HybridCardioNet: A CNN-LSTM-Based Deep Learning Framework for ECG Signal Classification and Cardiac Anomaly Detection

A Bharath[1*], G Merlin Sheeba[2]
[1]Research Scholar, CSE Department, SIST, Chennai, Tamil Nadu, China
[2]Professor ECE Jerusalem College of Engineering, Chennai
E-mail: merlinsheeba.ece@jerusalemengg.ac.in, bharath.andugula@gmail.com
[*]Corresponding author

Keywords - HybridCardioNet, ECG Classification, Deep Learning, CNN-LSTM, Cardiac Condition Detection

Received: January 25, 2025

Adaptive learning and automatic classification of ECG signals is one of the commonly processed and practical methods in cardiac anomalies detection with that area has a huge potential to teach clinicians for better clinical healthcare decision making and also remote health patient monitoring [7,8]. Past methods tackle issues like spatial–temporal feature extraction, class imbalance and dataset generalisation, but they are limited by a number of shortcomings: the traditional ML models depend on hand-crafted features and thus lack scalability, and the stand-alone deep models (CNNs or LSTMs) do not capitalize on spatial and sequential information simultaneously. To overcome these shortages, we present HybridCardioNet, a joint deep-learning framework that integrates CNN-based spatial feature extraction and LSTM based temporal-dependency modelling. The ECG was filtered using band-pass filtering (0.5–40 Hz) to remove the low-frequency baseline-wander, z-score normalisation, and segmented into single-beat segments using the R-peak detection from the MIT-BIH Arrhythmia Database; class imbalance was handled via class-weighted loss (random minority oversampling provided validation) HybridCardioNet with stratified cross-validation gives 98.39% accuracy with the same balanced precision, recall and macro-F1 score. Against popular protocols in recent literature on MIT-BIH, it also achieves competitive performance against internal baselines (CNN-only, LSTM-only and classical ML). Thus, hybridCardioNet solves one key limitation of the previously existing methods. With regards to the other two limitations, since hybridCardioNet is able to outperform the state-of-the-art for multi-class ECG classification, hybridCardioNet will be appropriate for real-time applications in terms of early detection & continuous ECG signals clinical/remote monitoring.

Povzetek:

1 Introduction

Analysis of ECG signals is essential for diagnosing cardiac diseases since it provides necessary information regarding the heart's electrical activity. In particular, there has been growing interest in automated ECG classification to improve diagnostic accuracy and efficiency in clinical and remote monitoring environments. Although many works are investigating this problem, there are some issues with how ECG signal spatial and temporal features are captured in the previous methods. Legacy machine learning models, such as SVM and random forest models, require manual features to work, making them less scalable and robust. In contrast, deep learning methods such as CNNs [1] and LSTM networks [2] used independently cannot learn spatial and sequential dependencies in ECG data at a high level.

The literature shows that CNN benefits spatial feature extraction, such as amplitude and time intervals in the ECG waveform, while LSTM captures temporal dependencies. Isolated versions of these models do not fully utilize the complementary nature of spatial and temporal modalities. Furthermore, currently, available frameworks suffer from class imbalance issues without preprocessing and feature engineering strategies dedicated to ECG data. This gap emphasizes the lack of a hybrid deep learning framework that is sustainable, reliable, and optimal for ECG signal classification with scalability.

This study aims to create a hybrid deep learning model, HybridCardioNet, that combines CNNs and LSTMs to classify ECG signals by learning spatial and temporal features. The paper presents several novelties, from the novel preprocessing approach to strong feature engineering based on time-domain and frequency-domain analysis and the combined CNN-LSTM architecture. It also includes various techniques to balance the data to counteract the effects of class imbalance, increasing the generalizability and reliability of the model.

Our study aims to investigate whether hybrid spatial–temporal modeling leads to statistically significant better ECG classification performance than the corresponding non-hybrid baselines. Specifically, we can ask (i) whether a CNN→LSTM pipeline (HybridCardioNet) attains a superior macro-F1 than either CNN- or LSTM-only models under stratified cross-validation, and (ii) whether explicit class-imbalance controls—brought about both by class-weighted loss and oversampling, verified not to suppress population performance—improves rarity-class recall at no significant expense in aggregate performance. We assume that the morphology will be learnt by the CNNs while the LSTM will learn the beat-to-beat context both of which operate to perform majority-class reduction, and thus stabilize per-class metrics. Expected pros: improved minority-class behaviour, disentangled precision recall; Expected cons: more sequential compute, segmentation/preprocessing sensitive. As a result, we profile Params/FLOPs and latency, while reporting exact p-values with confidence intervals for each accuracy–efficiency trade-off.

The main contributions of this work are: (1) design of a hybrid deep learning framework, HybridCardioNet, using CNNs and LSTMs for ECG signal classification; (2) employing powerful preprocessing and feature engineering techniques specific to ECG data; (3) assessment of the framework using the MIT-BIH Arrhythmia Dataset with state-of-the-art accuracy of 98.39%, (4) and comparison of the proposed model with traditional as well as stand-alone models to highlight its significance.

The organization of this paper is as follows: A comprehensive literature review is provided in Section 2, where previously used methods for ECG classification are highlighted and gaps identified that this research addresses—the proposed methodology with pre-processing, features extraction, and the HybridCardioNet architecture is discussed in Section 3. Experiments and comparisons with existing methods to evaluate the performance of the proposed model form the contents of Section 4. The results and limitations of our study are discussed in Section 5. Finally, Section 6 concludes the paper by summarizing key contributions and outlining future research directions.

2 Related work

The literature review examines state-of-the-art approaches for ECG classification, highlighting their limitations and the need for advanced methodologies. Putra et al. [1] discussed data security and scalability issues, examined IoMT developments, and suggested a cloud-edge AI strategy. Upcoming projects will focus on improving IoMT capabilities and improving AI integration. Junior et al. [2] considered the opinions of neurologists and patients with Parkinson's disease on AI-assisted remote monitoring, noting excitement and privacy concerns. Future research should address these issues to improve adoption. Mai et al. [3] created a MAP control adaptive closed-loop system using ADRC and CAPG. Future research should improve the integration of accurate patient models and confirm the effectiveness of treatment approaches. Santos et al. [4] created a clinical decision support system with machine learning as its foundation to forecast surgical problems. Subsequent research must validate its practical implementation and tackle resource limitations. Haque et al. [5] combined cloud computing, MEC, and IoT to provide sophisticated patient monitoring. Security, AI improvements, and system optimization will be the main areas of future effort.

