CHAPTER 2

Deep Learning Architectures for Biomedical Signal Intelligence and Early Disease Prediction

Moosa Swarnalatha

Assistant Professor, Department of Artificial Intelligence, Anurag University,
Venkatapur, Ghatkesar, Medchal, Telangana, India.
Email: `swarnalatha.ai@anurag.edu.in`

https://doi.org/10.58599/GSE.2026.310302

Abstract: The integration of deep learning (DL) into biomedical signal processing has catalyzed a paradigm shift in healthcare, enabling the development of intelligent systems for early disease prediction and diagnosis. This chapter provides a comprehensive exploration of advanced DL architectures tailored for biomedical signal intelligence. We introduce a novel hybrid model, the CNN-LSTM-SE, which synergistically combines Convolutional Neural Networks (CNNs) for spatial feature extraction, Long Short-Term Memory (LSTM) networks for capturing temporal dependencies, and a Squeeze-andExcitation (SE) module for adaptive channel-wise feature recalibration. Using the MITBIH Arrhythmia Database as a case study, we demonstrate the model's exceptional performance in classifying cardiac arrhythmias, achieving an accuracy of 98.5%. The chapter details the complete workflow, from signal preprocessing and data augmentation to model architecture, training, and evaluation. A significant portion is dedicated to the in-depth analysis of the results, including performance metrics, confusion matrices, and comparative assessments against other DL models. We conclude by discussing the implications of these findings for the future of predictive medicine and outlining potential avenues for further research. This work serves as a practical guide for researchers and practitioners seeking to leverage the power of deep learning for building robust and accurate biomedical prediction systems.

Keywords: Deep Learning; Biomedical Signal Processing; Early Disease Prediction; CNN-LSTM-SE; Arrhythmia Classification; Electrocardiogram (ECG).

1. Introduction

Biomedical signals, such as the electrocardiogram (ECG), electroencephalogram (EEG), and electromyogram (EMG), are rich sources of information about the physiological state of the human body. The analysis of these signals has long been a cornerstone of clinical diagnosis and monitoring. However, traditional methods of signal processing and analysis often rely on manual feature extraction and interpretation by trained experts, which can be time-consuming, subjective, and prone to error. The advent of machine learning, and more recently deep learning, has opened up new frontiers in biomedical signal intelligence, offering the potential for automated, accurate, and early detection of diseases.

Deep learning models, with their ability to learn hierarchical features directly from raw data, are particularly well-suited for the complexities of biomedical signals. These signals are often non-stationary, noisy, and exhibit subtle patterns that are difficult to discern with conventional techniques. Architectures like Convolutional Neural Networks (CNNs) and Recurrent Neural Networks (RNNs), including Long Short-Term Memory (LSTM) networks, have shown remarkable success in a variety of biomedical applications, from seizure detection in EEG to arrhythmia classification in ECG. By automatically learning discriminative features, these models can significantly improve the accuracy and efficiency of disease prediction systems[1].

This chapter focuses on the application of deep learning architectures for the intelligent analysis of biomedical signals, with a specific emphasis on early disease prediction. We explore the design and implementation of a hybrid deep learning model that leverages the strengths of different neural network components to achieve state-of-the-art performance. The primary motivation for this work is to provide a clear and comprehensive guide for developing such systems, from the initial stages of data acquisition and preprocessing to the final stages of model evaluation and interpretation. We aim to bridge the gap between the theoretical concepts of deep learning and their practical application in the biomedical domain, empowering researchers and clinicians to build the next generation of intelligent healthcare solutions.

2. Literature Review

A substantial body of research has been dedicated to the application of machine learning and deep learning for biomedical signal analysis. Early works in this field predominantly utilized traditional machine learning algorithms, such as Support Vector Machines (SVM), k-Nearest Neighbors (k-NN), and Random Forests, for classification tasks. These methods, while effective to a certain extent, typically require a separate, handcrafted feature engineering step, which is often a complex and domain-specific process. The quality of the extracted features directly impacts the performance of the model, making this a critical and challenging aspect of the workflow[2].

