A Hybrid CNN-Random Forest Framework for Interpretable ECG Arrhythmia Detection

[1]Muazu Abubakar Muhammad*, [2]Tarun Kumar Agrawal
[1, 2] Department of Computer Science & Engineering, Vivekananda Global University, Jaipur India
[1]abubakarmuhammadmuaz@gmail.com
[2]tarun.agrawal@vgu.ac.in

Abstract. An automated clinical decision support based on cardiac arrhythmia detection from electrocardiogram (ECG) signals is still a burning issue, especially where high sensitivity, robustness under realistic signal conditions, and post-hoc interpretability are required. This paper proposes a Hybrid Convolutional Neural Network-Random Forest (CNN-RF) framework that addresses these challenges by employing a two-stage pipeline that uses an optimized 1-D CNN, which is trained by ECG-specific data augmentation that serves as a deep feature extractor and outputs only compact 64-dimensional morphological embeddings from raw ECG segments. Then the embeddings are classified by a regularized Random Forest ensemble. Shapley Additive Explanations (SHAP) is used for the Random Forest component to give clinically meaningful post-hoc feature attribution. The proposed Hybrid CNN-RF framework achieved 95.1% accuracy, 95.2% sensitivity, 94.9% specificity, a 95.1% F1-score, and an AUC-ROC of 0.950, matching the standalone CNN while providing interpretable features that a pure deep learning model alone cannot provide. The LSTM baseline showed critical performance degradation with the sensitivity and accuracy of 13.3% and 56.2%, respectively. These indicate that the proposed framework provides a reliable, robust, and interpretable alternative for automated arrhythmia detection, which can be improved and integrated for real-world wearable and clinical deployment scenarios.

Keywords: ECG arrhythmia detection; convolutional neural network; random forest; hybrid deep learning; SHAP interpretability; cardiac classification; wearable health monitoring.

1 Introduction

Cardiovascular diseases (CVDs) continue to be the leading cause of morbidity and mortality all over the world, causing an estimated 17.9 million deaths annually according to the World Health Organization [1]. Among the most clinically important manifestations of CVD can be listed the cardiac arrhythmias, which are abnormal electrical conduction patterns such as atrial fibrillation (AF), premature ventricular contractions (PVC), and ST-segment elevation (STE), which indicate acute myocardial infarction.

B. Singh et al. (eds.), *Proceedings of the International Conference on Advances in Computing Technology and Artificial Intelligence (COMPUTATIA 2026)*, Atlantis Highlights in Intelligent Systems 18,
https://doi.org/10.2991/978-94-6239-713-2_13

The need for early and efficient detection of these conditions is paramount. Fast intervention of arrhythmic events has been shown to significantly curtail the risk of catastrophic incidents such as stroke, sudden cardiac arrest, and heart failure decompensation [2]. The electrocardiogram (ECG) is the most useful non-invasive tool for cardiac rhythm diagnosis, as it encodes the electrical activity of the heart as a time series waveform from which clinically useful morphological features, such as the absence of P waves, QRS complex widening, and ST segment deviation, can be extracted and analyzed. However, manual ECG interpretation by trained clinicians is time-consuming, prone to inter-observer variability, and not practical for continuous ambulatory ECG monitoring on a scale that could meet the clinical need, leaving a huge clinical gap for automated, reliable, and interpretable ECG analysis systems [3].

The use of machine learning and deep learning in the automated classification of ECG arrhythmias has been widely researched in the last decade. Convolutional Neural Networks (CNNs) have been described to show a sufficient efficiency in extracting hierarchical morphological features from raw ECG data without the need for manual feature engineering, with reported accuracies up to 99% on controlled benchmark datasets [4]. Hybrid architecture consisting of a combination of CNN and some recurrent network variants, such as LSTM, has been further studied for joint spatial and temporal pattern extraction [5], [6]. However, despite the current trajectory in modern technology and improvement in these systems, there is still a mountain of limitations in the current studies. First, a large proportion of previous studies report performance on clean, well-curated benchmark datasets, which are not representative of the diagnostic ambiguity, physiological variability, and label uncertainty inherent in the ECG data of real clinical practice. Second, most deep learning models used in ECG classification are black box predictors with little post-hoc interpretability, a trait that is increasingly demanded for the regulatory approval and clinical adoption of AI-based diagnostic systems [6]. Third, the relative performance of CNN-based hybrid classifiers against standalone recurrent models under identical and realistic experimental conditions remains inconsistently evaluated in the literature, leaving important architectural design questions unresolved [4], [5].