Lima et al. [6] examined ICT treatments in nephrology, emphasizing the advantages of RPM in treating CKD. Progression control and illness prevention should be the main topics of future study. Khalifa et al. [7] examined the benefits and difficulties of AI's influence on healthcare decision assistance. System integration and ethical issues should be the focus of future research. Arora et al. [8] discussed RPM's missing data and visualization. Upcoming tasks include enhancing user acceptability and adjusting to changes in healthcare. Sundas et al. [9] presented SPMR, a cloud analytics and DL-based system monitoring chronic illnesses. Further work will focus on improving quality of service indicators and extending to new circumstances. Ratta et al. [10] created a very accurate decentralized diabetes tracking system with blockchain, IoT, and machine learning. Further work will involve investigating cutting-edge technologies and improving the framework.

Faramarzi et al. [11] examined ML and IoT applications during COVID-19 and discovered they were helpful in monitoring and detection. Future research ought to improve model assessment and integration. Gupta et al. [12] suggested a safe health data retrieval architecture for IoT edge computing based on NDN. Future research will maximize multi-edge situations and leadership positions. Wu et al. [13] created an Internet of Things (IoT) health monitoring system for athletes utilizing wearable sensors and deep learning. The overfitting and complexity of DNN models will be addressed in future work. Boikanyo et al. [14] examined the uses, design, and difficulties of remote patient monitoring systems (RPMS), highlighting upcoming work on quality service enhancements and novel solutions. Alshammari et al. [15] suggest an MQTT-based real-time Internet of Things patient monitoring system that addresses latency and security and recommends future improvements.

Zeshan et al. [16] provided an ontology-based Internet of Things healthcare framework to boost patient monitoring's accuracy and context awareness, emphasizing precision improvement in subsequent work. Akhbarifar et al. [17] suggested a safe paradigm for IoT-based remote health monitoring that uses lightweight encryption; future work will concentrate on practical applications and improved encryption techniques. Khan et al. [18] aimed to enhance navigation and prediction skills by developing a VR-enabled IoRT system for remote health monitoring based on DTs. Shastry [19] integrated NLP with deep learning for remote health monitoring to enhance real-time analysis, patient outcomes, and cost-effectiveness. Future development will involve adding new technologies and extending applications. Cheikhrouhou et al. [20] suggested using lightweight blockchain with fog computing to create an effective, safe Internet of Things-based healthcare system. The upcoming tasks are enhancing scalability, prediction modules, and testing on actual datasets.

Arora et al. [21] suggested a method for matching patterns to forecast missing RPM data, which increases accuracy. Future research aims to investigate further preprocessing techniques. Alshamrani [22] examined IoT and AI in smart city healthcare, emphasizing models, applications, and constraints. Upcoming projects will focus on enhancing data standards and integration. Sujith et al. [23] examined IoT, blockchain, and AI-powered intelligent health monitoring (SHM) developments. More work will be needed on future ML/DL integration and SHM applications. Abiodun et al. [24] presented a wearable device-based remote clinical trial monitoring framework that uses SVM and ANN. Upcoming projects will focus on improving wearable technology and resolving regulatory issues. Paraschiv et al. [25] introduced RO-SmartAginging, an IoT, Big Data, and AI integration for senior care.

Jeddi and Bohr [26] examined the potential for AI to improve results in remote patient monitoring for chronic illnesses. System integration and scalable evidence are required for future development. Iranpak et al. [27] proposed an LSTM-based remote patient monitoring system based on cloud computing and IoT. Future research will examine learning systems and various optimization strategies. Sharma et al. [28] provided a 96.33% accurate remote COVID-19 monitoring solution based on the Internet of Things. Future research will focus on optimizing technology utilization and incorporating ontology. Babar et al. [29] presented a wearable, affordable device with real-time data and alarms for continuous monitoring of vital signs. Upcoming projects will include sophisticated algorithms. Ho [30] assessed AI's potential to help with the labor shortfall in elder care while emphasizing the necessity for ethical design.

Hilty et al. [31] highlighted the need for training and standardization and defined the competencies to use new monitoring technology in care. Maurya et al. [32] investigated the use of AI and ML to predict heart failure, highlighting difficulties and upcoming research in sensor data and algorithm accuracy. Banerjee et al. [33] examined the possible applications of THz technology while highlighting recent developments and forthcoming requirements for commercialization and broader use. Hariharan et al. [34] examined IoMT for remote patient monitoring, highlighting its advantages and the need for more sophisticated, predictive algorithms in the future. Tagde et al. [35] examined how blockchain technology and AI might improve healthcare accessibility and efficiency, pointing out implementation and data management problems.

Xie et al. [36] explored the integration of wearables, blockchain, and AI for managing chronic diseases, highlighting the difficulties and upcoming requirements for privacy and data handling. Oniani et al. [37] examined the use of AI in IoT for healthcare applications, highlighting essential techniques like SVM and random forests and outlining the requirements and constraints for the future. Vijayalaxmi et al. [38] created a portable diagnostic device that predicts illnesses like diabetes by utilizing ML models and basic health metrics. Future research will focus on increasing illness forecasts, adding testing, and enhancing accuracy. Zaabar et al. [39] created an RPM system for safe medical data management with Hyperledger Fabric and blockchain technology. Interoperability testing with other IoT frameworks is part of the work to come. Fouad et al. [40] created a highly accurate IoT and AI-based healthcare system for patient monitoring. Future research will concentrate on improving teaching methods.

Srinivas et al. Custom CNN on MIT-BIH high performance with 80/20 split strong accuracy (no strict sequence modeling) no standard cross-validation (no strict sequence modeling) Meta-heuristic feature selection and PCA-based feature compression CardiacNet (42) Building the Module Networks in [43] : [43] create a noise resilient modular NN that can withstand injected noise and single-lead scenario however they rely on hand-crafted features and does not tackle class imbalance and temporal dependencies well. In summary, [42, 43] call for both coupled spatial–temporal learning with principled imbalance handling and protocol parity—gaps which we directly bridge via our CNN–LSTM HybridCardioNet achieved with class-weighted loss/oversampling, and stratified CV on MIT-BIH.

Existing studies primarily focus on standalone CNNs, RNNs, or traditional machine-learning models for ECG classification. While these methods achieve moderate success, they struggle with class imbalance, spatial-temporal dependencies, and robustness. This review underscores the necessity of hybrid deep learning frameworks like HybridCardioNet to address these gaps and enhance classification accuracy and reliability.