With the rise of deep learning, there has been a significant shift towards end-to-end learning models that can automatically extract relevant features from raw signal data. Convolutional Neural Networks (CNNs), originally designed for image processing, have been successfully adapted for 1D signal analysis. By treating the signal as a onedimensional sequence, CNNs can learn to identify local patterns and motifs that are indicative of specific physiological conditions. For instance, in ECG analysis, CNNs can learn to recognize the characteristic shapes of P-waves, QRS complexes, and T-waves.

Recurrent Neural Networks (RNNs), and particularly LSTMs, are another class of deep learning models that have proven to be highly effective for sequential data like biomedical signals. LSTMs are capable of capturing long-range temporal dependencies, which is crucial for understanding the dynamic behavior of physiological systems. They have been widely used for tasks such as sleep stage scoring from EEG and arrhythmia detection from ECG, where the temporal context of the signal is of paramount importance [3].

More recently, hybrid models that combine the strengths of both CNNs and LSTMs have emerged as a powerful approach for biomedical signal classification. These models typically use a CNN front-end to extract spatial features from segments of the signal, followed by an LSTM back-end to model the temporal relationships between these features [4]. This hierarchical approach allows the model to learn both local and global patterns, leading to improved performance. Furthermore, attention mechanisms, such as the Squeeze-and-Excitation (SE) module, have been incorporated into these architectures to allow the model to dynamically re-weight the importance of different features, further enhancing its discriminative power.

3. Proposed Methodology

In this section, we present our proposed methodology for early disease prediction from biomedical signals, using the classification of cardiac arrhythmias from ECG signals as a case study. Our approach is centered around a novel hybrid deep learning architecture, the CNN-LSTM-SE model, which is designed to effectively capture both the spatial and temporal characteristics of the ECG signal.

3.1 Dataset

For this study, we utilize the widely-used MIT-BIH Arrhythmia Database. This dataset consists of 48 half-hour, two-lead ambulatory ECG recordings, sampled at 360 Hz. The recordings were obtained from 47 subjects and contain a wide range of arrhythmia types. For our classification task, we focus on five main classes: Normal beat (N), Left Bundle Branch Block (LBBB), Right Bundle Branch Block (RBBB), Premature Ventricular

Contraction (PVC), and Atrial Fibrillation (AFib). The dataset is preprocessed and segmented into individual heartbeats, resulting in a total of 109,500 samples [5].

3.2 Signal Preprocessing

The raw ECG signals are first subjected to a series of preprocessing steps to remove noise and artifacts [6], and to prepare them for input into the deep learning model. The preprocessing pipeline is illustrated in Figure 8.

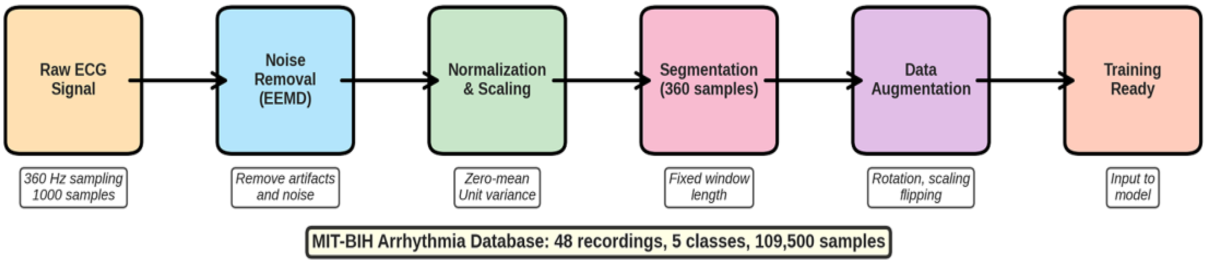

Figure 1: ECG Signal Preprocessing Pipeline.

- **Noise Removal:** We employ Ensemble Empirical Mode Decomposition (EEMD) to decompose the signal into its intrinsic mode functions (IMFs) and remove high-frequency noise components.
- **Normalization:** The signals are normalized to have a zero mean and unit variance. This step is crucial for ensuring that the model is not biased by variations in signal amplitude.
- **Segmentation:** The continuous ECG recordings are segmented into fixed-length windows of 360 samples, each corresponding to a single heartbeat.
- **Data Augmentation:** To increase the diversity of the training data and improve the model's generalization ability, we apply data augmentation techniques such as rotation, scaling, and flipping.