To overcome these shortcomings, a hybrid CNN-Random Forest (RF) architecture for automatic ECG arrhythmia detection, which jointly optimizes predictive performance and robustness under realistic signal conditions, as well as post-hoc explainability, is proposed in this paper. The proposed framework uses a regularized 1-D CNN as a trainable feature extractor that generates short 64-dimensional morphological embeddings from raw ECG segments that are then classified by a regularized Random Forest ensemble. The RF component supports native probabilistic outputs and usage of Shapley Additive Explanations (SHAP)-based post-hoc feature attribution that can provide transparent and clinically meaningful explanations of individual classification decisions. In order to simulate realistic diagnostic environments, physiological heart rate variability, three arrhythmia types (AF, PVC, and STE), and 10% label noise are actually added to the experimental data. A standalone LSTM network is then evaluated as a recurrent baseline under the same conditions to obtain a rigorous architectural comparison. The rest of this paper is structured as follows: Section II reviews the related work; Section III presents the methodology proposed in this work; Section IV presents the results and analysis; and Section V concludes the paper with a discussion of limitations and future directions.

2 Related Work

Automated cardiac arrhythmia detection has been given much emphasis over the past decade due to the development of wearable biosensor technology, the Internet of Medical Things (IoMT), and deep learning. This part is a review of the main contributions in four thematically related fields, namely, ECG-based arrhythmia and myocardial infarction detection, multi-signal cardiovascular monitoring, metabolic diseases prediction, and other AI-enabled wearable health systems. A tabulated summary of all reviewed studies is displayed in Table 1.

2.1 ECG-Based Arrhythmia and Myocardial Infarction Detection

A hierarchical, event-driven Random Forest classifier for Myocardial Infarction (MI) detection with wearable ECG validated on the PTB ECG database. Their system achieved over 90% diagnostic accuracy and provides a stunning 2.6x improvement in battery life through energy-efficient event-driven inference. Building upon the real-time event-driven classification, a CNN-LSTM hybrid architecture for early MI detection on wearable systems [6]. By using the spatial feature extraction of CNNs together with the temporal modeling of LSTMs, the system attained high diagnostic accuracy rates using clinical datasets. However, the study was limited by small sample sizes, which made it difficult to understand the generalizability of the results.

AF and heart failure (HF) detection. Their system achieves a sensitivity of about 95% and a specificity of 94% for atrial fibrillation when operating using cloud-based and remote sensing architecture [7]. The results were validated against clinical trial data. Lately, authors [8] used a CNN-LSTM hybrid model trained on ECG, heart rate variability (HRV), and PPG signals for wearable AI for cardiovascular monitoring and presented accuracy between 85% and 98% on benchmark data, although restricted by the imbalance of data and a lack of real-world trials on large populations. Authors [9] examined machine learning algorithms to manage arrhythmias based on data from implantable devices and showed a significant reduction in the number of inappropriate therapeutic shocks but noted that there are still open questions regarding the issue of regulation and data privacy.

2.2 Multi-signal and Multimodal Cardiovascular Monitoring

The CNN-LSTM framework for wearable health monitoring using ECG, temperature, and PPG signals with an accuracy of 85-95% for public physiological datasets [10]. The advanced multi-disease detection with ResNet and U-Net with transfer learning on imaging and EHR data and achieved a sensitivity of more than 90%, but the reliance of the framework on hospital-based imaging hindered its wearable deployment. Authors [11] investigated the detection using ACS with ANN, SVM, and ensemble classifiers with a sensitivity of 88 to 91% under a cloud-based architecture, with a noted risk of overfitting due to limited dataset diversity.

2.3 Metabolic Disease Detection and General AI-Enabled Wearable Monitoring

An IoMT-enabled real-time blood glucose prediction system based on CNN and LSTM models on the continuous glucose monitoring (CGM) data, implemented under the hybrid edge-cloud architecture [9]. The system achieves an elevated predictive correlation and improved latency as compared to the regression baseline. However, the hurdle of adequate clinical generalizability and multi-center validation still looms large on the horizon. A narrative review of the methods of AI in the diagnosis of early CVD, with accuracy ranging from 85% to 99% across different methods and different diseases [11]. The discriminative power of CNN and RNN architectures for Human Activity Recognition (HAR) using inertial sensor data with accuracy exceeding 90% but provided the foundational evidence for the usage of convolutional architectures on wearable time series data [7].