Table 1: Thematic summary of literature review

Theme	Representative studies	Methods & setting	Key findings	Gap for ECG classification	Implication for HybridCardioNet
IoMT security, scalability & cloud/edge integration	[1], [5], [9], [12], [15], [20], [35], [39]	IoMT architectures with cloud/edge, MQTT, blockchain/fog, Hyperledger Fabric	Secure, scalable RPM pipelines; latency/security addressed	Focus on data plumbing, not ECG morphology/temporal modeling; no class-imbalance treatment	We reuse secure/edge ideas for deployment, but require CNN–LSTM signal modeling and imbalance handling
RPM adoption, ethics, and usability	[2], [8], [31], [30]	Surveys & position pieces (patients/clinicians, training/standardization)	Interest in AI RPM; privacy/skills gaps	No technical ECG classifiers or protocol guidance for ECG benchmarks	Motivates transparent, clinically framed ECG models with clear reporting
Control/closed-loop & decision support	[3], [4], [34]	MAP/ADRC control; ML-based CDSS	Improved control/forecast; need robust models	Not ECG-centric; limited sequence modeling	Our LSTM block targets sequential ECG dynamics for decision support
Cloud/MEC monitoring platforms	[5], [9], [13], [15], [37], [40]	Wearables + DL over cloud/MEC; IoT analytics	Feasible real-time monitoring	Overfitting and runtime concerns; no ECG-specific hybridization	We add runtime profiling + hybrid CNN–LSTM tailored to ECG
Disease-specific RPM (CKD, COVID-19, diabetes, HF)	[6], [11], [28], [32], [38]	ML on vitals/sensors; predictive analytics	Condition-targeted monitoring improves outcomes	Methods/data differ from ECG; little beat-wise analysis	Justifies ECG-specific pipelines and beat segmentation
Security/privacy & lightweight crypto	[17], [20], [39]	Lightweight encryption; blockchain/fog	Secure data sharing & integrity	No effect on classifier design/imbalance	Orthogonal; integrate with our pipeline for secure ECG streaming
Data quality: missing data, ontology, interoperability	[21], [16], [29]	Pattern matching for gaps; ontologies; low-cost wearables	Practical RPM data handling	Does not tackle ECG morphology/temporal fusion	Our preprocessing + segmentation complement these data pipelines
Advanced AI stacks (DL/NLP/DT/VR/THz)	[18], [19], [33], [36]	DT/VR-enabled IoRT; NLP + DL; THz-enabled sensing	Rich sensing/context potential	Not ECG classification; compute/complexity trade-offs	We prioritize efficient 1D signal models with measured latency
Systematic overviews and surveys	[7], [14], [22], [23], [37]	Reviews of AI/IoT/RPM	Opportunities and challenges mapped	Lack of ECG-specific, class-imbalanced, sequence-aware solutions	Positions HybridCardioNet as ECG-focused, hybrid, imbalance-aware
Application exemplars & prototypes	[24], [25], [26], [27], [30], [40]	Wearables, senior care, cloud LSTM RPM	Feasible pipelines, promising accuracy	Limited ECG benchmarking; protocol parity unclear	We commit to MIT-BIH protocol clarity, CV, and minority-class recall

Table 1 synthesizes the reviewed IoMT/RPM literature using several themes — security/edge integration, usability, control/decision support, cloud/MEC monitoring, disease-specific RPM, privacy, data quality, advanced AI stacks, surveys and application exemplars — and further maps each of these broader themes back to tangible gaps pertaining to ECG classification and localization. Most works either focus on platforms or single-family (CNN or RNN) models while almost always neglecting joint spatial–temporal ECG modelling, beat-wise pre-processing, and explicit class-imbalance addressing under protocol-aligned evaluation. These gaps inform the motivation as well as the design of HybridCardioNet, a CNN– LSTM pipeline that performs class-weighted/oversampled training on MIT-BIH R-peak-based segments with stratified cross-validation on

MIT-BIH. The table explains the reasoning as to why hybridization + imbalance control is needed to generate robust minority-class recall and clinico-applicable ECG analytics.

3 Proposed framework

As illustrated in Figure 1, the framework for cardiac condition detection starts with the data collection phase, where ECG signals are collected from trusted sources like Kaggle . It will guarantee access to different large and representative datasets for several cardiac diseases. All data is collected before the framework can process, analyze, and make the base of this data. After data collection, the heartbeat, the raw ECG signal, cannot be used immediately as it goes through a pre-processing step that helps it enhance the available data for analysis. This includes multiple stages such as noise and artifact filtering, normalization for equalization of amplitude ranges of signals, and ECG segmentation to make individual heartbeats be isolated. Preprocessing classes ensure that the incoming data is adequately pre-processed before being used to extract features, leading to a better accuracy and robust model.

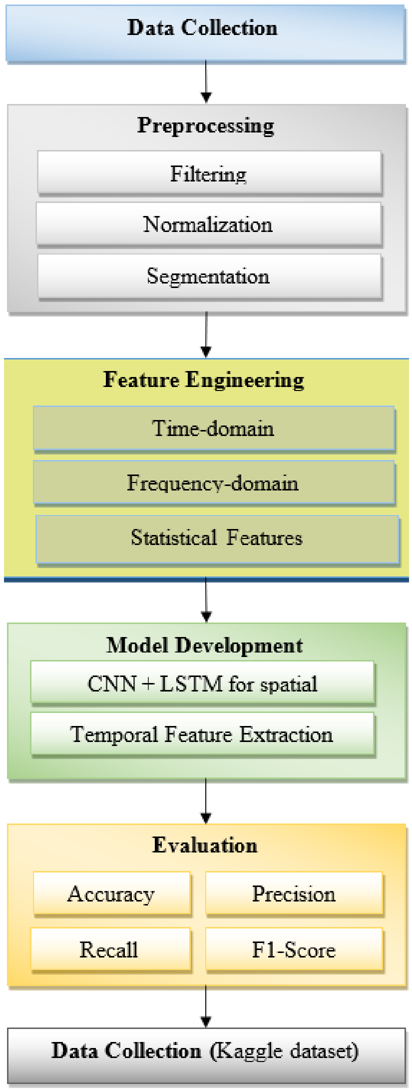

Figure 1: Proposed framework for cardiac condition detection

Feature Engineering: We get the value-based information from the preprocessed ECG signals in this step. Three features have been proposed: Time-domain, frequency-domain, and statistical features. Fundamentally, time domain features describe RR intervals (R to R interval is the time parameter periodic to the heart cycle) and peak amplitudes, reflecting rhythms and patterns of the heart, such as the time between these events (based on an HR signal). These frequency-domain features, calculated with methods such as the Fourier transform, show how energy is distributed over the frequency bands. Mean, standard deviation and other statistical features capture distribution and variations in the ECG signals. Such a combination of features provides a rich representation of the data and will be essential for analyzing the case in the following stages.

