Edge-Deployable Deep Neural Network for Atrial Fibrillation Detection in Ambulatory ECG

Atitaya Phoemsuk, Fatemeh Alidostdargah, Vahid Abolghasemi, Abdolrahman Peimankar, Devender Kumar, Helena Dominguez, Sadasivan Puthusserypady

***Abstract*—Ambulatory electrocardiogram (ECG) monitoring is essential for detecting long-term atrial fibrillation (AFib), but it remains challenging due to severe noise, motion artefacts, and signal non-stationarity. This paper proposes a lightweight deep learning architecture for AFib classification from ambulatory ECG, designed for efficient deployment on resource-constrained hardware. The model combines one-dimensional convolutional encoding, residual feature refinement, and a compact self-attention mechanism to capture both local ECG morphology and long-range temporal rhythm irregularities. Global average pooling produces a compact representation for low-complexity classification. The proposed method is evaluated on the Contextualized Ambulatory Electrocardiography Arrhythmia Dataset (CACHET-CADB), a long-term real-world ambulatory ECG dataset with expert annotations and naturally occurring noise. Experimental results demonstrate an average classification accuracy of 97.31%, outperforming existing deep learning-based approaches while maintaining low computational complexity, with a model size of only 2.67 MB, making it suitable for edge and wearable deployment.**

***Index Terms*—Atrial Fibrillation, Ambulatory electrocardiogram signals, Edge AI, Lightweight deep learning**

I. Introduction

Atrial fibrillation (AFib) is the most common sustained cardiac arrhythmia characterised by disorganised electrical activity in the atria, leading to an irregular and often rapid heart rhythm [1]. This abnormal atrial activity compromises effective cardiac contraction and blood flow, thereby increasing the risk of thromboembolic events, including stroke. Electrocardiogram (ECG) remains the primary pre-screening tool for diagnosis of AFib and other cardiovascular diseases, as it provides direct information on the cardiac electrical activity and enables the identification of characteristic rhythm abnormalities that are reflective of cardiac disorders. AFib may cause irregular heart rate, shortness of breath, and chest pain. Automated ECG-based AFib diagnosis enables early detection, supports long-term health monitoring, and reduces clinician workload, whilst facilitating the prioritisation of patients for timely clinical assessment and treatment, thereby contributing to the prevention of adverse cardiac events, including sudden cardiac arrest, in large-scale screening and resource-constrained healthcare settings.

With advances in artificial Intelligence (AI), automated AFib diagnosis has been widely used as an AI-assisted tool for ECG-based pre-screening and long-term monitoring. These approaches enable the detection of both symptomatic and asymptomatic AFib episodes, supporting clinical treatment and post-treatment follow-up through the timely identification of AFib recurrence. Traditional Machine Learning (ML) methods for AFib detection typically depend on handcrafted feature extraction from ECG signals, including time-domain, frequency-domain, and morphological characteristics. In contrast, Deep Learning (DL) methods perform end-to-end classification by automatically learning discriminative temporal features directly from raw ECG signals, reducing the need for manual feature extraction and improving scalability for large-scale screening and continuous monitoring. Various ML-based AFib diagnosis approaches have been explored, including Support Vector Machines (SVM), Random Forests (RF), Decision Trees and XGBoost [2].

DL-based approaches are increasingly preferred over traditional ML-based models because they can learn directly from raw ECG signals. Recent DL-based ECG detection methods have focused on convolutional neural network (CNN) architectures due to their effective balance between classification performance and computational efficiency [3]–[8]. In addition to CNN-based models, hybrid DL architectures combinithat combine CNNs with long short-term memory (LSTM) networks [9], [10] and attention-based mechanisms [11] have been explored for AFib detection to capture both local morphological characteristics and temporal dependencies ingnals. Several lightweight DL-based architectures have been proposed to improve efficiency without compromising performance. Compact models for ECG tasks such as R-peak detection demonstrate that accurate temporal feature extraction can be achieved with substantially reduced complexity [12], [13]. Lightweight classification frameworks have also been explored, including a cascade CNN–BiGRU model for single-lead and 12-lead ECG classification [14] and CardioNet for myocardial infarction detection using time–frequency representations [15]. However, such design choices increase computational overhead and limit real-time deployment on embedded devices.

