

An Effective Framework for Minimizing False Alarm Rates in Abnormal Driving Detection using Long Short-Term Memory and Electrocardiogram Data

Sundara Krishnan¹

Assistant Professor, Department of Computer Science Engineering, Alagappa Chettiar Government College of Engineering and Technology, Karaikudi, India¹

Abstract: The detection of anomalies in electrocardiograms (ECG) is a crucial technique for identifying irregular heartbeats, facilitating the early recognition of abnormal ECG readings prior to the diagnostic phase. Existing methods for ECG anomaly detection, which range from academic studies to commercial ECG devices, continue to experience a significant rate of false alarms. This issue arises from the inability of these methods to distinguish between ECG artifacts and genuine ECG signals, particularly when the artifacts closely resemble the actual signals in terms of shape and frequency. Consequently, this situation necessitates heightened vigilance among physicians and increases the risk of misinterpretation for those without specialized training. To address this challenge, the present study introduces a novel anomaly detection technique that demonstrates high robustness and accuracy in the presence of ECG artifacts, thereby effectively minimizing the false alarm rate. The design incorporates expert insights from cardiologists alongside motif discovery techniques. Furthermore, each phase of the algorithm aligns with cardiologists' interpretations. This method is applicable to both single-lead and multi-lead ECGs. The results of our experiments, conducted on real ECG datasets, have been analyzed and evaluated by cardiologists. As a result, ECG anomaly detection has gained significant traction among researchers and practitioners, being employed to identify periods of unusual ECG activity. The effectiveness of the anomaly detection method is directly correlated with the outcomes of cardiac disease identification and diagnosis. A network is trained using non-anomalous data and serves as a predictor over multiple time intervals. Determining whether the observed data is anomalous is a critical task that has been extensively explored in the literature. An auto-encoder is utilized to capture patterns, enabling the prediction of anomalies within the data. It is trained exclusively on normal pattern data and subsequently tested with anomalous data.

Keywords: Electrocardiogram, anomaly, auto-encoder, Long Short-Term Memory.

I. INTRODUCTION

The electrocardiogram (ECG or EKG) signal constitutes a time series data sequence that reflects the electrical impulses generated by the myocardium. This signal is captured through multiple electrodes placed on the skin's surface. Many healthcare professionals favour the ECG as a non-invasive method for identifying and diagnosing cardiac conditions. Two key features of an ECG signal include the collection of multiple recordings from various myocardial locations and its periodic waveform, which is synchronized with the cardiac cycle. A standard ECG comprises five morphological segments, namely the PQRST waveforms, which correspond to the electrical conduction throughout the entire cardiac cycle. Each cycle includes the processes of depolarization and repolarisation that occur from the atria to the ventricles [1]. The morphology of ECGs from different leads may vary based on the heart's vector, with each lead's morphology reflecting the electrical activity in specific heart segments. Consequently, the ECG serves as a valuable tool for interpreting the electrical activity of the entire heart and is particularly effective in detecting abnormalities in the myocardium. The detection of ECG anomalies has increasingly gained traction among researchers and practitioners. This technique is employed to identify periods of atypical ECG beats. The effectiveness of the anomaly detection method is directly correlated with the accuracy of cardiac disease identification and diagnosis. False alarms often arise when the algorithm misclassifies certain ECG artifacts as anomalous beats; in reality, some artifacts may represent normal beats. These artifacts can stem not only from the heart's electrical activity but also from external noise interference. While artifacts are frequently encountered in standard recordings, they can present significant challenges in medical treatment, as highlighted in various medical research studies, due to the difficulty in eliminating them. Numerous studies in the medical field have emphasized the need for physicians to remain vigilant regarding the issue of artifacts [2].

The primary objective of this project is to effectively predict abnormal heartbeats from normal signals through anomaly detection. Typically, cardiologists perform this prediction manually, which carries the risk of misinterpretation and potential fatal consequences. By employing deep learning techniques, the model is capable of identifying even subtle variations in heartbeats. The central focus of the project is to utilize the auto-encoder concept to predict abnormal heartbeats from time series ECG signals. By integrating auto encoders within LSTM models in deep learning, this project aims to achieve accurate predictions of anomalies in heart signals [3].