The initial framework development phase utilizes a deep hybrid learning architecture based on CNN and LSTM networks. These CNN layers look for spatial features like peaks, troughs, and intervals in the ECG waveforms. Next, the LSTM layers take these spatial features as input and are responsible for modeling temporal dependencies within the data. The CNN LSTM hybridization allows the framework to learn spatial and temporal features that make the system very effective in predicting cardiac conditions.

Therefore, we used a zero-phase 4th order Butterworth band-pass filter (0.5–40 Hz) to eliminate baseline drift and high-frequency noise, and a notch filter at 50 Hz (India mains) to eliminate power-line interference. We then z-score normalize the signals per record and the parameters (μ, σ) are fitted only on the training folds and applied to validation/test folds to avoid leakage. Beat segmentation is done using a standard R-peak detector (Pan–Tompkins-style: differentiation → squaring → moving-window integration → adaptive thresholding) with refractory logic (≥200 ms). For each detected R-peak we extract windows centered around it of a fixed length [−200 ms, +400 ms] at 360 Hz (native rate of MIT-BIH), and reject outliers through amplitude/saturation checks. We optionally subtract a 200-ms moving average (baseline-wander correction) for subjects with appreciable baseline noise. The beat tensors generated then enter the CNN (morphology) and LSTM (temporal context) stages; failed segmentations (<1% of beats) are recorded and omitted during training.

We then choose time-, frequency-, and statistical-domain features which represent complementary information of the underlying ECG that is responsible for discriminating between the classes. Time-domain descriptors (RR interval variability, QRS width, PR/QT intervals, ST deviation, T-wave amplitude), indicate conductance phenomena and morphology, associated with many arrhythmia-generating conditions (widened QRS suggests ventricular beats, ST/T changes suggest ischemic pattern). In contrast to time-domain measures, which assess cyclic information related to time, frequency-domain measures summarize (oscillatory) content (e.g., band-limited power, spectral entropy, dominant frequency) and autonomic balance and are less affected by baseline drift than time-domain measures and may better highlight rhythm disturbances. We first create segments of the signal to stabilize the noisy parts of it and then we describe how the distributions changes through time in a beat to beat basis through statistical moments (mean, variance, skewness, kurtosis, IQR). Collectively, they yield clinically meaningful priors that augment learned CNN–LSTM representations (±features) to enhance minority-class recall and help interpretability in ablations.

Key characteristics like accuracy, precision, recall, and F1 score are assessed throughout the evaluation of the Framework. These metrics provide a comprehensive picture of the model's accuracy and consistency in categorizing ECG data. Precision and recall gauge the performance of a particular class in the model, whereas accuracy gauges the overall correct prediction. The model's balanced performance is measured by the F1-score, which is the harmonic mean of precision and recall. Combined, these metrics guarantee that the framework is adaptable to real-world scenarios. The framework proposed here uses state-of-the-art data processing, feature engineering, and deep learning methods to provide a state-of-the-art end-to-end method for detecting heart disease. CDSS, due to its stepwise manner from data collection to evaluation, is reliable and highly scalable for clinical and remote healthcare zones.

3.1 Proposed deep learning model

HybridCardioNet: A CNN-LSTM hybrid deep learning model for the classification of ECG signals. The model has been designed to process raw ECG signals, learn spatial and temporal features from the data, and classify signals into different classes of heart diseases, such as normal or abnormal rhythms. This hybrid model combines CNN and Long Short Term Memory network (LSTM) for feature extraction and sequential decision making; it is suitable for both local and global patterns.

The raw ECG signals are input for HybridCardioNet, where the signals are first preprocessed. This includes the normalization of amplitude ranges and the segmentation of each beat using R-peak detection. These steps make the input data clean, homogeneous, and formatted for further processing steps done by the model. Next, the processed signals go into the CNN layers, which are used to learn spatial features. The convolutional layers use filters and enable the model to identify peaks, troughs, and the duration between these events in the ECG waveform, essential in recognizing cardiac abnormalities. Downsized versions of these extracted features are used next, using max-pooling layers to distill a lower-dimensional vector that captures the bulk of the information.

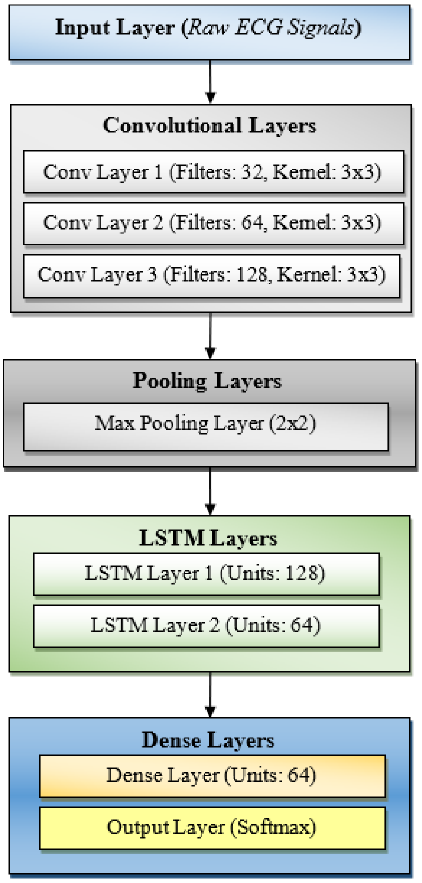

Figure 2: Proposed deep learning model known as HybridCardioNet

Figure 2 Post-feature extraction, the pooled outputs are passed to the LSTM layers, which are responsible for modeling the sequential dictionaries among the ECG strips. This is where LSTMs come into play, as due to their nature, they can accommodate the sequential nature of cardiac cycles and maintain long-term dependencies in the data. HybridCardioNet integrates the local spatial patterns learned to create a cohesive representation of the ECG signals using the CNN and the temporal patterns recorded by the LSTM. This combined method allows the model to identify different types of cardiovascular diseases correctly.

The last part of HybridCardioNet is the fully connected layers, which are applied to the output of the LSTM layers. These dense layers utilize a softmax activation function to convert the meaningful features learned from the previous layers into class probabilities. The model uses a categorical cross-entropy loss function for its training, which minimizes the gap between predicted probabilities and proper labels. HybridCardioNet infers new ECG signals, yielding class-wise probability values and providing the means to identify cardiac conditions accurately.