Most existing approaches rely on clean ECG segments collected in controlled clinical environments rather than real-world ambulatory ECG, which contains motion artefacts, noise, and posture changes [2]. This reliance on clinically clean data can limit generalisation to free-living recordings

A. Phoemsuk, F. Alidostdargah, and V. Abolghasemi are with the School of Computer Science & Electronic Engineering, University of Essex, CO4 3SQ, UK emails: {ap19698, fa23653, v.abolghasemi}@essex.ac.uk

A. Peimankar, and D. Kumar are with the Maersk Mc-Kinney Moller Institute, Faculty of Engineering, University of Southern Denmark, Campusvej 55, 5230 Odense M, Denmark.

H. Dominguez is with Bispebjerg and Frederisksberg Hospital, Bispebjerg Bakke 23, DK-2400 Copenhagen, Denmark. She is also with Department of Biomedical Sciences, Faculty of Health and Medical Sciences, University of Copenhagen, Blegdamsvej 3b, Panum Building 12.4, DK-2200, Denmark.

S. Puthusserypady is with the Department of Health Technology, Technical University of Denmark, Kongens Lyngby, 2800, Denmark.

with lower signal quality and increased artefacts. As a result, model robustness under ambulatory recording conditions and its potential for deployment in resource-constrained environments remain unexplored. Furthermore, deployment in resource-constrained or mobile environments demands models that are not only robust to noise but also computationally lightweight.

To address these gaps, we propose a novel lightweight deep learning architecture for accurate AFib classification in noisy single-lead ambulatory ECG data. Our model balances high diagnostic performance with low computational complexity, making it suitable for wearable and point-of-care deployment. The architecture employs hierarchical feature learning with convolutional encoding, residual connections, and a compact self-attention block. Our contributions are threefold. First, we explore AFib detection on the challenging Contextualized Ambulatory Electrocardiography Arrhythmia Dataset (CACHET-CADB)—an underexplored dataset reflecting realistic, noisy monitoring conditions. While CACHET-CADB has been utilised in recent studies [16]–[20], these typically rely on extensive preprocessing and feature engineering. In contrast, our approach processes raw ECG signals with a compact end-to-end architecture. Second, we demonstrate strong and consistent performance in both binary (AFib vs. Normal) and 3-class (AFib vs. Normal vs. Noise) classification tasks. Third, we achieve this with a highly compact architecture, underscoring its practicality for resource-constrained edge and wearable devices.

II. Methodology

A. Data preparation

The CACHET-CADB dataset comprises 259 days of contextualised ambulatory single-lead ECG recordings from 24 subjects, including 1,602 manually annotated 10-s rhythm segments [21]. Signals were sampled at 1024 Hz, with recording durations ranging from 24 h to 3 weeks. Although the dataset includes contextual information (physical activity, acceleration, posture), this study uses only ECG signals, which are provided as pre-segmented ambulatory recordings preserving realistic noise characteristics.

The dataset includes 747 AFib, 615 NSR (Normal Sinus Rhythm), 221 noise, and 19 other non-AFib arrhythmia segments, all of which were used. Noise segments correspond to intervals with severe signal corruption (e.g., motion artefacts and poor electrode contact). Representative AFib, NSR, and noise examples are shown in Fig. 1, illustrating the irregular rhythm and absent P-waves in AFib, regular rhythm with clear P-waves in NSR, and high-frequency, low-amplitude artefacts in noise segments.