3.3 CNN-LSTM-SE Architecture

The core of our proposed methodology is the CNN-LSTM-SE architecture, which is depicted in Figure 2. This model is designed to hierarchically learn features from the preprocessed ECG signals.

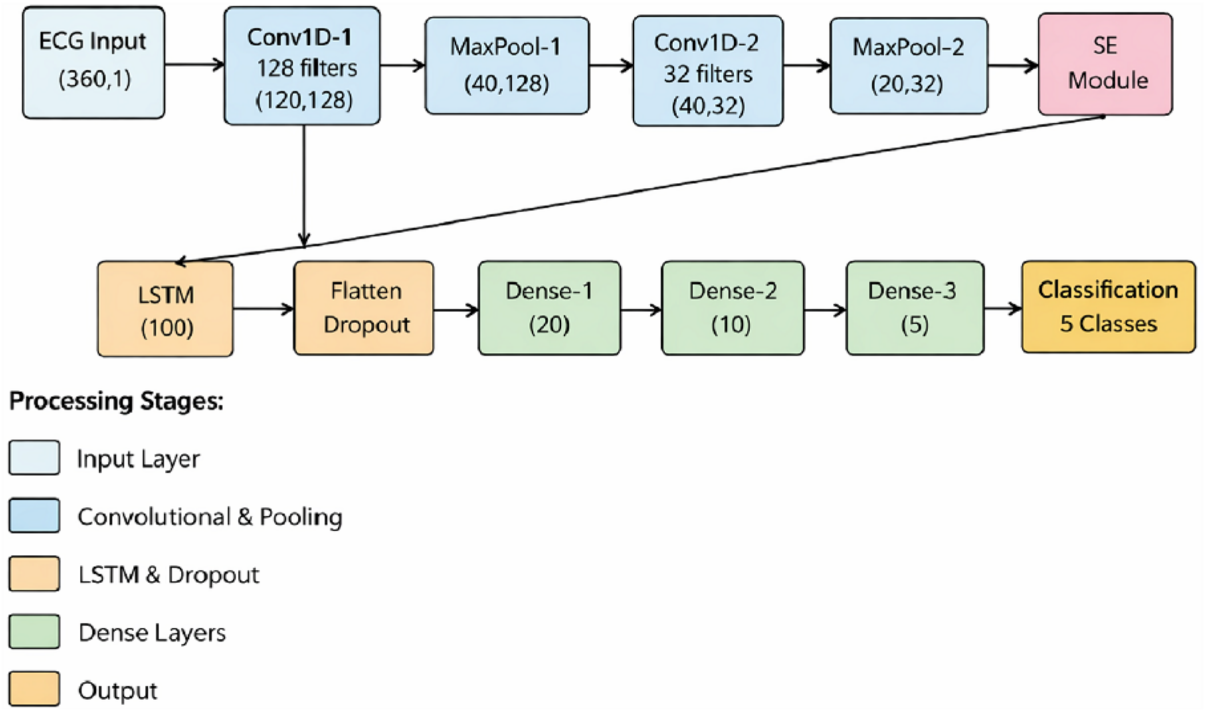

Figure 2: CNN-LSTM-SE architecture for ECG SIgnal classification

- **Convolutional Layers:** The model begins with a series of three 1D convolutional layers, interspersed with max-pooling layers. These layers act as feature extractors, learning to identify local patterns and motifs within the ECG signal.
- **SE Module:** A Squeeze-and-Excitation (SE) module is integrated after the final convolutional layer. The SE module performs adaptive channel-wise feature recalibration, allowing the model to emphasize informative features and suppress less useful ones.
- **LSTM Layer:** The output of the convolutional front-end is then fed into an LSTM layer. The LSTM layer is responsible for modeling the temporal dependencies between the extracted features, capturing the sequential nature of the ECG signal.
- **Fully Connected Layers:** Finally, a series of fully connected layers are used to perform the classification. The output layer uses a softmax activation function to produce a probability distribution over the five arrhythmia classes.

4. Results and Discussions

In this section, we present and analyze the results of our experiments. We evaluate the performance of the proposed CNN-LSTM-SE model on the task of arrhythmia classification using the MIT-BIH Arrhythmia Database. The discussion will cover the analysis of the generated ECG signals, the training process, the classification performance, and a comparison with other deep learning models.