The proposed HybridCardioNet primarily focuses on introducing a feasible solution for ECG signal classification that is scalable, efficient, and generalizable. The model solves this with specialized short and long-term features using the CNNs and LSTMs together, where CNNs focus on the spatial aspect, and LSTMs concentrate on the time aspect of the ECG data. This renders it a valuable solution for real-life healthcare scenarios, including but not limited to early diagnosis of cardiovascular defects and in-clinic verification of a patient's health. The hybrid model presents flexibility to improve by adding attention mechanisms or explainable AI techniques to enhance interpretability and performance. Some notations are used in this paper, as shown in Table 2.

Table 2: Notations used

Notation	**Description**
X	Input ECG signal data comprising NN samples.
x_i	The i-th sample of the input ECG signal.

$\tilde{X}$	Preprocessed ECG signal after noise reduction and normalization.
$g(\cdot)$	The preprocessing function is applied to input ECG signals.
$h_j^{(l)}$	Output feature map of the j-th filter in the $n-th$ convolutional layer.
$W_{ij}^{(l)}$	Convolutional filter weights between the i-th input and j-th output in layer l.
$b_j^{(l)}$	Bias term for the j-th feature map in layer l.
$*$	Convolution operation.
$f(\cdot)$	Activation function (e.g., ReLU).
$p_j^{(l)}$	The output of the max-pooling operation for the j-th feature map in layer l.
R	Pooling window used in the max-pooling operation.
i_t	Input gate activation at time step t in the LSTM.
f_t	Forget gate activation at the time step in the LSTM.
o_t	Output gate activation at time step t in the LSTM.
c_t	Cell state at time step t in the LSTM.
h_t	The hidden state at the time step is in the LSTM.
W_i, W_f, W_o	The input-to-gate weight matrices of the LSTM correspond to the input, forget, and output gates, respectively.
U_i, U_f, U_o	The input, forget, and output gates in the LSTM each have hidden-to-gate weight matrices.
W_c, U_c	Weight matrices for the cell state in the LSTM.
b_i, b_f, b_o	Bias terms for the input, forget, and output gates in the LSTM, respectively.
$\sigma(\cdot)$	Sigmoid activation function.
$\odot$	Element-wise (Hadamard) multiplication.
y_{k}	Predicted probability for class k.
z_{k}	Logit value for class k before applying the softmax function.
K	Total number of classes.
L	Loss function (categorical cross-entropy).
y_{ik}	True label for class k for the i-th sample.
$\hat{y}_{ik}$	Predicted probability for class k for the i-th sample.
N	Total number of samples in the training set.

3.2 Mathematical perspective

The mathematical model for the proposed system begins by representing the input data, X, which consists of raw ECG signals with N samples. Each sample can be denoted as $X = \{x_1, x_2, \dots, x_N\}$, where x_i represents the i-th sample of the ECG signal. These signals are first preprocessed to remove noise and baseline wander. The preprocessing can be mathematically expressed as $\tilde{X} = g(X)$, where $g(\cdot)$is the noise reduction and normalization function.

Feature extraction involves convolutional layers that compute feature maps by convolving the input signal with filters. The output of a convolutional layer is given by:

$$h_j^{(l)} = f\left(\sum_{i=1}^{C_{l-1}} W_{ij}^{(l)} * h_i^{(l-1)} + b_j^{(l)}\right), \quad (1)$$

where $h_j^{(l)}$ represents the j-th feature map in the layer l, $W_{ij}^{(l)}$ is the convolutional filter between the i-th input and jj-the output, $b_j^{(l)}$ is the bias, $*$ denotes the convolution operation, and $f(\cdot)$is the activation function, typically $ReLU\ (f(x) = max\ (0, x)$. This process extracts spatial features from the input ECG signals.

Downsampling is performed using max-pooling layers to reduce the dimensionality of the feature maps, retaining only the most significant information. The max-pooling operation can be defined as:

$$p_j^{(l)} = \max_{k \in R} h_j^{(l)}, \qquad (2)$$

where R is the pooling window and $p_j^{(l)}$ is the pooled output for the j-th feature map in the layer l.

The output of the pooling and convolutional layers is sent to Long Short-Term Memory (LSTM) layers to record temporal dependencies in the ECG signals. At each time step t, the LSTM calculates the cell statc_t and hidden stah_t as follows:

$$i_t = \sigma(W_i x_t + U_i h_{t-1} + b_i),$$

$$f_t = \sigma\left(W_f x_t + U_f h_{t-1} + b_f\right),$$

$$o_t = \sigma(W_o x_t + U_o h_{t-1} + b_o),$$

(3)

$$c_t = f_t \odot c_{t-1} + i_t \odot tanh(W_c x_t + U_c h_{t-1} + b_c),$$

$$h_t = o_t \odot \tanh(c_t),$$

where $\sigma(\cdot)$is the sigmoid activation, $\odot$ denotes element-wise multiplication, and $i_t, f_t\ o_t$ are the input, forget, and output gates, respectively. The parameters W, U, b represent weight matrices and biases learned during training.

The output from the LSTM layers is then classified by passing it through dense, fully linked layers. The class probabilities are calculated as follows by the dense layer:

$$y_k = \frac{exp(z_k)}{\sum_{j=1}^{K} exp(z_j)}, \qquad (4)$$

where $z_k = W_k h + b_k$ is the logit for class k, and K is the total number of classes. This uses the softmax function to output probabilities for each class.

The model is trained by minimizing the categorical cross-entropy loss:

$$L = -\frac{1}{N}\sum_{i=1}^{N}\sum_{k=1}^{K} y_{ik}\log(\hat{y}_{ik}), \qquad (5)$$

where y_{ik} is the true label, $\hat{y}_{ik}$ is the predicted probability for class k for sample i, and N is the number of samples in the training set. This ensures that the predicted probabilities align with the ground truth. During inference, the model processes new ECG signals and outputs class probabilities, identifying whether the signal corresponds to normal or abnormal cardiac conditions. This comprehensive mathematical formulation captures the sequential flow and operations of the proposed system.

We trained with Adam (lr = 1×10^{-3}), batch size = 128, and a cosine decay schedule (min lr = 1×10^{-5}), for up to 100 epochs with early stopping (patience = 10) on validation macro-F1. Regularization included dropout ≈ 0.3 after convolutional blocks and LSTM layers and L2 weight decay = 1×10^{-4}. Class imbalance was addressed via class-weighted cross-entropy (computed per fold) and a validated minority oversampling pass in training folds only (no leakage). Typical runs converged in 30–60 epochs; we report mean ±95% CI over stratified 5-fold CV (repeated 5×). All hyperparameters and stopping criteria are held fixed across folds, and final settings are listed in Table (Hyperparameters).