B. Proposed Architecture

Figure 2 illustrates the overall framework of the proposed method. The convolutional (CONV) encoding block, as the initial stage of the proposed method, is designed to extract local ECG features that support the discrimination of AFib, NSR, and noise. A one-dimensional convolutional layer with 50 filters and a kernel size of 28 captures short-term temporal patterns, including irregular rhythm characteristics associated

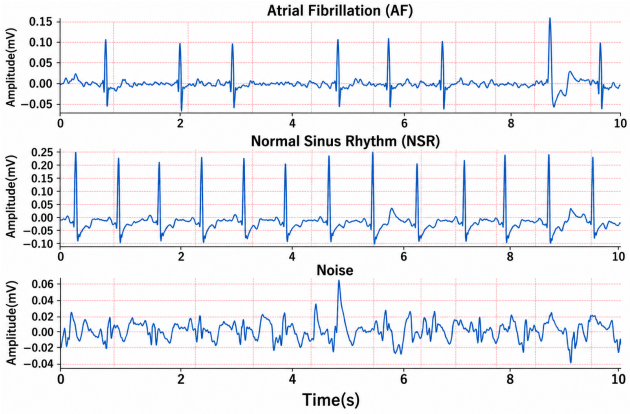

Fig. 1. Examples of 10-second ECG segments, including AFib (top), NSR (middle), and Noise (bottom) from the CACHET-CADB dataset [21].

with AFib, stable waveform structures in normal ECG, and non-physiological patterns caused by noise and artefacts. A Rectified Linear Unit (ReLU) activation function is applied to enhance feature discrimination. Batch normalisation is used to stabilise feature distributions and improve training robustness when handling noisy ambulatory ECG recordings. A dropout layer with a rate of 0.1 is incorporated to reduce overfitting and encourage the learning of shared and robust feature representations across different classification settings. Finally, max-pooling reduces temporal resolution while preserving the most informative features relevant to AFib, normal, and noise classification.

A residual CONV block is employed to further refine the extracted ECG features while preserving information from earlier layers. The block consists of two sequential one-dimensional convolutional layers with 128 filters and a kernel size of 16. Each convolution is followed by batch normalisation to stabilise feature distributions during training. Residual learning is achieved through a shortcut connection that performs an identity mapping, with a 1×1 convolution used when required to match feature dimensions before element-wise addition. The residual mapping can be expressed as shown in (1).

$$\mathbf{y} = \mathcal{F}(\mathbf{x}, \{W_i\}) + \mathbf{x}, \tag{1}$$

where $\mathbf{x}$ denotes the input ECG signal, $\mathcal{F}(\mathbf{x}, \{W_i\})$ represents the output of the convolutional layers parameterised by the learned weights $\{W_i\}$, and $\mathbf{y}$ corresponds to the output of the residual block. After residual feature fusion, a ReLU activation function is applied, followed by max-pooling to reduce temporal resolution and computational complexity. This residual block is repeated twice to progressively enhance feature abstraction.

Then, a self-attention encoder is introduced to capture long-range temporal dependencies in the ECG signal by modelling relationships between temporally distant segments across multiple cardiac cycles. The self-attention encoder begins with layer normalisation, followed by a multi-head self-attention mechanism with four attention heads and a feature dimension of 64. The self-attention operation is defined as shown in (2).

$$\mathbf{Attention}(Q, K, V) = \mathbf{softmax}\left(\frac{QK^T}{\sqrt{d_k}}\right)V, \tag{2}$$

where $Q = W_Q x$, $K = W_K x$, and $V = W_V x$ denote the query, key, and value projections from the ECG input x, with W_Q, W_K, and W_V as learnable weight matrices. The