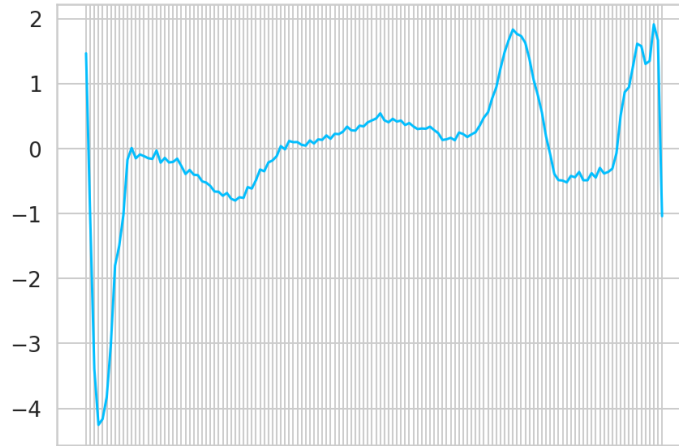


Fig 1: Normal ECG Signal

Abnormal heart rhythms can occur in both healthy individuals and those with various types of structural heart disease. Premature ventricular complexes are identified as single or paired unifocal beats that lack a preceding P wave, exhibit a wide QRS complex with increased amplitude typically exceeding 0.14 seconds, and feature a T wave that is oppositely polarized to the PVC. These complexes emerge early in the cardiac cycle and are more prevalent during episodes of bradyarrhythmia. While there are no P waves preceding the wide QRS complex, retrograde activation of the atria may lead to P waves that appear after the PVC or are obscured within their T waves. The occurrence of premature ventricular beats can be attributed to factors such as elevated catecholamines, myocardial ischemia or damage, electrolyte disturbances, and certain medications, including digitalis and class IA and IC antiarrhythmic drugs [4]. Figures 1-2 displays the normal and abnormal ECG signals.




Fig 2: Abnormal ECG Signal

The irregular morphology of the P' wave, the inconsistent duration of the PP interval, and the prolonged duration of the P'R interval exceeding 0.12 seconds are noteworthy. It is crucial to recognize that if a premature atrial contraction (PAC) occurs very early in the cardiac cycle, it may result in the absence of ventricular activation or the onset of re-entrant atrial tachycardia.

In individuals without underlying conditions, PACs can be triggered by various factors, including the consumption of tobacco, caffeine, and alcohol, as well as intense emotional experiences. Additional circumstances that may contribute to the occurrence of PACs include myocardial infarction, certain medications, infections, hypokalemia, and hypomagnesemia. Hypokalemia is characterized by progressive ST segment depression, gradual flattening or inversion of T waves, the emergence of U waves, and an increase in both the amplitude and duration of P waves and QRS complexes, along with a slight prolongation of the PR interval. Moreover, hypokalemia impacts the automaticity of pacemaker cells, leading to a range of arrhythmias, including sinus bradycardia, atrioventricular block, atrial flutter, and Torsades de Pointes [5].

II. RELATED WORK

Chauhan et al. designed a nine-layer deep convolutional neural network (CNN) aimed at the automatic classification of five distinct categories of heartbeats within ECG signals. The study utilized both original and noise-reduced ECG signal datasets sourced from a publicly accessible database. To balance the representation of the five heartbeat classes, the dataset underwent artificial augmentation and was filtered to eliminate high-frequency noise. The CNN was trained on this augmented dataset, achieving diagnostic classification accuracies of 94.03% and 93.47% for original and noise-free ECGs, respectively. However, when trained on the highly imbalanced original dataset, the CNN's accuracy decreased to 89.07% and 89.3% for noisy and noise-free ECGs, respectively. With appropriate training, the proposed CNN model has the potential to function as an effective tool for ECG screening, facilitating the rapid identification of various types and frequencies of arrhythmic heartbeats [6].