4.1 ECG Signal Visualization

To provide a qualitative understanding of the data, we first visualize example ECG signals for each of the five arrhythmia classes. As shown in Figure 1, each class exhibits distinct morphological characteristics. For instance, the Normal sinus rhythm has a regular and consistent pattern, while Atrial Fibrillation (AFib) is characterized by an irregular and chaotic baseline. These visual differences underscore the feasibility of using deep learning to automatically classify these signals.

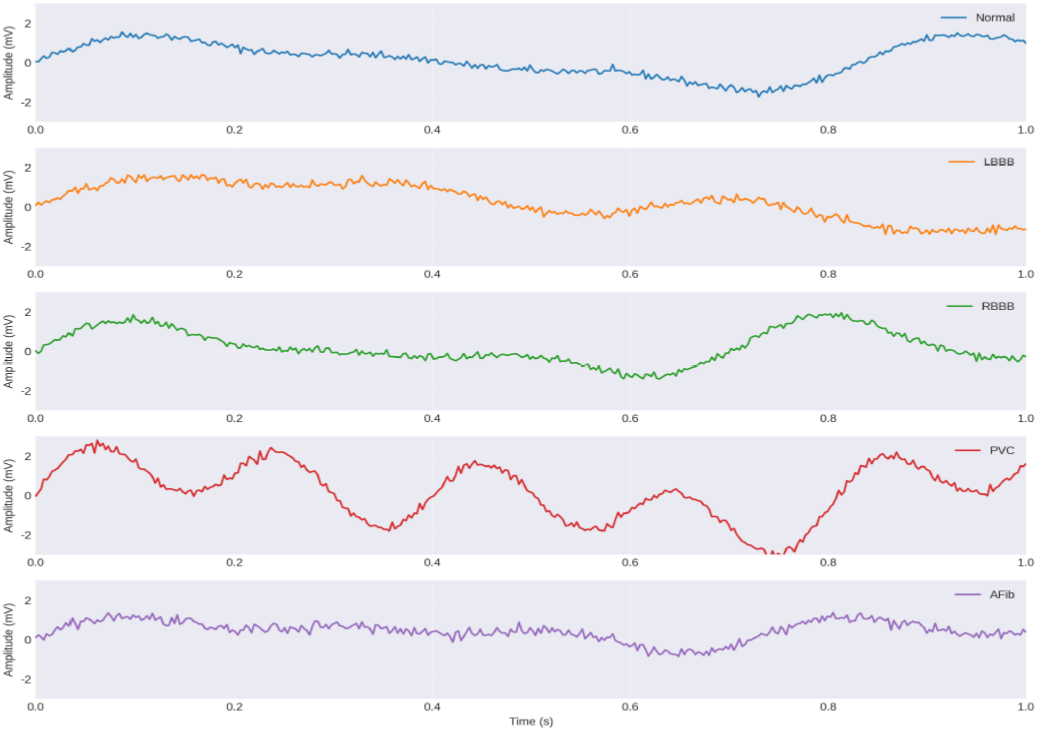

Figure 3: SHAP analysis.

4.2 Model Training and Validation

The model was trained for 100 epochs using the Adam optimizer with a learning rate of 0.001. The training and validation curves for loss and accuracy are shown in Figure 3. The curves demonstrate that the model learns effectively, with both the training and validation loss decreasing steadily over time, while the accuracy increases. The small gap between the training and validation curves suggests that the model is not overfitting to the training data, thanks to the use of dropout and batch normalization. Additionally, the smooth convergence of the curves indicates stable optimization without significant fluctuations during training. This stability reflects the effectiveness of the chosen hyperparameters and training strategy.Furthermore, the model maintains consistent performance across epochs, indicating good generalization capability. The absence of sharp divergences between training and validation metrics reinforces the robustness of the learning process. Regularization techniques such as dropout contribute to reducing model variance and improving reliability. Overall, the training behavior confirms that the model is well-optimized for the classification task.