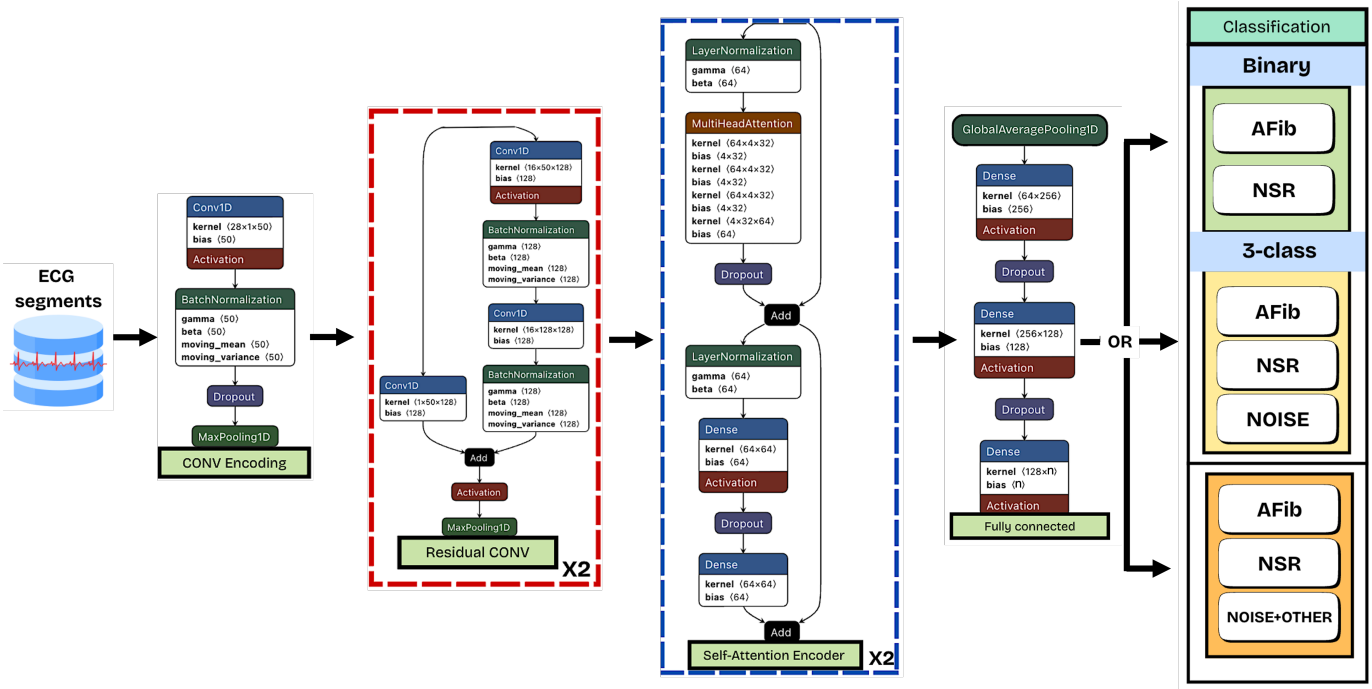

Fig. 2. Illustration of the proposed model for ambulatory ECG analysis. CONV and residual CONV blocks capture local ECG morphology, while a self-attention encoder handles long-range temporal dependencies and signal variability. Global average pooling aggregates the learned temporal features into a compact representation, which is then passed to a low-complexity fully connected classifier for binary or three-class ECG classification.

scaling factor $\sqrt{d_k}$ stabilizes gradients during training. Using four attention heads on 32-dimensional subspaces of the 64-dimensional input, the multi-head mechanism captures diverse temporal dependencies, enabling the model to attend to AFib rhythm patterns, NSR structures, and noise artifacts. Dropout is applied to the attention output, followed by a residual connection with x and layer normalization. The normalized features pass through a position-wise feed-forward network of two dense layers (64 dimensions, ReLU), with dropout after the first layer and a residual connection after the second. Two stacked self-attention encoder blocks progressively refine long-range temporal dependencies for discriminating AFib, NSR, and noise.

A one-dimensional global average pooling operation aggregates temporal features into a fixed-length vector, capturing rhythm irregularity across multiple cycles. The pooled feature is then processed by a classification head: a dense layer (64 $\rightarrow$ 256, ReLU, dropout), a second dense layer (256 $\rightarrow$ 128, ReLU, dropout), and a final output layer projecting to n classes, supporting binary or 3-class AFib/NSR/noise classification. This transforms high-level temporal representations into discriminative decisions for robust AFib detection under noisy ambulatory conditions.

III. Experimental Results and Discussion

All experiments were implemented in Python and executed on Google Colab Pro with GPU acceleration. An allocation of up to 100 compute units was used. The proposed model architecture and experimental configuration were maintained across binary and 3-class ECG classification tasks, with only the final output layer modified to the number of classes.

