incorporate data scripts.

As shown in the Figure 4, the user interface provides an intuitive and accessible means for users to interact with the tool. Through this interface, users can select the patient's identification number and the desired lead for the ECG analysis and customize the scale of the visualization to suit their preferences. Upon submitting their preferences, users are provided with several informative outputs generated by the tool. First, the visualization of the waveform for the selected lead is displayed, allowing users to examine the ECG signal's characteristics in detail.

Next, the tool presents its prediction results, showcasing the three most probable outcomes, each accompanied by a percentage score. The sum of these percentage scores amounts to 100 percent, ensuring a comprehensive representation of the model's confidence in its predictions.

Finally, users are given access to a table containing an overall analysis of the ECG. This table consolidates valuable insights and offers a more in-depth understanding of the ECG's features, aiding users in interpreting the results and making informed decisions.

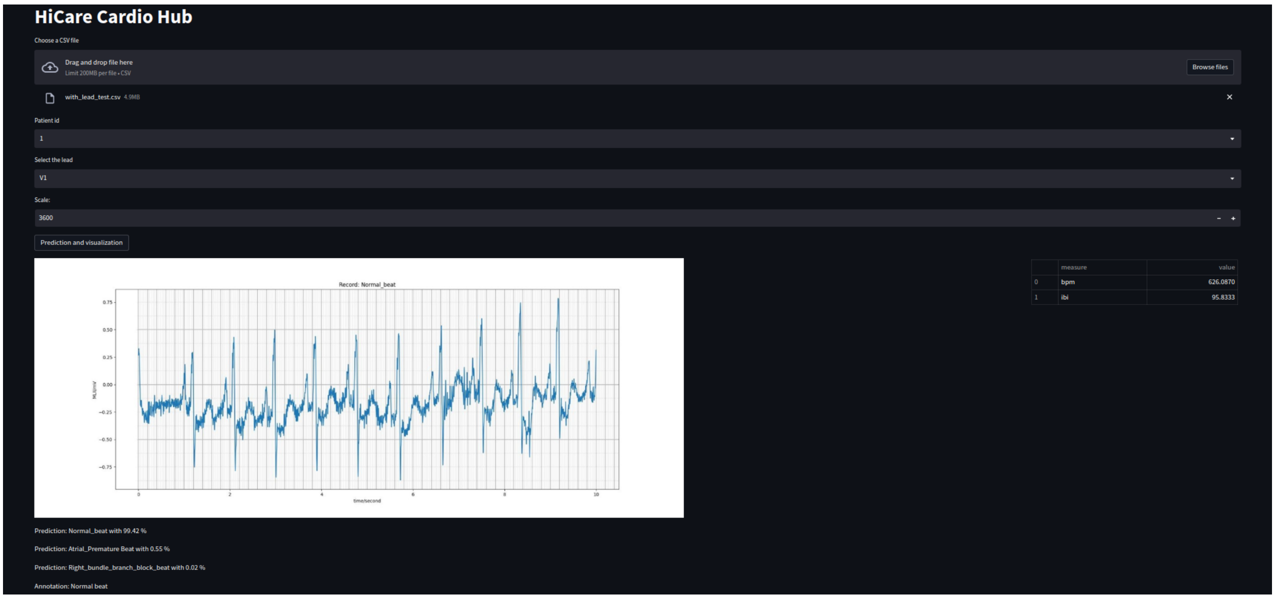

Figure 4: Preview of the Tool

3 Results

We developed python based tool that will be integrated into the healthcare system as an assistant for the doctors. It is automatically analyzing the patients ECG.

The model's performance is evaluated across multiple epochs during the training process. For a total of 10 epochs, with a learning rate of 0.0003, throughout the training process, we recorded the average training loss, training accuracy, test accuracy, validation loss, and test loss at each epoch. The results indicate that the model's performance improved over time. The training loss decreased from 0.1776 in epoch 1 to 0.0526 in epoch 10, while the training accuracy increased from 0.6 to 0.8. Similarly, the test loss dropped from 0.15522 in epoch 1 to 0.08340 in epoch 10, and the test accuracy increased from 0.3333 to 1.0.

The validation loss also showed a general downward trend, decreasing from 0.14903 in epoch 1 to 0.06784 in epoch 10. The learning rate remained constant at 0.0003 throughout the training process.

It is important to note that the test accuracy reached 1.0 in epochs 4, 5, 9, and 10, indicating that the model was able to correctly classify all arrhythmia cases in the test dataset during these epochs. However, there were fluctuations in the test accuracy across different epochs, which might be attributed to the limited size of the training and test datasets. Also, this might mean that we have an overfitting problem.

In summary, the results demonstrate the effectiveness of our deep CNN model in detecting arrhythmias using long-duration ECG signals. The model showed significant improvements in training and test performance over the course of the ten epochs.