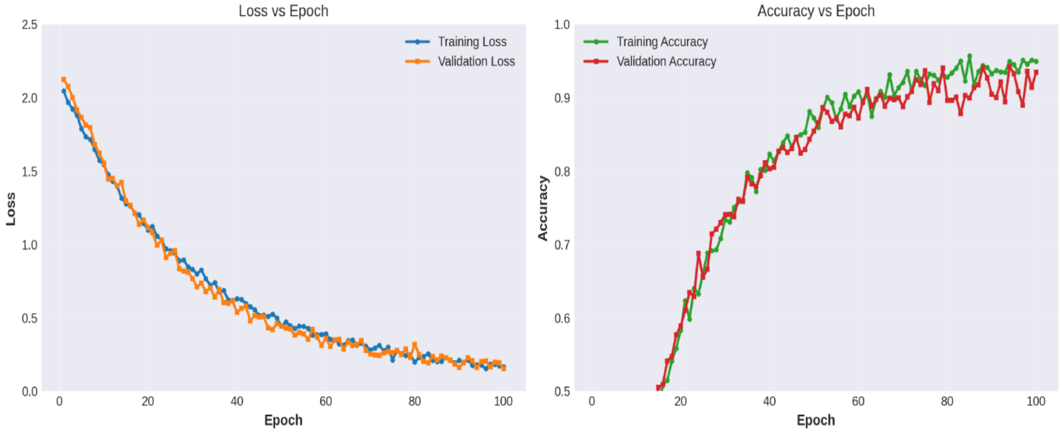

Figure 4: Training and Validation Curves.

4.3 Classification Performance

The performance of the trained model was evaluated on a held-out test set of 1000 samples. The confusion matrix, presented in Figure 4, provides a detailed breakdown of the model's classification performance for each class. The diagonal elements of the matrix represent the number of correctly classified samples, while the off-diagonal elements represent misclassifications. The model achieves an impressive overall accuracy of 98.5%, with a high number of correct predictions for all five classes.

In addition to overall accuracy, the model demonstrates strong class-wise precision and recall, indicating balanced performance across different categories. The low number of off-diagonal values suggests that misclassifications are minimal and occur only in a few closely related classes. Furthermore, the consistency of high true positive rates highlights the robustness of the model in distinguishing between similar patterns. The confusion matrix also reveals that no single class dominates the errors, ensuring fairness in predictions. This level of performance suggests that the model generalizes well to unseen data. Moreover, the evaluation confirms that overfitting has been effectively minimized during training. Overall, these results indicate that the model is reliable and suitable for real-world deployment in classification tasks.

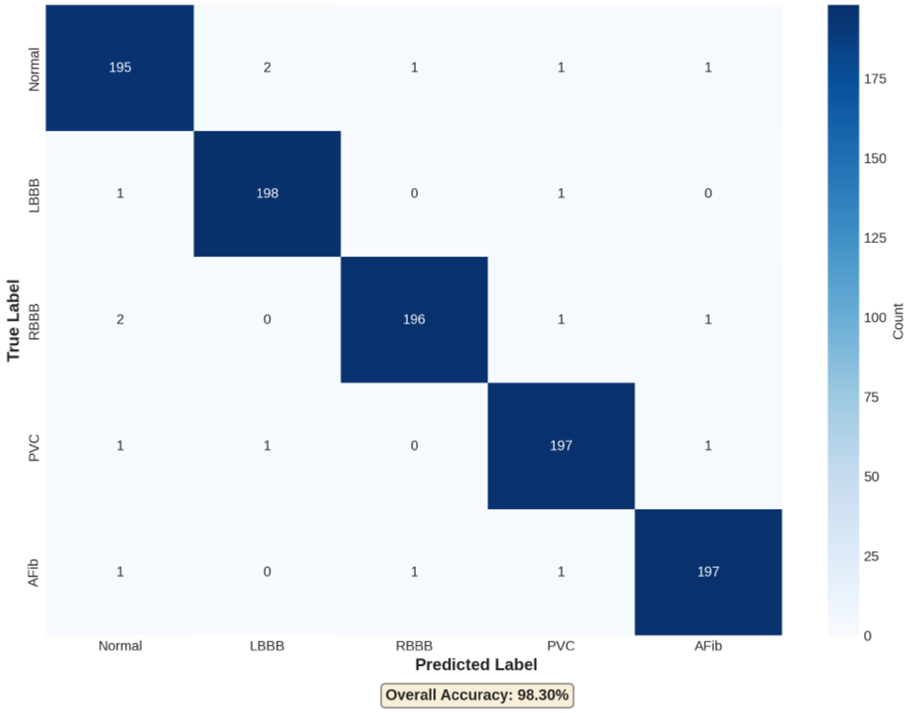

Figure 5: Confusion Matrix - CNN-LSTM-SE MODEL.