The dataset was partitioned into training and test sets, where 70% of the data was used for model development, and the remaining 30% was held out as an unseen test set. Model evaluation was then performed on the training portion using stratified 10-fold cross-validation. Within each fold, 10% of the training data was reserved for validation. The model was trained for 20 epochs using a batch size of 16, the Adam optimiser with a learning rate of 0.00025, and categorical cross-entropy loss. These hyperparameters were selected through a random search optimisation strategy.

To evaluate the model, performance metrics including *Accuracy*, *Precision*, *Recall*, and *F1-Score* were used to measure the performance of the binary classification model via:

$$\begin{aligned} \text{Accuracy} &= \frac{\text{TP}+\text{TN}}{\text{TP}+\text{TN}+\text{FP}+\text{FN}}, \quad \text{Precision} = \frac{\text{TP}}{\text{TP}+\text{FP}}, \\ \text{Recall} &= \frac{\text{TP}}{\text{TP}+\text{FN}}, \quad \text{F1-Score} = 2 \times \frac{\text{Precision} \times \text{Recall}}{\text{Precision}+\text{Recall}}. \end{aligned} \tag{3}$$

where TP, TN, FP, and FN denote correctly identified AFib, correctly identified NSR, NSR misclassified as AFib, and AFib misclassified as NSR, respectively. For 3-class classification, precision, recall, and F1-score are computed per class.

Table I reports the performance of the proposed model on the unseen test set. The model demonstrates strong and consistent performance across both classification settings. For the binary classification, it achieves an accuracy of 97.31% with an F1-score of 97.29%, demonstrating consistent discrimination between AFib and NSR. For the 3-class classification setting, which distinguishes AFib, NSR, and noise, the model achieves an accuracy of 94.11%. When the other arrhythmia and noise classes were merged into a single category (due to the very small number of segments in the "Other" class), the model achieved an accuracy of 93.14%, indicating that performance remained stable despite the increased variability within the merged class. Overall, the precision and recall are

well balanced, indicating reliable discrimination across the three classes.

TABLE I
PERFORMANCE OF THE PROPOSED MODEL EVALUATED ON THE UNSEEN TEST SET FOR BINARY AND 3-CLASS CLASSIFICATION.

Classification	**Acc (%)**	***F1****	***Precision****	***Recall****
Binary	97.31	97.29	97.21	97.40
3-class (AFib,NSR,Noise)	94.11	94.41	94.04	94.93
3-class (AFib,NSR,(Noise+Other))	93.14	93.37	93.62	93.24

* Precision, recall, and F1 are reported as macro averages.

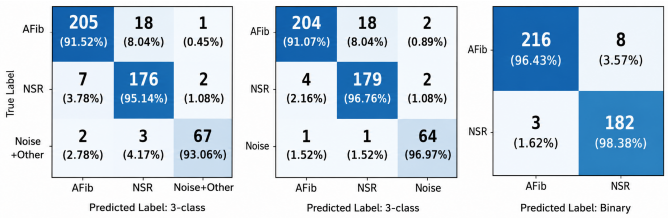

Fig. 3. Proposed method: Confusion matrices evaluated on the unseen data.

Figure 3 illustrates the confusion matrices of the proposed model evaluated on the unseen test set for both binary and 3-class classification. In the binary setting shown in Figure 3 (right), 216 out of 224 AFib segments are correctly classified, corresponding to a correct classification rate of 96.43%, while 8 segments, or 3.57%, are misclassified as NSR. For NSR, 182 out of 185 ECG segments are correctly identified, achieving a classification rate of 98.38%, with only 3 segments, or 1.62%, incorrectly identified as AFib. In both 3-class settings, shown in Figure 3 (left and middle) , similar consistent performance is observed despite some degree of imbalance within classes.

Furthermore, merging the Other arrhythmia class with Noise reduces the impact of severe class imbalance, improves the stability of performance metrics, and better reflects realistic ambulatory ECG conditions. The observation that most AFib misclassifications are assigned to NSR rather than noise indicates that these errors are associated with similarities in rhythm characteristics rather than noise-related effects. Furthermore, the high correct classification rate for noisy segments indicates that the proposed model can effectively separate poor-quality ECG signals, supporting its applicability to real-world monitoring scenarios.