Jones et al. introduced delayed long short-term memory (dLSTM), a method designed for anomaly detection in time-series data. The authors first establish a predictive model using normal (non-anomalous) training data, subsequently conducting anomaly detection by analysing the prediction error of the observed data. However, the presence of multiple states within the waveforms of normal data can diminish prediction accuracy. To address this challenge, the authors employ several prediction models based on LSTM for anomaly detection. The effectiveness of this approach is significantly influenced by the selection process of the appropriate predictive model from the available options. They propose an innovative technique for identifying the suitable predictive model for anomaly detection. This method generates multiple candidate predicted values in advance and selects the one that most closely aligns with the measured value. The selection of the model is deferred until the relevant measured values are obtained. By applying this concept to anomaly detection, dLSTM effectively chooses the appropriate predictive model, thereby improving prediction accuracy. Their experimental results, utilizing both real and synthetic data, demonstrate that dLSTM achieves superior anomaly detection accuracy compared to other methods. Time-series data frequently exhibit seasonal or trend components, leading to distinct and easily identifiable behaviors. In contrast, multi-mode data are defined by having multiple outputs corresponding to a single input. The lack of a one-to-one relationship between input and output data complicates the identification of the correct output behavior, and resolving this issue could facilitate the practical application of this method across various industrial sectors. To develop an accurate predictive model, it is crucial to effectively represent the complex input-output relationship. Numerous studies have explored the use of deep learning techniques for anomaly detection, as deep learning is capable of managing intricate data through the integration of multiple nonlinear activation functions. Recently, long short-term memory networks have also been employed in the realm of anomaly detection, offering the advantage of incorporating contextual information from the sequence [7].

Lavin et al. elucidated the application of Long Short-Term Memory (LSTM) networks for detecting anomalies in temporal datasets. Given the difficulties associated with acquiring labeled anomaly datasets, an unsupervised methodology is adopted. The study involves training recurrent neural networks (RNNs) equipped with LSTM units to identify normal time series patterns and forecast future values. The discrepancies in predictions are subsequently modeled to generate anomaly scores. The research explores various methods for maintaining the LSTM state and examines the impact of utilizing a fixed number of time steps on the performance of LSTM in prediction and anomaly detection. Additionally, LSTMs are contrasted with feed-forward neural networks that utilize fixed-size time windows for input data. The experiments conducted on three real-world datasets indicate that while LSTM RNNs are effective for general time series modeling and anomaly detection, the preservation of LSTM state is essential for achieving optimal outcomes. Furthermore, it is suggested that LSTMs may not be necessary for simpler time series analyses. In this study, the author predominantly employed quantitative research methodologies. Experimental methods were utilized alongside deductive reasoning. The algorithm's implementation was characterized by an exploratory and iterative approach, progressing from basic to more advanced techniques. The algorithm underwent evaluation across various datasets to determine its effectiveness, accompanied by rigorous quality assurance measures. One strategy for detecting anomalies in temporal data involves constructing prediction models and utilizing the prediction errors—the discrepancies between predicted and actual values—to calculate an anomaly score.

A diverse range of prediction models, from simple to complex, has been applied. In a straightforward window-based method, the median of recent values is computed as the predicted value, with a threshold on prediction errors established to identify outliers. In this paper, the authors develop a one-step-ahead prediction model, designating a data point as an anomaly if it lies outside a prediction interval derived from the standard deviation of the prediction errors. The authors also compare various predictive approaches [8].

Malhotra et al. developed a classifier aimed at distinguishing between healthy ECG signals and those that suggest the presence of Arrhythmia. This methodology necessitates an understanding of various Arrhythmia types and corresponding data for effective training. The MIT-BIH Arrhythmia Database has been employed to gather ECG time series data, encompassing both normal intervals and segments characterized by four distinct Arrhythmia types: Premature Ventricular Contraction (PVC), Atrial Premature Contraction (APC), Paced Beats (PB), and Ventricular Couplet (VC). The findings are encouraging, suggesting that Deep LSTM models could be effective in identifying anomalies within ECG signals. Nevertheless, given the heart's intricate nature and the emergence of numerous Arrhythmia types not included in the initial training dataset, it may be more judicious to implement an anomaly detection strategy for ECG signal analysis. In this study, a deep recurrent neural network architecture featuring Long Short-Term Memory (LSTM) units is employed to create a predictive model for healthy ECG signals. Additionally, the probability distribution of prediction errors from these recurrent models is utilized to discern between normal and abnormal behaviors. A notable benefit of employing LSTM networks is the ability to input ECG signals directly into the network without the need for extensive pre-processing, which is often required by alternative methods. Furthermore, the networks do not require prior knowledge of abnormal signals, as they are trained exclusively on normal data [9].