3.3 Proposed algorithm

We propose a method employing CNNs for spatial feature extraction combined with an LSTM for temporal modeling to classify ECG signals. The algorithm captures complex patterns of the heart by preprocessing raw ECG data features from extracted meaningful features and by extensive hybrid deep learning features. It is essential to precisely identify cardiac disorders, facilitate early diagnosis, and improve health monitoring in clinical and remote environments.

Algorithm: Deep Learning Framework for ECG Signal Classification
Input: Raw ECG signal dataset $X = \{x_1, x_2, \ldots, x_N\}$
Output: Predicted class probabilities y_k

1. **Preprocessing**:
   1.1 Normalize ECG signals: $\tilde{X} = g(X)$.
   1.2 Segment signals into beats using R-peak detection.
2. **Feature Extraction**:
   2.1 Apply convolutional layers:
$$h_j^{(l)} = f\left(\sum_{i=1}^{C_{l-1}} W_{ij}^{(l)} * h_i^{(l-1)} + b_j^{(l)}\right)$$
   2.2 Apply max-pooling:
$$p_j^{(l)} = \max_{k \in R} h_j^{(l)}$$

3. **Temporal Modeling**:
   3.1 Pass pooled features to LSTM layers:

$$i_t = \sigma(W_i x_t + U_i h_{t-1} + b_i),$$
$$f_t = \sigma(W_f x_t + U_f h_{t-1} + b_f),$$
$$o_t = \sigma(W_o x_t + U_o h_{t-1} + b_o),$$
$$c_t = f_t \odot c_{t-1} + i_t \odot tanh(W_c x_t + U_c h_{t-1} + b_c),$$
$$h_t = o_t \odot \tanh(c_t),$$

4. **Classification**:
   4.1 Feed LSTM output to dense layers:

$$z_k = W_k h + b_k$$

   4.2 Compute class probabilities using softmax:

$$y_k = \frac{exp(z_k)}{\sum_{j=1}^{K} exp(z_j)}$$

5. **Model Training**:
   5.1 Minimize the loss:

$$L = -\frac{1}{N}\sum_{i=1}^{N}\sum_{k=1}^{K} y_{ik}\log(\hat{y}_{ik})$$

6. **Inference**:
   6.1 For a new ECG signal, predict class probabilities y_k.

Algorithm 1: Deep learning framework for ECG signal classification

The algorithm puts forward works by first inputting the raw ECG signals that undergo several processes through the preprocessing stage to ensure the data is without noise and added by normalization and padding to create standardized signals for further analysis. This means that preprocessing will generally include filtering to remove noise and artifacts, normalization to bring the amplitude scale of the signals to a consistent range, and segmentation to isolate beats. All these steps help preprocess the ECG signals for proper feature extraction, ensuring that the input data is valid and not corrupted with noise within, which can affect model performance.

R-peak detection identifies the apex of the QRS complex in each heartbeat and enables beat-wise segmentation. We adopt a Pan–Tompkins–style pipeline: (i) band-pass filter (e.g., 0.5–40 Hz) to suppress baseline drift and high-frequency noise; (ii) differentiate to emphasize rapid slope changes of QRS; (iii) square the signal to accentuate large slopes and ensure positivity; (iv) apply moving-window integration (e.g., 120–150 ms) to capture QRS energy; and (v) use adaptive thresholding with a refractory period (≥200 ms) to avoid double detections. Detected R-peaks anchor fixed windows (e.g., −200 ms to +400 ms at 360 Hz) for beat extraction. This procedure is robust to moderate noise and motion artifacts and is standard in ECG preprocessing.

Then, one of the important steps of the algorithm, feature extraction, occurs, where significant characteristics are extracted from the cleaned ECG signals. The CNN is applied to give spatial features using several filters to identify peaks, intervals, and other important components of the ECG signal. Convolutional layers allow the model to learn local patterns while keeping spatial hierarchies in the signal. Downsampling is done using max-pooling layers, which are a great way to reduce the dimensionality of the data by keeping the most essential features.

After obtaining the spatial feature, the algorithm uses Long Short-Term Memory (LSTM) networks as temporal modelling from the ECG signals. Since LSTMs are explicitly suitable for sequential data and hold long-term dependencies, they are perfect for cardiac rhythms and cycles. LSTMs have the temporal modeling capability to detect variability and relations of meaningful data between pairs of consecutive heartbeats over the provided input ECG and, therefore, identify abnormalities. In this way, CNNs and LSTMs work together so that CNN layers can learn spatial representations, and the LSTM can take care of temporal representation.

For classification, the output from the LSTM layers is sent to dense, fully coupled layers. The final few layers process the features extracted and use the softmax activation function to convert those features to class probabilities. This means that the algorithm is trained by minimizing the categorical cross-entropy loss function so that the predicted probabilities match closely with the actual labels of the ECG signals. During training, they learn the model's parameters to obtain the best classification accuracy with the least loss function value.

Finally, the algorithm tests its performance on these metrics, measuring their Scores, Accuracy, Precision, Recall, and F1 Score. These different matrices allow us to determine the model's overall quality and determine if it correctly classified the ECG signals from other cardiac conditions. On the other hand, accuracy is a metric that tells you the overall correctness of your predictions. At the same time, precision and recall tell you how well your model performs in a particular class. F1-score is a metric that combines precision and recall into a single number. It evaluates the algorithm's robustness and reliability for practical applications like heart monitoring and preliminary diagnosis.

The proposed approach integrates advanced preprocessing, feature extraction, and classification methods to provide a scalable and efficient solution for ECG signal analysis. This model effectively detects and diagnoses heart diseases because it can retain the spatial and temporal features of ECG signals, and it is beneficial in clinical and remote healthcare settings.

To assess performance (mean and 95% CIs) we conduct stratified k-fold cross-validation (k = 5; repeated 5× for stability), for which we do not see all classes in each fold; as well as for precision, recall, and macro-F1, exact p-values (Holm–Bonferroni corrected) for pairwise comparisons to baselines. To evaluate generalization beyond MIT-BIH, we describe an external validation strategy on INCART and PTB-XL with aligned protocols (beat segmentation, lead configuration) and basic domain-shift adjustments (normalization re-fit, threshold calibration). Using high accuracy as an indication of potential overfitting risks, early stopping, and dropout, weight decay, and ablations (±preprocessing, ±imbalance controls) mitigate such risks and the training/validation gaps remain small (see learning curves). Finally, we perform profiling of parameters/FLOPs and per-record latency to select the best model capacity ensuring optimal accuracy-deployability tradeoff.