Table II provides an ablation study to evaluate the effectiveness of different developments leading to the proposed architecture. The baseline model [11] is an attention-driven deep network used for ECG signals classification. However, it is considerably large due to the use of a Flatten layer followed by fully connected layers, resulting in 14.58M parameters and a model size of 55.8 MB. As seen from Table II, the baseline model achieved an accuracy of 71% in the binary classification setting and 66% in the 3-class setting, which is not promising. However, the proposed model, which integrates CONV encoding, residual CONV, and a self-attention encoder, achieved the best overall trade-off between performance and efficiency (Table II). It obtained accuracies of 97.31% and 94.74% for binary and 3-class classification, respectively, while maintaining a compact model size at around 2.6 MB. Compared with the baseline, this represents a substantial reduction in model complexity alongside improved classification performance, confirming the suitability of the proposed architecture for deployment in resource-constrained and real-time ECG monitoring scenarios.

Table III compares the proposed model with classical machine learning and deep learning baselines evaluated under the same protocol as mentioned at the beginning of this section. The traditional classifiers, including SVM, KNN, and logistic regression, show limited performance, with accuracies ranging from 48.66–55.26% for binary classification and 38.11–48.63% for 3-class classification. They also store substantial portions of the training data, which accounts for their unexpectedly large size. Three deep learning baselines were evaluated: (i) a simple 1D CNN adapted from [22], (ii) a deep bidirectional LSTM (DBLSTM) inspired by [10], and (iii) a hybrid CNN–LSTM following the design in [9]. Among the baseline models, DBLSTM achieved the strongest performance (62.35% accuracy) with the smallest model size, while the 1D CNN yielded moderate results (58.19%), suggesting that modeling temporal dependencies in ECG signals improves class discrimination. The 1D CNN's large footprint stems from its Flatten layer, which transforms feature maps into a high-dimensional vector and leads to a parameter-heavy fully connected layer. In contrast, DBLSTM remains compact due to its use of only 64 units and global pooling, resulting in far fewer parameters and a smaller saved model size. The proposed method, however, delivers markedly superior performance while requiring just 2.671 MB of storage, achieving 97.31% accuracy for binary classification and 94.11% for the 3-class task. This substantial performance gap underscores the effectiveness of the proposed architecture, which learns both discriminative morphological features and temporal dependencies in ECG signals, enabling more reliable classification across different class configurations.

Finally, our model was benchmarked against the most comparable existing work, FHO-BiLSTM [16], which was also trained and evaluated on the CACHET-CADB dataset. While FHO-BiLSTM reported 96.74% accuracy and 94.44% F1-score on unseen data, our method achieves a competitive 97.31% accuracy and 97.29% F1-score. Crucially, our model offers the additional advantage of a simpler end-to-end architecture that processes raw ECG signals directly, eliminating the need for the extensive handcrafted feature extraction pipeline required by FHO-BiLSTM.

IV. CONCLUSION

This study presents a lightweight hybrid model for robust atrial fibrillation detection using single-lead ECG. The model achieves 97.31% accuracy in binary classification and 94.11% in the 3-class task (AFib, normal sinus rhythm, noise) with a compact footprint of 2.67 MB—a 22-fold compression over a dense-layer baseline. Our results confirm robust performance in noisy ambulatory conditions, supporting the feasibility of deploying the proposed lightweight model for screening and continuous monitoring on edge devices in resource-constrained environments. Future integration of contextual data, such as

TABLE II
ABLATION STUDY: COMPARING BASELINE MODEL WITH PROPOSED MODEL USING DIFFERENT LAYER CONFIGURATIONS FOR BINARY AND 3-CLASS.