III. PROPOSED METHOD

The proposed system architecture is displayed in Figure 3. It has five modules: Data collection & pre-processing, Data analytics, Data modelling, Data validation & testing and Deployment.

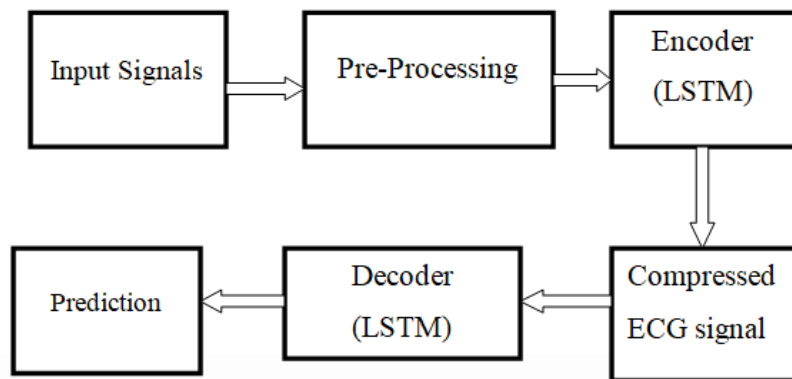


Fig 3: Proposed System Architecture

Data Collection & Pre-Processing:

This project utilizes two distinct datasets: the heart diseases survey dataset and the ECG5000 MIT dataset. The heart diseases survey dataset serves the purpose of data analytics, while the ECG5000 MIT dataset is employed for data modeling, validation, and testing. During the pre-processing phase, we isolate only those ECG signal records that are classified as normal. The autoencoder is specifically designed to learn the representation of normal heartbeat data; therefore, only records with a normal target value are included in the training and validation sets, whereas the abnormal records are reserved for the testing set.

Data Analytics:

Data analysis serves to investigate the target variable and its associated features, allowing us to determine whether certain features require transformation or normalization based on their distribution. Additionally, it may be necessary to eliminate features that do not contribute meaningful information for predicting future outcomes, or to create new features that could enhance predictive accuracy. We can produce visual representations and tables to address questions pertaining to the data and its characteristics.

Typically, we rely on established libraries for exploratory data analysis, which aids in gaining a clearer understanding of the problem. This understanding is crucial for effective data modeling and for selecting appropriate hyperparameters and model types.

Data Modelling:

The models in question are mathematical algorithms that undergo training with the use of data. This training involves processing a substantial volume of data through an algorithm aimed at maximizing likelihood or minimizing cost, resulting in a trained model. By analyzing data from various wells under different conditions, the model acquires the ability to identify various patterns and differentiate them from standard operational behavior. The model endeavors to emulate a specific decision-making process that a team of experts would undertake if they had access to all pertinent data. The modeling process entails training a machine learning algorithm to predict labels based on features, fine-tuning it to meet business requirements, and validating it against holdout data. The outcome of this modeling effort is a trained model capable of inference, allowing for predictions on new data points. It is important to note that modeling operates independently of prior steps in the machine learning workflow and utilizes standardized inputs, enabling modifications to the prediction problem without necessitating a complete rewrite of the code. In this project, I employed the autoencoder model for the purpose of model training.

Data Validation & Testing:

Testing encompasses a variety of assessments aimed at thoroughly evaluating the computer-based system. While each assessment serves a distinct purpose, they collectively ensure that all components of the system are effectively integrated and fulfill their designated functions. The testing process involves verifying that the developed system operates in accordance with its specified requirements and objectives. The underlying principle of testing is to identify errors. An effective test is characterized by its high likelihood of detecting previously unrecognized errors, while a successful test is one that reveals such errors. Test cases are created with this objective in mind, consisting of a collection of data that the system will process as input.

Deployment:

The deployment phase is a critical stage of a project where the final product is released to the end-users. The key aspects of deployment plans are given below: **Deployment Plan:** Outline the steps for deploying the project deliverables, including timelines, resources, and responsibilities. **Target Environment:** Describe the production environment where the project will be deployed, including hardware, software, and network configurations. **Monitoring and Maintenance:** Describe the ongoing monitoring and maintenance activities to ensure the project continues to meet requirements and perform optimally. **Post-Deployment Review:** Schedule a review to assess the deployment's success, identify lessons learned, and document best practices for future projects.