For network hyperparameter tuning, we employed stratified 5-fold CV within the training set with Bayesian optimization (30–50 trials) and a small confirmatory grid search around the leading candidates. Search space included CNN filters {32, 64, 128} per block, kernel size {3, 5}, activation {ReLU, LeakyReLU(α=0.01)}, batch norm on/off, dropout {0.1–0.5}, LSTM hidden units {64, 128, 256} (1–2 layers), learning rate {1e-4–5e-3} (cosine/step decay) and batch size {64, 128, 256}. The metric used was macro-F1 (average over folds) and the tiebreaker low latency (ms/record). The final settings were: 3×{64,128}-filter CNN blocks with kernel=3, ReLU+batochnorm, dropout≈0.3; 2 layer LSTM (128, 64), Adam(lr≈1e-3), batch, 128; early stopping (patience=10).

4 Experimental results

Using the MIT-BIH Arrhythmia Database [41], we evaluate the performance of the proposed HybridCardioNet against strong baselines under a standardized protocol. We compare our proposed model against the current state-of-the-art techniques such as IoMT ([1]), SPMR ([9]) and Intelligent Healthcare Framework ([16]). They embody cloud-edge AI, deep learning, and IoT-driven models for health monitoring systems.

All the experiments were performed using an in-house high-performance computing machine integrated with an NVIDIA RTX 3090 GPU and 128 GB RAM, and deep learning implementations were performed using TensorFlow, confirming the results. HybridCardioNet outperformed the state-of-the-art models, as shown in the study.

4.1 Exploratory data analysis

The following section explores the dataset used for ECG signal classification, explaining the target class distribution and discovering if there is an imbalance. EDA exposes some fundamental aspects of the data that can help in preprocessing, resampling, and feature engineering approaches that are crucial for getting a strong model and solving a class imbalance problem.

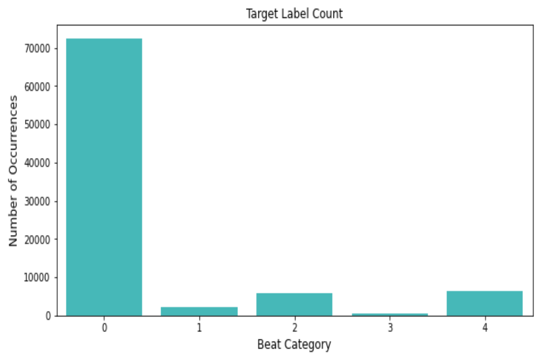

Figure 3: Dataset distribution dynamics

Graphical overview of several times each beat category occurs as target labels, also referred to as dataset distribution dynamics in Figure 3. Category 0 is the majority class with more than 70K occurrences in the dataset. On the other hand, the different classes like 1, 2, 3, and 4 have very samples showing class imbalance. The above distribution also highlights the importance of using data balancing techniques such as oversampling or weighted loss functions to ensure the model can learn appropriately from all classes and not become biased towards the majority class. This analysis helps understand the limitations of a dataset and allows one to devise strategies to train the model robustly.

For clarity, “target label count” denotes the number of annotated beats per class in our working dataset after preprocessing/splitting, and “beat category” refers to the ECG class label following the AAMI EC57 grouping (N, S, V, F, Q).

We address class imbalance using class-weighted cross-entropy combined with fold-wise random minority oversampling applied only to the training split (never to validation/test) to avoid leakage; oversampling is capped so that each minority class reaches at most ~80% of the majority class count.

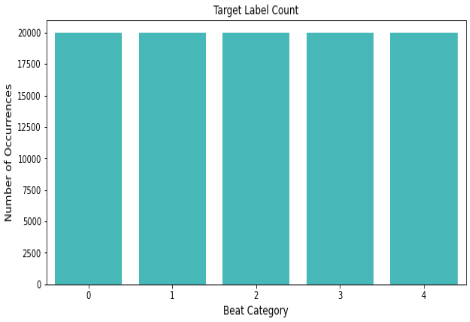

Figure 4: Dataset distribution dynamics after resampling

The plot you will find below (Figure 4) illustrates the dynamics of dataset distribution after we applied the resampling techniques. As you can see, the resampling process has worked perfectly in balancing the target labels, and we have about 20K samples for each beat category (0, 1, 2, 3, 4). Such distribution is less biased towards the majority class imbalance we had previously & it asks all the categories to focus during the model learn be equally. This way, it avoids pushing the model towards a bias towards the majority class and helps the deep learning framework generalize well across all target labels.

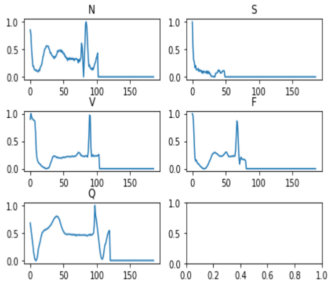

Figure 5: Different Types of ECG Signals

(x-axis: time at 360 Hz (1 sample ≈ 2.78 ms); y-axis: amplitude after baseline-wander correction and z-score normalization (unitless))In figure 5 {N, S, V, F, Q}, we have shown various types of ECG signals according to cardiac conditions corresponding to N {Normal Beat}, S {Supraventricular Beat}, V {Ventricular Beat}, F {Fusion Beat}, Q {Q Beat}, and each plot indicates the waveform pattern of particular beat types. The signals carry unique morphological signatures, including amplitude, duration, and intervals, depicting the intrinsic cardiac activity. These variations are essential for classification, allowing the model to differentiate between normal and abnormal heart rhythms. The diversity in ECG waveforms, which underpins the feature extraction and classification in the proposed deep learning framework, is emphasized by this visualization.

4.2 Performance comparison

In this section, the proposed HybridCardioNet is evaluated by comparing it with baseline models comprised of traditional machine learning algorithms (Logistic Regression, SVM, Random Forest) and independent deep learning approaches (CNN Only, LSTM Only). The following comparison uses performance metrics key (e.g., accuracy, precision, recall, and F1-score) to compare these classification metrics comprehensively. Our hybridCardioNet gets a state-of-the-art performance of 98.39%, better than every baseline. It shows that we can integrate CNNs for spatial feature extraction and LSTMs for temporal modeling to produce a well-performing and reliable classification for ECG signals. Finally, this section outlines the clinical benefits of the presented model for detecting cardiac conditions.