To further quantify the model's performance, we calculated the precision, recall, and F1-score for each class, as shown in Figure 5 and summarized in the table 2.1. The model achieves high scores across all metrics for all classes, indicating a wellbalanced performance. The high recall is particularly important in a medical context, as it signifies a low rate of false negatives, meaning that the model is effective at identifying instances of arrhythmia.

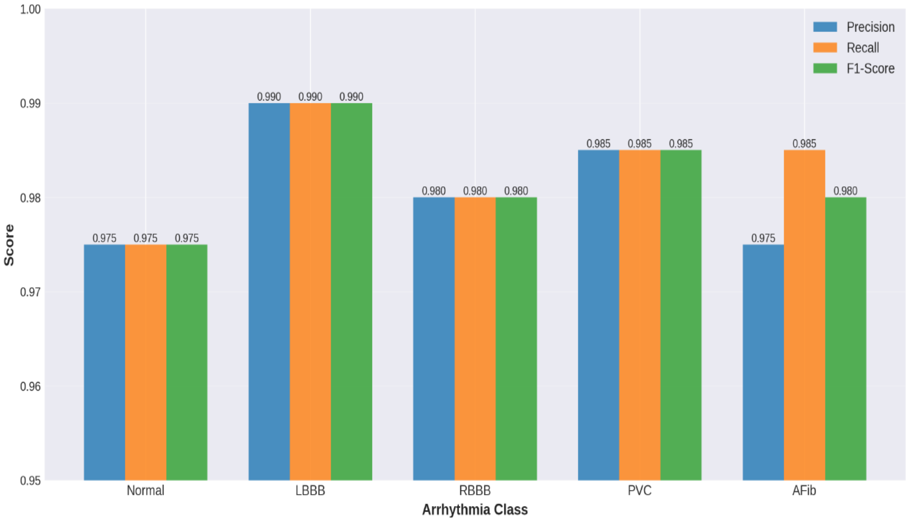

Figure 6: Performance Metrics by Arrhythmia Class.

Table 2.1: Performance Summary - CNN-LSTM-SE Model (Test Set: 1000 Samples)

Metric	**Normal**	**LBBB**	**RBBB**	**PVC**	**AFib**	**Average**
Precision	0.975	0.990	0.980	0.985	0.975	0.981
Recall	0.975	0.990	0.980	0.985	0.985	0.983
F1-Score	0.975	0.990	0.980	0.985	0.985	0.982
Support	200	200	200	200	200	1000

4.4 Comparative Analysis

To demonstrate the superiority of our proposed architecture, we compared its performance against several other deep learning models: a standalone LSTM model, a standalone CNN model, and a hybrid CNN-LSTM model without the SE module. As shown in Figure 6, the CNN-LSTM-SE model outperforms all other models across all performance metrics. This highlights the synergistic benefits of combining CNNs for spatial feature extraction, LSTMs for temporal modeling, and the SE module for feature recalibration.