2.4 Research Gap and Positioning of the Proposed Work

The reviewed literature remains consistent in showing the superiority of CNN-based architectures and their hybrid variants in achieving superior diagnostic performance for the ECG-based cardiac classification problem in comparison to standalone recurrent models and classical ML approaches [9–12]. However, three critical gaps exist that are not being addressed. First of all, most previous works on hybrid deep learning do not include the use of interpretable post-hoc attribution techniques like SHAP, a major shortcoming given explainability requirements for clinical AI systems [11]. Second, few studies rigorously assess model robustness to deliberate label noise or physiological variation, leaving unanswered questions about model generalizability in the real world [7], [9]. Third, the long-short-term memory (LSTM) versus the convolution neural network (CNN)-based approach has still been inconsistently reported on fixed-length ECG beat classification tasks [6]. This work aims to address these challenges, particularly through a hybrid-RF framework with controlled label perturbation and physiological noise simulation, SHAP-based explainability, and direct empirical comparison against an LSTM baseline.

Table 1. Summary of Related Work on AI-Based Cardiac and Metabolic Disease Detection

Disease	Data	Model(s)	Key Result	Limitation	Source
MI	ECG	Random Forest	>90% Acc.; 2.6× battery life	Single modality, moderate dataset	[6]

Diabetes	CGM	CNN, LSTM	High predicreative correlation	Limited multi-center validation	[9]
AF & HF	ECG, PPG	ML Wearable	Sens. ≥95%, Spec. ≥94%	Regulatory interoperability	[7]
Wearable Monitor.	ECG, PPG, and temp.	CNN, LSTM	Acc. 85–95%	Computational constraints	[10]
ACS	ECG, biomarkers	ANN, SVM, Ensemble	Sens. 88–91%	Possible overfitting	[11]
CVD (review)	ECG, Clinical	ML & DL (various)	85–99% across methods	Narrative review only	[11]
Arrhythmia & HF	ECG, HRV, PPG	CNN, CNN-LSTM	Acc. 85–98%	Data imbalance; limited trials	[2]
Multi-disease	Imaging, EHR	ResNet, U-Net, TL	>90% Sensitivity	No wearable deployment	[5]
Arrhythmia	ECG, Implantables	ML-based	Reduced inappropriate shocks	Regulatory & privacy challenges	[7]

3 Methodology

The proposed framework consists of 5 steps for conducting automated ECG arrhythmia detection: For synthetic data replication and generation, the MIT-BIH Dataset, signal preprocessing, CNN-based feature extraction, Random Forest (RF) classification, and SHAP-based interpretability analysis. A standalone LSTM model is trained and tested as a comparative baseline. Each stage is described below.

3.1 ECG Data Synthesis and Preprocessing

Using a parameterized Gaussian biophysical model, in accordance with the MIT-BIH Arrhythmia Database standard, we generated a synthetic ECG corpus consisting of 80,000 independent one-second segments (360 samples each) [12]. Normal beats were generated by adding P-wave, QRS complex, and T-wave components with physiologically determined heart rate variability ($hr \sim U(0.8, 1.2)$), baseline wander, and additive Gaussian muscle noise. Arrhythmic beats (50% of the dataset) represented three common rhythm disorders with equal frequency: Atrial Fibrillation (AF), marked by absent and/or sporadic P-waves and variable RR-intervals; Premature Ventricular Contraction (PVC), which has a wide QRS and inverted T-wave; and ST-Elevation (STE), signaling myocardial infarction. In order to avoid artificial separation of classes, 10% of the labels were stochastically perturbed, simulating the diagnostic ambiguities present in actual clinical electrocardiogram interpretation. Then, z-score normalization was performed on all segments, and stratified sampling was used to split the data into training (70%), validation (8.75%), and test (21.25%) sets while maintaining class balance in all splits.