Architecture	Binary			3-class		
	Acc (%)	#Params	Size (MB)	Acc (%)	#Params	Size (MB)
Baseline model [11]	71.00	14,582,388	55.796	66.00	14,582,517	55.797
CONV + 2 Residual CONV	93.64	531,002	2.120	95.79	531,131	2.121
CONV + 2 Residual CONV + 1 Self-Attention	89.73	572,800	2.317	89.68	572,929	2.317
CONV + 2 Self-Attention	61.12	145,822	0.675	65.68	145,951	0.675
CONV + 1 Residual CONV + 2 Self-Attention	94.38	688,838	2.769	89.05	688,967	2.769
Proposed model	97.31	656,006	2.671	94.74	656,135	2.672

* Precision, recall, and F1 are reported as macro averages.

TABLE III
PERFORMANCE COMPARISON OF SEVERAL MODELS. MODEL SIZE REPRESENTS THE STORAGE REQUIRED FOR THE TRAINED MODEL. ML METHODS MAY STORE SUBSETS OR ALL OF THE TRAINING DATA. DL MODELS STORE TRAINABLE PARAMETERS.

Architecture	Binary					3-class				
	Acc (%)	*Precision**	*Recall**	*F1**	*Size (MB)*	Acc (%)	*Precision**	*Recall**	*F1**	*Size (MB)*
SVM	55.26	63.13	58.26	52.03	73.073	48.63	68.04	47.53	42.44	86.516
KNN	50.12	45.66	47.22	43.18	37.473	41.68	36.65	41.26	35.94	43.529
Logistic Regression	48.66	48.43	48.42	48,40	0.314	38.11	35.76	37.43	35.74	0.470
1D CNN [22]	58.19	57.93	55.67	53.49	319.99	58.74	66.49	50.68	52.36	319.99
DBLSTM [10]	62.35	63.87	59.84	58.04	0.570	70.32	73.24	71.40	72.23	0.570
CNN–LSTM [9]	54.28	39.78	49.60	35.66	80.468	40.84	53.35	37.52	26.56	80.469
Proposed model	97.31	97.21	97.40	97.29	2.671	94.11	94.04	94.93	94.41	2.672

* Precision, recall, and F1 are reported as macro averages.

patient activity or medication records, could further improve robustness by distinguishing physiological variations from pathological rhythms. Challenges include data synchronization and privacy, but the model's efficiency enables deployment on wearable patches or edge devices. Future work will focus on external validation across diverse ambulatory datasets and real-time implementation on embedded hardware.

REFERENCES

[1] S. Yang, N. De Kruijf, F. Zhu, J. A. E. van Oortmerssen, N. M. S. de Groot, and M. Kavousi, "Trends in burden of atrial fibrillation over three decades: a population-based study," *Heart*, 2025.

[2] S. Islam, M. R. Islam, M. A. Abedin, T. Dökeroğlu, and M. Rahman, "Recent advances in the tools and techniques for AI-aided diagnosis of atrial fibrillation," *Biophys. Rev.*, vol. 6, no. 1, 2025.

[3] J. Lei, Y. Zhou, X. Tian, Q. Zhao, Q. Zhang, S. Geng, Q. Wu, and S. Hong, "A deep learning method for beat-level risk analysis and interpretation of atrial fibrillation patients during sinus rhythm," *Biomed. Signal Process. Control*, vol. 100, p. 107028, 2025.

[4] A. Khan, M. R. Mughal, S. A. Irtaza, M. Khan, M. Tahir, A. Ali, and Z. Saeed, "A deep learning based ultra-lightweight architecture for atrial fibrillation detection using single-lead ECG recordings," *IEEE Access*, 2025.

[5] X. An, T. Xia, J. Hu, H. Bo, L. Cheng, and Y. Ma, "Mini_ECGNet: a lightweight neural network with adaptive covariance pooling and dilated convolution for noise-robust ecg classification," *Biomed. Signal Process. Control*, vol. 113, p. 109232, 2026.

[6] A. Phoemsuk and V. Abolghasemi, "CADNet: A lightweight neural network for coronary artery disease classification using electrocardiogram signals," *IEEE J. Biomed. Health Inform.*, 2025.

[7] ——, "Coronary artery disease classification using one-dimensional convolutional neural network," in *IEEE Conf. Artif. Intell. (CAI)*, 2024, pp. 389–394.