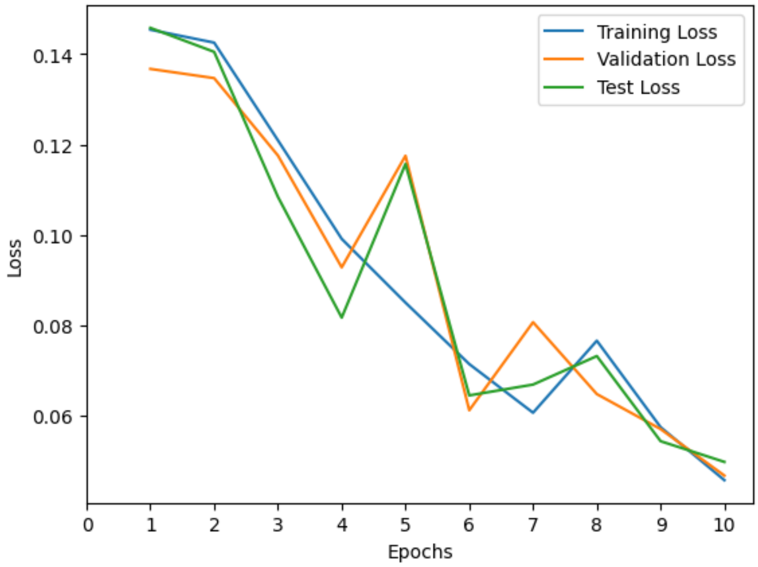

Figure 5: Train, Validation and Test Loss.

4 Discussion and Conclusions

In this study, we presented the use of a deep convolutional neural network (CNN) model for arrhythmia detection using long-duration ECG signals. The model demonstrated significant improvements in training and test performance over the course of several epochs, providing a strong foundation for the development of a reliable arrhythmia detection tool. However, there were concerns regarding the possibility of overfitting, which could impact the model's ability to generalize well on unseen data.

One limitation of our dataset for detecting arrhythmias on standard ECG is the availability of ECG signals from only a few leads (MLII, V1, V2, and V5), while clinical practice often relies on information from all 12 leads. In future work, incorporating data from all 12 leads could enhance the model's performance and provide a more comprehensive representation of the underlying cardiac activity, potentially improving arrhythmia detection accuracy for the standard ECG.

Additionally, our dataset suffers from class imbalance, where some arrhythmias have very small representations. This issue can hinder the model's ability to learn effectively from these underrepresented classes, leading to suboptimal performance in detecting rare arrhythmias. To address this limitation, acquiring a larger dataset with more balanced class representation for each arrhythmia type is recommended.

To mitigate the risk of misleading diagnoses and account for the nuances and responsibilities associated with the medical field, our tool presents the top three most probable diagnoses. This approach offers clinicians a wider perspective on the potential arrhythmia types, reducing the likelihood of misdiagnoses and providing better support for clinical decision-making.

In conclusion, our deep CNN model demonstrates promise as a tool for arrhythmia detection using long duration ECG signals. However, to achieve optimal performance and ensure the model's reliability in a clinical setting, future work should focus on addressing the limitations related to data quality and quantity, as well as implementing strategies to prevent overfitting. By addressing these challenges, our arrhythmia detection tool has the potential to significantly contribute to the early diagnosis and management of cardiac arrhythmias, ultimately leading to better patient outcomes and potentially saving lives.

References

Barrett, P. M., Komatireddy, R., Haaser, S., Topol, S., Sheard, J., Encinas, J., ... Topol, E. J. (2014). Comparison of 24-hour holter monitoring with 14-day novel adhesive patch electrocardiographic monitoring. *The American Journal of Medicine*, *127*(1), 95.e11-95.e17. Retrieved from https://www.sciencedirect.com/science/article/pii/S000293431300870X doi: https://doi.org/10.1016/j.amjmed.2013.10.003

Centers for Disease Control and Prevention. (2020). *Heart disease facts.* Retrieved from https://www.cdc.gov/heartdisease/facts.htm

E., G. A. A. L. G. L. . . S. H. (2000). Components of a new research resource for complex physiologic signals. Circulation [Online]. , e215-e220. doi: https://doi.org/10.13026/C2F305

George B. Moody. (2010). *Mit-bih arrhythmia database directory.* Retrieved from https://archive.physionet.org/physiobank/database/html/mitdbdir/intro.htm#leads

Ribeiro, A. H., Ribeiro, M. H., Paixão, G. M. M., Oliveira, D. M., Gomes, P. R., Canazart, J. A., ... Ribeiro, A. L. P. (2020). Automatic diagnosis of the 12-lead ECG using a deep neural network. *Nature Communications*, *11*(1), 1760. doi: https://doi.org/10.1038/s41467-020-15432-4

Yildirim, ., Pawiak, P., Tan, R. S., & Acharya, U. R. (2018, 09). Arrhythmia detection using deep convolutional neural network with long duration ecg signals. *Computers in Biology and Medicine*, *102*, 411-420. doi: 10.1016/j.compbiomed.2018.09.009