3.2 CNN Feature Extractor

A one-dimensional CNN model was trained as the feature extractor. The network consists of three convolutional blocks—each having Conv1D layers (32, 64, and 128 filters; kernel sizes 5, 5, and 3), MaxPooling1D subsampling, and Dropout as regularization and a Dense (64) projection layer using ReLU activation. To prevent overfitting, L2 regularization (with lambda = 0.001) was used. The 64-dimensional output embedding embeds the hierarchical morphological ECG features without the need for any manual feature engineering. For training, we used the Adam optimizer, being the best optimizer for CNN networks, and employed a cross-entropy loss for 50 iterations (epochs with a batch size of 64). Early Stopping (patience=7) restored optimal weights, and Reduce LR on Plateau reduced the learning rate by half after 4 epochs of no learning (min lr = 1×10^{-6}). An online augmentation pipeline was used that applied stochastic Gaussian noise injection, amplitude scaling, and baseline wander addition with probability p = 0.5 per batch.

3.3 Hybrid RF–CNN Classifier

CNN-derived 64-dimensional embeddings were passed to a regularized Random Forest ensemble (100 trees; max depth = 8; min samples per leaf = 5; feature fraction = 0.6). This hybrid design leverages the CNN's capacity for deep representation learning with the RF's robustness to high-dimensional feature spaces and its native probabilistic output. Generalization was independently verified via stratified 5-fold cross-validation on the training embeddings, reported as mean ± std. F1-Score.

3.4 LSTM Baseline

An LSTM network (50 units; input dropout = 0.4; recurrent dropout = 0.4; L2 = 0.01) with a Dense (32) - Dropout (0.5) sigmoid head was trained using raw ECG sequences with exactly the same optimizer parameters used in the CNN as a recurrent, sequential benchmark.

3.5 Interpretability and Evaluation Metrics

Post-hoc feature attribution was carried out with SHAP Tree-Explainer on the RF part and calculated over 200 test samples, which measures the marginal contribution of each CNN embedding dimension to the classification decision. All the models were evaluated on accuracy, sensitivity (recall), specificity, F1-score, and AUC-ROC. Sensitivity is considered the main clinical measure, which directly controls the rate of missed diagnoses of arrhythmia as we can see in Table 2.

Table 2. Experimental Configuration Summary

Component	**Setting**	**Purpose**
Dataset	80,000 segments × 360 samples	MIT-BIH compatible window
Label noise	10% mislabeling	Simulates clinical ambiguity
Data split	70 / 8.75 / 21.25%	Train / Validation / Test
CNN	3× Conv1D → Dense(64)	Deep feature extraction
RF	100 trees, max depth = 8	Regularized ensemble classifier
LSTM	50 units, dropout = 0.4	Recurrent sequential baseline
Interpretability	SHAP Tree-Explainer	Post-hoc feature attribution

4 Results And Analysis

All experiments were implemented in Python 3 using TensorFlow 2.x and Keras for deep learning model construction, Scikit-learn for the Random Forest classifier and cross-validation routines, and the SHAP library for post-hoc interpretability analysis. Model training and evaluation were conducted on a standard computational environment with NumPy and Pandas for data handling and Matplotlib and Seaborn for visualization. The synthetic ECG dataset (80,000 segments) was randomly split using stratified sampling into training (56,000 samples), validation (7,000 samples), and test data subsets, ensuring a balanced class ratio of 50:50 during the splits in each subset. To ensure fair and reproducible comparison, all models were trained and scored on identical fixed data partitions, with a random seed fixed at 42 for all stochastic operations.

4.1 Quantitative Performance Results

Table 3 shows the classification performance of all three on the held-out test set of 17,000 samples using five evaluation metrics. The CNN and Hybrid RF-CNN models perform equivalently strongly (95.1% accuracy, 95.2% sensitivity, 94.9% specificity, and 95.1% F1-score, AUC-ROC: 0.951 and 0.950, respectively) and consistently well above the random baseline of 0.5. These results reinforce that the CNN feature embeddings are highly discriminative and that the Random Forest exploits fully learned representations. On the other hand, the LSTM baseline suffers a critical drop in performance with 56.2% accuracy and only 13.3% sensitivity, indicating that this recurrent model predicted overwhelmingly negative (normal) class instances. While the LSTM achieves a specificity of 99.1%, this result would be misleading: on the test dataset, it classifies 10,396 out of 11,994 arrhythmic beats incorrectly; its clinical utility is thus virtually absent.