[8] ——, "Enhanced coronary artery disease classification through feature engineering and one-dimensional convolutional neural network," *IEEE Access*, 2025.

[9] G. Petmezas, K. Haris, L. Stefanopoulos, V. Kilintzis, A. Tzavelis, J. A. Rogers, A. K. Katsaggelos, and N. Maglaveras, "Automated atrial fibrillation detection using a hybrid CNN-LSTM network on imbalanced ECG datasets," *Biomed. Signal Process. Control*, vol. 63, p. 102194, 2021.

[10] Ö. Yildirim, "A novel wavelet sequence based on deep bidirectional LSTM network model for ecg signal classification," *Comput. Biol. Med.*, vol. 96, pp. 189–202, 2018.

[11] A. Phoemsuk and V. Abolghasemi, "Multi-disease cardiovascular detection from ECG signals using an attention-driven deep network," in *47th Annu. Int. Conf. IEEE Eng. Med. Biol. Soc. (EMBC)*, 2025, pp. 1–7.

[12] B. A. Dong, T. H. Dat, and N. H. A. Vy, "Harnessing TinyML for accurate ECG beat detection," in *10th IEEE Int. Conf. Integr. Circuits, Design, Verification (ICDV)*, 2025, pp. 13–18.

[13] J. Yao, Y. Zhang, and C. Dong, "A lightweight R peak detection algorithm for noisy ECG signals," in *22nd IEEE Int. Conf. Bioinform. Bioeng. (BIBE)*, 2022, pp. 255–260.

[14] J. Li, L. Huang, Y. Zhai, S. Ling, H. Ouyang, and L. Li, "MFF-IBL: A lightweight cascade network based on multibranch feature fusion and incremental broad learning for lead-independent ECG classification," *IEEE Trans. Instrum. Meas.*, vol. 74, pp. 1–15, 2025.

[15] K. Gupta, V. Bajaj, and I. A. Ansari, "CardioNet: A lightweight deep learning framework for screening of myocardial infarction using ecg sensor data," *IEEE Sens. J.*, vol. 25, no. 4, pp. 6794–6800, 2025.

[16] Y. S. Chichani and S. L. Kasar, "An efficient IoT enabled heart disease prediction model using finch hunt optimization modified BiLSTM classifier," *Biomed. Signal Process. Control*, vol. 100, p. 107170, 2025.

[17] D. Kumar, S. Puthusserypady, H. Dominguez, K. Sharma, and J. E. Bardram, "An investigation of the contextual distribution of false positives in a deep learning-based atrial fibrillation detection algorithm," *Expert Syst. Appl.*, vol. 211, p. 118540, 2023.

[18] N. Phukan, M. S. Manikandan, and R. B. Pachori, "Noise-aware atrial fibrillation detection for resource-constrained wearable devices," in *16th Int. Conf. Electron., Comput. Artif. Intell. (ECAI)*, 2024, pp. 1–6.

[19] D. Kumar, A. Peimankar, K. Sharma, H. Domínguez, S. Puthusserypady, and J. E. Bardram, "Deepaware: A hybrid deep learning and context-aware heuristics-based model for atrial fibrillation detection," *Comput. Methods Programs Biomed.*, vol. 221, p. 106899, 2022.

[20] A. E. Voinas, D. Kumar, J. Smeddinck, A. Stochholm, and S. Puthusserypady, "Atrial fibrillation detection from ambulatory ecg with accelerometry contextualisation: A semi-supervised learning approach," in *2025 47th Annu. Int. Conf. of the IEEE Eng. in Med. and Biol. Soc. (EMBC)*. IEEE, 2025, pp. 1–7.

[21] D. Kumar, S. Puthusserypady, H. Dominguez, K. Sharma, and J. E. Bardram, "CACHET-CADB: A contextualized ambulatory electrocardiography arrhythmia dataset," *Front. Cardiovasc. Med.*, vol. 9, p. 893090, 2022.

[22] X. Xu and H. Liu, "ECG heartbeat classification using convolutional neural networks," *IEEE Access*, vol. 8, pp. 8614–8619, 2020.