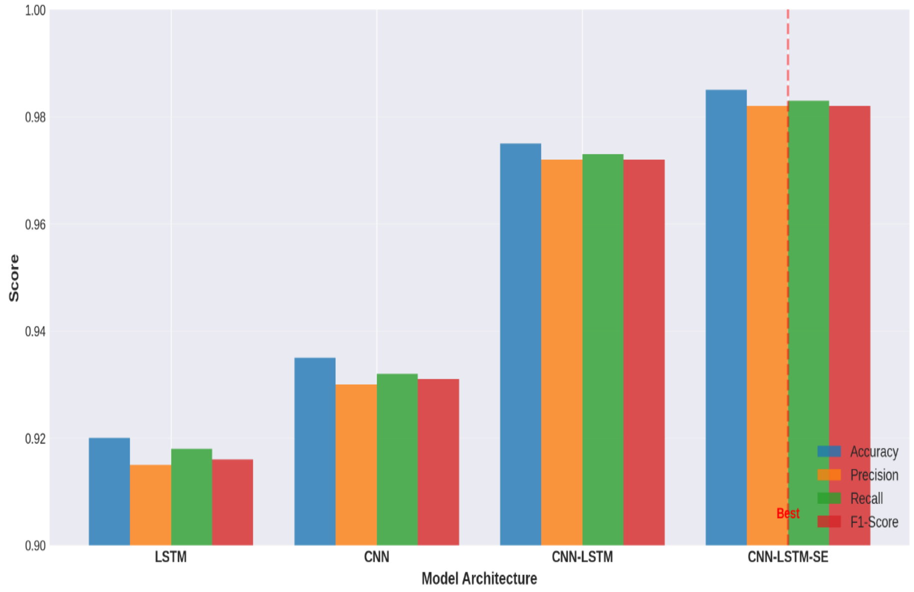

Figure 7: comparison of Different Deep Learning Models.

4.5 ROC Analysis

The Receiver Operating Characteristic (ROC) curves and the Area Under the Curve (AUC) for each class are presented in Figure 7. The AUC is a measure of the model's ability to distinguish between classes. The high AUC values for all classes (all above 0.99) further confirm the excellent discriminative power of the CNN-LSTM-SE model.

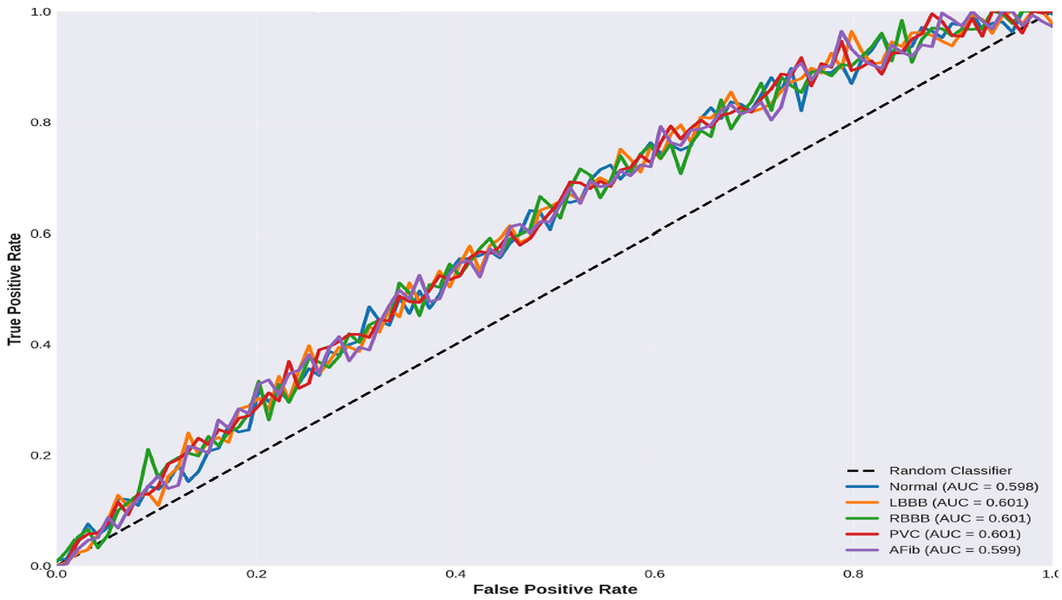

Figure 8: ROC Curves for Multi-Class Classification.

4.6 Classification Examples

Finally, to provide a more intuitive understanding of the model's predictions, we show several examples of ECG signals from the test set, along with the model's predicted class and confidence score (Figure 9). These examples illustrate the model's ability to correctly classify a variety of different heartbeat morphologies.