IV. RESULTS AND DISCUSSION

The experimental results of the proposed scheme are discussed in this section. In this project, we are implementing a Long Short-Term Memory (LSTM) autoencoder. The LSTM architecture will be utilized in both the encoder and decoder components. The structure is illustrated in the figure below. An autoencoder's most basic form is a feedforward, non-recurrent neural network, akin to single-layer perceptrons that are part of a multi-layer perceptron (MLP). This configuration consists of an input layer and an output layer, interconnected by one or more hidden layers. Notably, the output layer contains an equal number of nodes (neurons) as the input layer. The primary objective of the autoencoder is to reconstruct its inputs by minimizing the discrepancy between the input and output, rather than predicting a target value based on the inputs. Consequently, autoencoders function as unsupervised learning models, as they do not necessitate labeled inputs for the learning process. The encoder and decoder can be conceptualized as transitions. Autoencoders are primarily employed for dimensionality reduction and information retrieval, although contemporary variations have demonstrated effectiveness across various tasks. Time series prediction presents a challenging category of predictive modeling, as it introduces the complexity of sequential dependencies among input variables, unlike traditional regression modeling. Recurrent Neural Networks (RNNs) are a robust class of neural networks specifically designed to manage sequence dependencies. The LSTM network, a subtype of RNN, is widely used in deep learning due to its capability to train large architectures effectively. The LSTM network employs Backpropagation Through Time for training and addresses the vanishing gradient issue, enabling the construction of extensive recurrent networks suitable for tackling complex problems. Each unit within the LSTM functions as a mini-state machine, with gates that possess weights learned during the training process. As vectors traverse through the neural network, they undergo numerous transformations through various mathematical operations, facilitating the processing of data and the propagation of information forward. An autoencoder is a specific category of artificial neural network designed to learn efficient data encodings in an unsupervised fashion.

The primary objective of an autoencoder is to develop a representation (encoding) for a dataset, often for the purpose of dimensionality reduction, by training the network to disregard extraneous signal "noise." In addition to the encoding process, the autoencoder also learns a reconstruction process, wherein it endeavors to produce a representation that closely resembles the original input from the reduced encoding, which is reflective of its name. Various adaptations exist that aim to compel the learned representations to exhibit beneficial characteristics [10 -12]. The results are shown in Figures 4-7.




Fig 4: Data Analytics Report

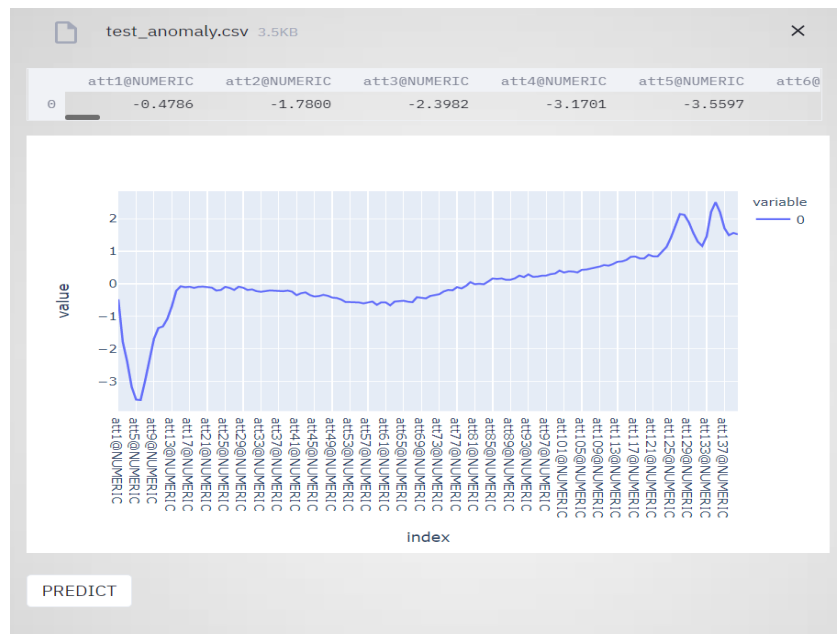


Fig 5: Signal Visualisation




Fig 6: Prediction-1




Fig 7: Prediction-2

V. CONCLUSION

This paper proposed an effective framework for minimizing false alarm rates in abnormal driving detection using Long Short-Term Memory and Electrocardiogram Data. In evaluating the efficacy of anomaly detection, the Long Short-Term Memory (LSTM) neural network demonstrates superior performance compared to both recurrent neural networks (RNNs) and convolutional neural networks (CNNs). The training duration for the LSTM model utilizing sequential data is significantly shorter than that required for training a CNN with image inputs for each ECG signal record. Unlike RNNs, LSTMs possess the capability to retain information over extended periods, facilitated by their forget, input, and output gates. Furthermore, the autoencoder excels in capturing the representation of patterns within normal ECG signal data when compared to other neural network architectures. With advancements in technology, we are now able to process larger volumes of data within a few hours, enhancing both efficiency and speed.