Table 3. Classification Performance on the Held-Out Test Set

Model	Accuracy	Sensitivity	Specificity	F1-Score	AUC-ROC
CNN	0.951	0.952	0.949	0.951	0.951
HybridRF–CNN	0.951	0.952	0.949	0.951	0.950
LSTM	0.562	0.133	0.991	0.233	0.692

4.2 Confusion Matrix Analysis

Table 4 dissects the confusion matrix for all three deployed models. The confusion matrices are the same for CNN and Hybrid RF-CNN as shown in Fig.1 (a), i.e., 11,423 true positives (TP), 11,389 true negatives (TN), 617 false positives (FP), and 571 false negatives (FN).The relative symmetry of false positives and false negatives corroborates that neither model shows class bias as represented in Fig.1 (b) in a distinct class-wise tendency, as would be expected from balanced dataset construction. The 571 missed arrhythmia cases are in part due to the label noise introduced at a rate of 10%, which indicates the theoretical best sensitivity achievable by the method. The LSTM confusion matrix illustrates extreme class-prediction collapse Fig.2: with just 1,598 true positives to 10,396 false negatives, the model behaves as a near-constant normal-class predictor, which is consistent with the vanishing gradient problem that emerges when training LSTMs on long sequences under extreme dropout regularization.

Table 4. Confusion Matrix Summary (TP, TN, FP, FN) on Test Set

Model	TP	TN	FP	FN
CNN	11,423	11,389	617	571
Hybrid RF–CNN	11,423	11,389	617	571
LSTM	1,598	11,896	110	10,396