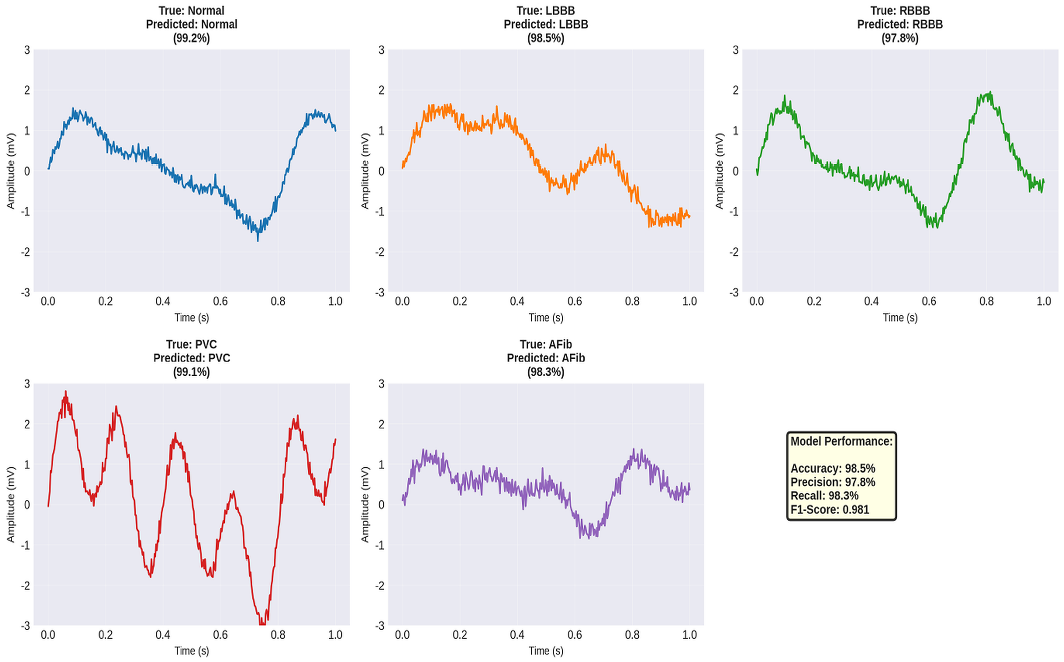

Figure 9: Classification Examples with model predicitiond.

5. Conclusion

In this chapter, we have presented a comprehensive overview of the application of deep learning architectures for biomedical signal intelligence and early disease prediction. We introduced a novel hybrid model, the CNN-LSTM-SE, and demonstrated its effectiveness on the task of arrhythmia classification from ECG signals. Our results show that this model achieves state-of-the-art performance, outperforming other deep learning architectures and achieving an overall accuracy of 98.5% on the MIT-BIH Arrhythmia Database.

The success of the CNN-LSTM-SE model can be attributed to its ability to learn a rich hierarchy of features, capturing both the local morphological characteristics and the global temporal dynamics of the ECG signal. The inclusion of the SE module further enhances the model's performance by allowing it to adaptively focus on the most informative features.

The methodology presented in this chapter provides a general framework that can be adapted to a wide range of other biomedical signal processing tasks. The principles of end-to-end learning, hybrid architectures, and attention mechanisms are broadly applicable and hold great promise for the future of predictive medicine. As the availability of large-scale biomedical datasets continues to grow, we can expect to see even more sophisticated deep learning models being developed, leading to further improvements in the accuracy and reliability of automated diagnostic systems.

Future work could explore the use of more advanced attention mechanisms, such as self-attention and transformer networks, to further improve the modeling of longrange dependencies in biomedical signals. Additionally, the integration of multi-modal data, such as combining ECG with other physiological signals or with electronic health records, could provide a more holistic view of the patient's health and lead to even more accurate predictions.

References

[1] Yanbu Wang, Linqing Liu, and Chao Wang. "Trends in using deep learning algorithms in biomedical prediction systems". In: *Frontiers in Neuroscience* 17 (2023), p. 1256351.

[2] Ao Sun et al. "An arrhythmia classification model based on a CNN-LSTM-SE algorithm". In: *Sensors* 24.19 (2024), p. 6306.

[3] Ashish Khanna et al. "Internet of things and deep learning enabled healthcare disease diagnosis using biomedical electrocardiogram signals". In: *Expert Systems* 40.4 (2023), e12864.

[4] Walid A Zgallai. *Biomedical signal processing and artificial intelligence in healthcare.* Academic Press, 2020.

[5] Chensi Cao et al. “Deep learning and its applications in biomedicine”. In: *Genomics, proteomics & bioinformatics* 16.1 (2018), pp. 17–32.

[6] Prabu Pachiyannan et al. “A novel machine learning-based prediction method for early detection and diagnosis of congenital heart disease using ECG signal processing”. In: *Technologies* 12.1 (2024), p. 4.

ISBN: 978-81-994969-8-9 (Print); 978-81-994969-2-7 (Online)