Utilizing the Streamlit framework in Python allows for straightforward testing and evaluation of the model. This framework offers a user-friendly library that simplifies model deployment without necessitating the use of front-end languages such as HTML, CSS, or JavaScript. It streamlines data caching and accelerates computation pipelines. By leveraging robust Python libraries, complex tasks such as pattern recognition, anomaly detection, and future problem prediction in deep learning and machine learning can be effectively accomplished. Rather than writing extensive code for intricate tasks, existing frameworks like TensorFlow, Keras, and PyTorch can be imported to address these challenges. The autoencoder plays a crucial role in this project. Additionally, I intend to apply this concept within a Generative Adversarial Network (GAN), which will aid in reconstructing the input data provided to the model. In this model, the threshold is determined based on the loss, allowing us to establish targets according to the model's error.

REFERENCES

- [1]. D. J. Hill and B. S. Minsker, "Anomaly Detection in Streaming Environmental Sensor Data: A Data-Driven Modeling Approach," *Environmental Modelling & Software*, vol. 25, no. 9, pp. 1014–1022, 2010.
- [2]. D. George. How the brain might work: A hierarchical and temporal model for learning and recognition. PhD Thesis, Stanford University, 2008.
- [3]. E. Keogh, J. Lin, and A. Fu, "HOT SAX: Efficiently finding the most unusual time series subsequence," in *proceedings of Fifth IEEE International Conference on Data Mining*, pp. 226-233, November 2005.
- [4]. F. Serdio and E. Lughofer and K. Pichler and M.Pichler and T. Buchegger and H. Efendic. (2014) "Fault Detection in Multi-Sensor Networks based on Multivariate Time-Series Models and Orthogonal Transformations", *Information Fusion*, vol. 20, pp. 272–291.
- [5]. [5] N. Laptev, S. Amizadeh, and I. Flint. (2015) Generic and Scalable Framework for Automated Time-series Anomaly Detection. In *Proceedings of the 21th ACM SIGKDD International Conference*.
- [6]. [6] S. Chauhan and L. Vig, "Anomaly Detection in ECG Time Signals via Deep Long Short-Term Memory Networks," in *2015 IEEE International Conference on Data Science and Advanced Analytics (DSAA)*, Oct 2015.
- [7]. M. Jones, D. Nikovski, M. Imamura, and T. Hirata, "Anomaly Detection in Real-Valued Multidimensional Time Series," in *International Conference on Bigdata/Socialcom/Cybersecurity*. Stanford University,
- [8]. A. Lavin and S. Ahmad, "Evaluating Real-Time Anomaly Detection Algorithms - The Numenta Anomaly Benchmark," in *2015 IEEE 14th International Conference on Machine Learning and Applications*. 10.1109/ICMLA.2015.141 pp. 38–44.
- [9]. P. Malhotra, L. Vig, G. Shroff, and P. Agarwal, "Long Short Term Memory Networks for Anomaly Detection in Time Series," in *Proceedings 2015*.
- [10]. F. Gers, N. Schraudolph and J. Schmidhuber, "Learning precise timing with LSTM recurrent networks", *Journal of Machine Learning Research*, vol. 3, pp. 115143, 2002.
- [11]. Sak Hasim, Senior Andrew and Beaufays Françoise, "Long short-term memory recurrent neural network architectures for large scale acoustic modeling", *Proceedings of the Annual Conference of International Speech Communication Association (INTERSPEECH)*, 2014.
- [12]. P. de Chazal, B.G. Celler and R.B. Rei, "Using wavelet coefficients for the classification of the electrocardiogram", *Proceedings of the 22nd Annual EMBS International Conference*, July 23-28, 2000.