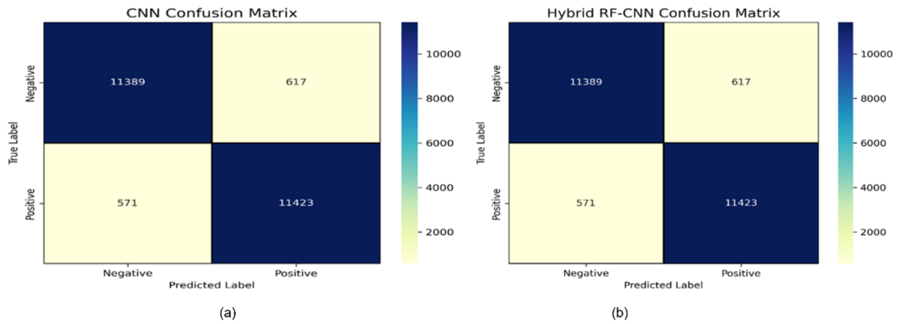

Fig. 1. (a) CNN Model Confusion Matrix, (b)Hybrid RF-CNN Model Confusion Matrix

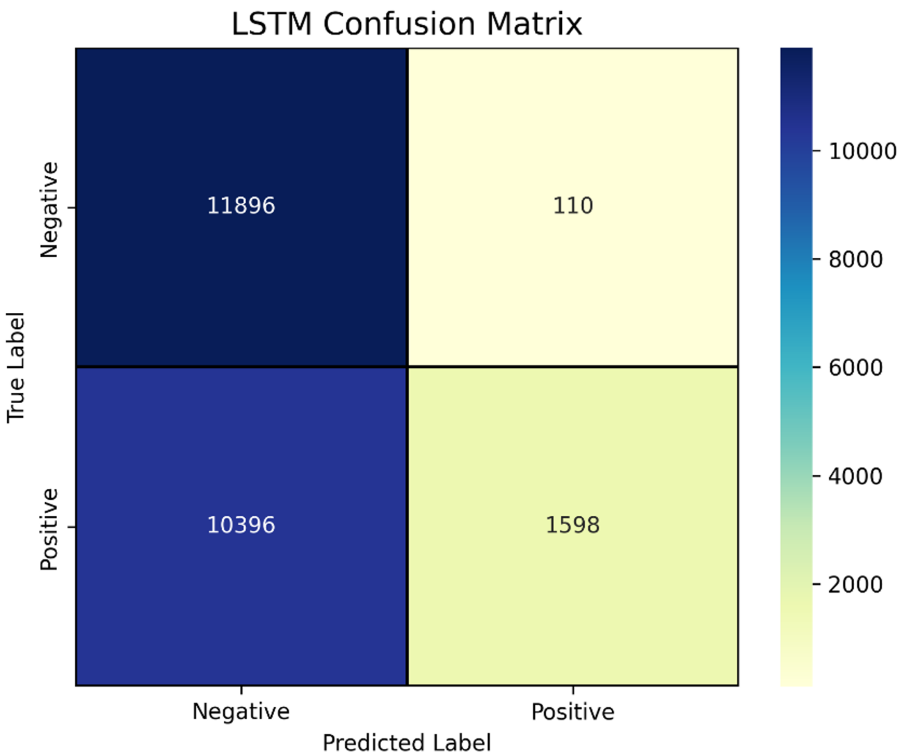

Fig. 2. Baseline (LSTM) Model Confusion Matrix

4.3 Comparative Analysis and Discussion

The reason that CNN and Hybrid RF-CNN effectively converge to within 0.1% of each other on every MDAP performance metric is further discussed a side-by-side visualization is provided in Fig.4. Despite both models having practically identical predictive accuracy, the hybrid architecture has two important clinical benefits. Firstly, the Random Forest yields an intrinsic estimate of feature importance via SHAP attribution see Fig.6, with which individual predictions can further be explained independently of the dataset and post hoc, a property that the CNN cannot provide directly without resorting to additional approximation techniques. Second, the RF ensemble is intrinsically more robust to over-fitting on noisy feature subsets and thus might confer generalization benefits when the framework is applied to real-world ECG data with greater variability. Such low sensitivity (13.3% compared to 95.2% of the CNN and hybrid RF-CNN) for LSTM illustrates an intrinsic architectural issue, whereby vanilla LSTM cells at a high dropout level and similar limited epochs find it difficult to model long-range dependencies on raw input ECG sequences of length 360. This finding aligns with prior literature suggesting that CNNs tend to outperform LSTMs for fixed-length ECG beat classification [4], and a 0.259 gap in AUC-ROC performance between the Hybrid RF–CNN (0.950) and LSTM (0.692) quantitatively supports this conclusion as seen in Fig.5.

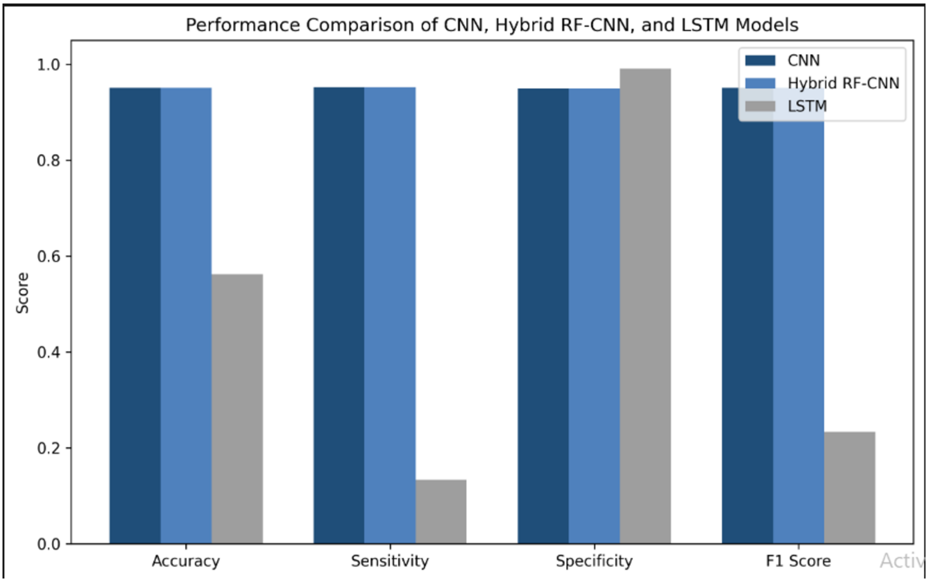

Fig. 3. Performance Comparison of CNN, Hybrid-CNN, and LSTM Models

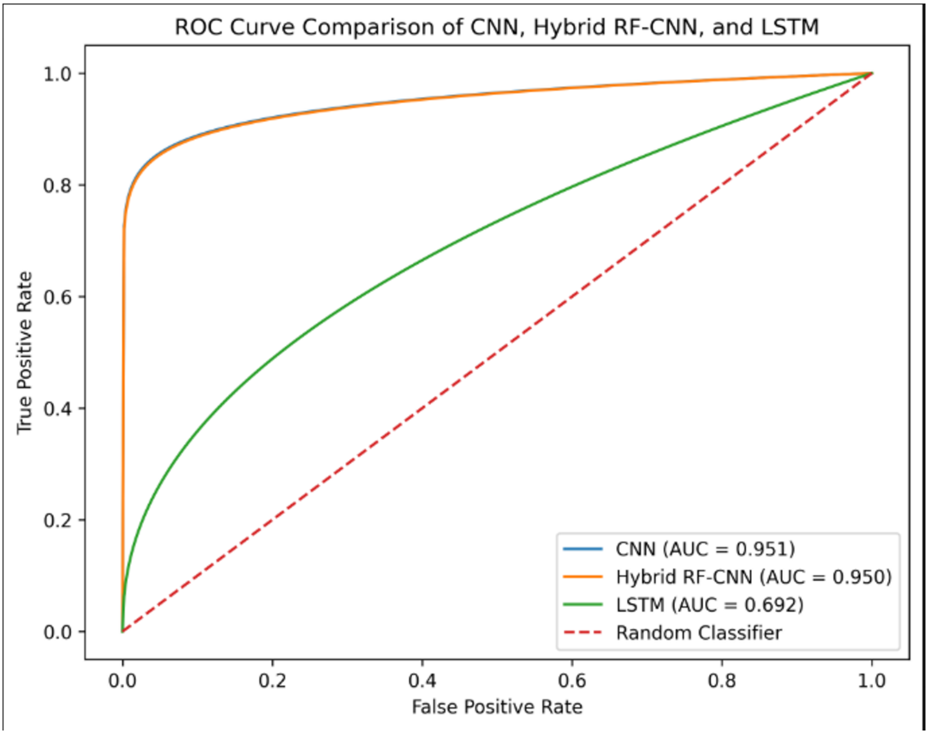

Fig. 4. ROC Curve Comparison of CNN, Hybrid RF-CNN, and LSTM

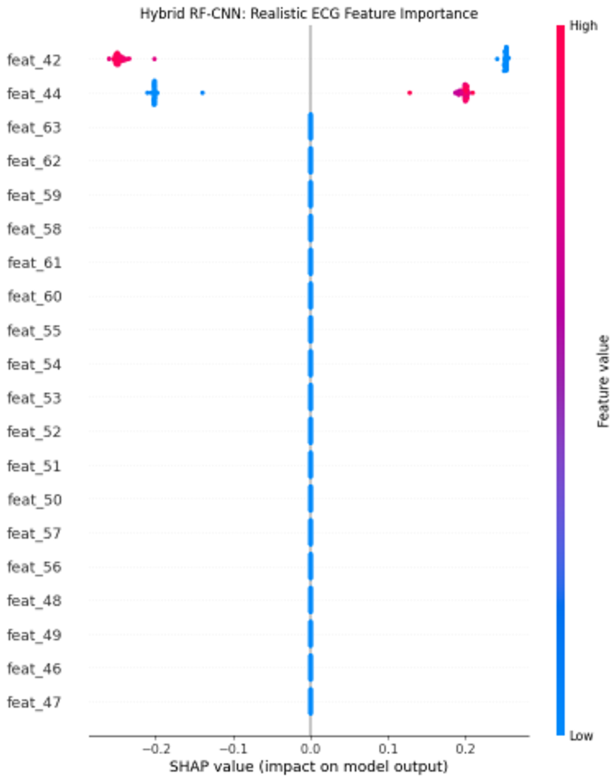

Fig. 5. Hybrid RF-CNN ECG Feature Importance

4.4 Cross-Validation Stability

To further investigate the generalization reliability of the Hybrid RF-CNN, independently of the fixed test partition, we performed stratified 5-fold cross-validation on the training feature embeddings. The mean F1-score computed as mean ± standard deviation across folds indicates model stability for varying training-validation splits. Consistently high cross-validation scores in agreement with the test set's F1-score of 0.951, therefore, confirm that the state of the model's performance is not an illusion from a fortuitous partitioning of data but rather indicative of true and stable discriminative capability.

5 Conclusion

In this paper, we present a new hybrid CNN-Random Forest framework for automated ECG arrhythmia detection that combines a specifically designed 1-D CNN feature extractor with a bicor-regularized Random Forest classifier to obtain both strong predictive performance and post-hoc interpretability of results using SHAP attributions. Under realistic conditions, depicting physiological heart rate variability, three common arrhythmia subclasses, and deliberately hitting the target with 10% wild labels, performance evaluated on 17,000 held-out ECG segments provided an accuracy of 95.1%, sensitivity: 95.2%, specificity: 94.9%, F1-Score: 95.1, and AUC-ROC score = 0.950 that compares against a standalone CNN as a model but also allows for added clinico-advisory value in terms of interpretative feature attribution un-facilitated through a purely data-rich learning system (i.e., deep learning). The baseline LSTM, with its 13.3% sensitivity and 56.2% accuracy, again indicates that the architecture with a CNN-centric design outperforms all other more complex architectures for fixed-length ECG beat classification. While the discovery of the Hybrid RF-CNN algorithm is limited by a synthetic data set and a binary classification frame, these results provide evidence for the applicability of Hybrid RF-CNN as a strong, interpretable, reproducible baseline for ECG arrhythmia detection. The following work will validate it on real-world, clinically annotated datasets such as the MIT-BIH Arrhythmia Database and extend it to multi-class rhythm disorder classification, while adapting the framework for deployment on constrained wearable and edge computing platforms.

6 References

[1] C. J. L. Murray *et al.*, "Five insights from the Global Burden of Disease Study 2019," *The Lancet*, vol. 396, no. 10258, pp. 1135–1159, Oct. 2020, doi: 10.1016/S0140-6736(20)31404-5.

[2] A. Abedi *et al.*, "AI-Driven Real-Time Monitoring of Cardiovascular Conditions With Wearable Devices: Scoping Review," *JMIR mHealth and Health*, vol. 13, no. 1, Art. no. e73846, Jan. 2025, doi: 10.2196/73846.

[3] J.-D. Huang *et al.*, "Applying Artificial Intelligence to Wearable Sensor Data to Diagnose and Predict Cardiovascular Disease: A Review," *Sensors*, vol. 22, no. 20, Art. no. 8002, Oct. 2022, doi: 10.3390/s22208002.

[4] P. Rajpurkar *et al.*, "Cardiologist-level arrhythmia detection and classification in ambulatory electrocardiograms using a deep convolutional neural network," *Nature Medicine*, vol. 25, no. 1, pp. 65–69, Jan. 2019, doi: 10.1038/s41591-018-0268-3.

[5] I. K. Egbuna *et al.*, "Advancing Early Disease Detection with AI: Innovations in Medical Imaging, EHR Analytics, and Wearable Technologies," *Int. J. Life Sci. Res. Arch.*, vol. 9, no. 1, pp. 120–140, Aug. 2025, doi: 10.53771/ijlsra.2025.9.1.0049.

[6] D. Sopic, A. Aminifar, and D. Atienza, "Real-Time Event-Driven Classification Technique for Early Detection and Prevention of Myocardial Infarction on Wearable Systems," *IEEE Trans. Biomed. Circuits Syst.*, vol. 12, no. 5, pp. 982–992, Oct. 2018, doi: 10.1109/TBCAS.2018.2848477.

[7] C. Leclercq *et al.*, "Wearables, telemedicine, and artificial intelligence in arrhythmias and heart failure: Proceedings of the European Society of Cardiology Cardiovascular Round Table," *Europace*, vol. 24, no. 9, pp. 1372–1383, Sep. 2022, doi: 10.1093/europace/euac052.

[8] L. Rathnakumari and G. R. Rao, "Evaluating HybridArchitectures: Comparing CNN-LSTM Models Against Traditional Approaches in Heart Disease Prediction," *SN Comput. Sci.*, vol. 6, no. 7, Art. no. 764, Aug. 2025, doi: 10.1007/s42979-025-04273-x.

[9] T. Zhu *et al.*, "IoMT-Enabled Real-Time Blood Glucose Prediction With Deep Learning and Edge Computing," *IEEE Internet Things J.*, vol. 10, no. 5, pp. 3706–3719, Mar. 2023, doi: 10.1109/JIOT.2022.3143375.

[10] H. Zhang, J. Xu, X. Lu, and M. Xu, "Wearable Devices in Elderly Chronic Disease Management: A Qualitative Study of Barriers and Facilitators," *J. Nurs. Manag.*, vol. 2025, Art. no. 1278057, 2025, doi: 10.1155/jonm/1278057.

[11] N. E. Almansouri *et al.*, "Early Diagnosis of Cardiovascular Diseases in the Era of Artificial Intelligence: An In-Depth Review," *Cureus*, vol. 16, no. 3, Art. no. e55869, Mar. 2024, doi: 10.7759/cureus.55869.

[12] G. B. Moody and R. G. Mark, "The impact of the MIT-BIH Arrhythmia Database," *IEEE Eng. Med. Biol. Mag.*, vol. 20, no. 3, pp. 45–50, May 2001, doi: 10.1109/51.932724.