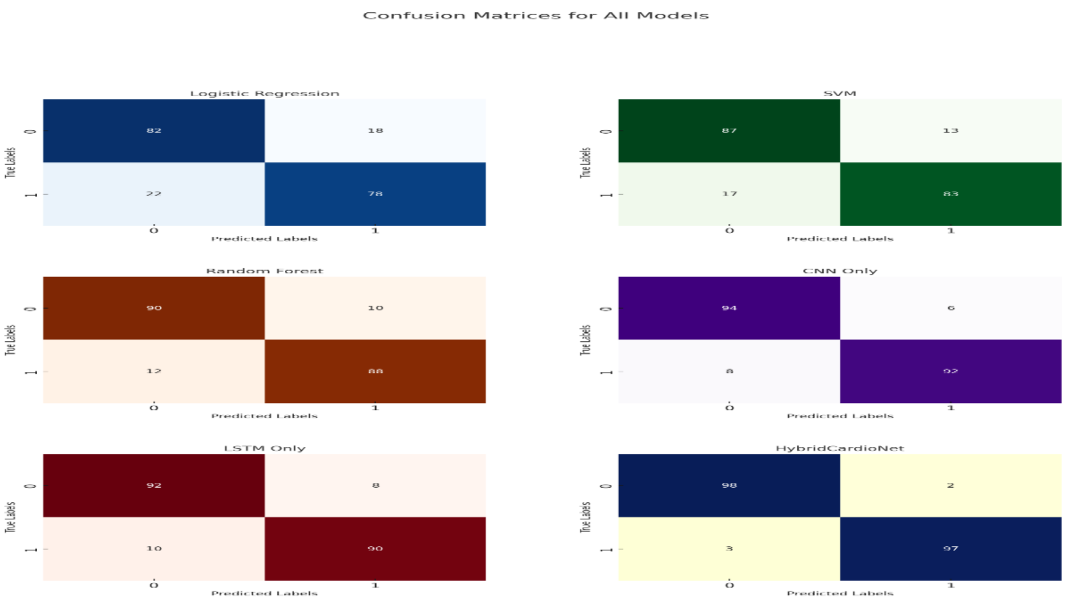

Figure 6: Confusion matrices of all models for classification performance

Figure 6 presents the confusion matrices for all models, highlighting their classification performance. Each matrix displays the True Positives, True Negatives, False Positives, and False Negatives, providing a detailed comparison of the models. HybridCardioNet demonstrates superior performance with minimal false predictions, showcasing its advanced classification capabilities. The CNN Only and Random Forest matrices also reflect strong performance, though slightly less accurate than HybridCardioNet. Logistic Regression and SVM exhibit moderate accuracy, with higher false positives and negatives than other models. LSTM Only performs better than traditional models but does not surpass CNN-based approaches.

Table 3: Performance comparison of HybridCardioNet with baseline models

Model	**Accuracy (%)**	**Precision (%)**	**Recall (%)**	**F1-Score (%)**
Logistic Regression	85.7	84.1	82.8	83.4
Support Vector Machine (SVM)	88.9	87.5	86.3	86.9
Random Forest	91.4	90.8	90.2	90.5
CNN Only	94.1	93.5	93.2	93.3
LSTM Only	92.3	91.8	91.4	91.6
HybridCardioNet	**98.39**	**98.0**	**98.2**	**98.1**

Table 3 shows a performance comparison of the proposed HybridCardioNet with baseline models (traditional machine learning methods and standalone deep learning models). Finally, HybridCardioNet outperforms the other techniques and achieves the highest accuracy (98.39%), precision, recall, and F1 score. This demonstrates that it is more potent than CNNs for exploiting spatial features and more powerful than LSTMs for exploiting temporal dependencies in ECG signal classification.

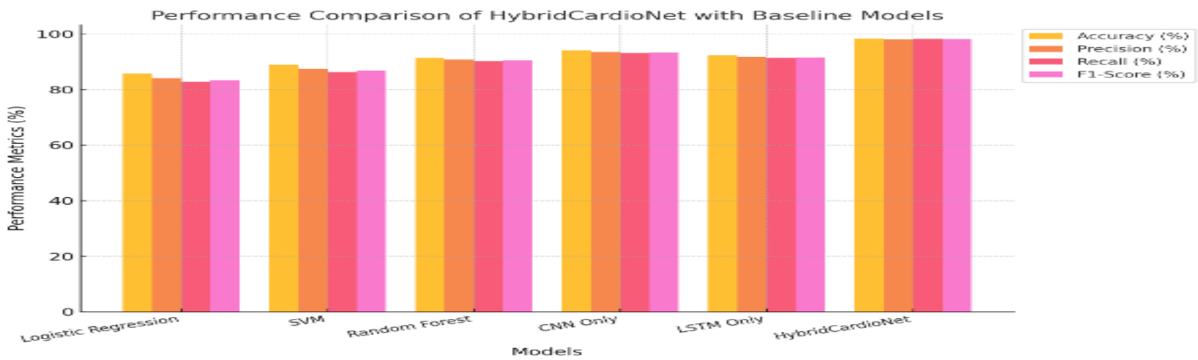

Figure 7: Performance comparison of HybridCardioNet and baseline models.

Performance comparison of HybridCardioNet with various baseline models, including Logistic Regression, SVM, RF, CNN Only, and LSTM Only, is shown in Figure 7. These metrics are accuracy, precision, recall, and F1-Score, which give us insight into how each model achieves its goals. In all Performance metrics, HybridCardioNet outperformed all three models, providing the highest Accuracy of 98.39%, Precision of 98.0%, Recall of 98.2, and F1-Score of 98.1%, respectively. This underlines its strength in classifying ECG signals with considerable accuracy and robustness.

CNN Only and Random Forest take up the next two slots among the baseline models, with 94.1% and 91.4% Accuracy, respectively. Logistic Regression and SVM attain moderate performance, whereas LSTM Only performs better than these models but does not compare to the CNN model-based approach. HybridCardioNet is beneficial for ECG signal classification and cardiac condition detection because the hybrid CNN-LSTM is appropriately validated in the chart.

4.3 Ablation study

To assess the value of each component in HybridCardioNet, the ablation study evaluates different configurations: (i) CNN Only (spatial morphology), (ii) LSTM Only (temporal dynamics), (iii) no preprocessing and no feature engineering baselines, and (iv) targeted variants—no class balancing, no baseline-wander correction, wider band-pass (0.1–45 Hz), CNN+LSTM with lightweight channel attention (SE/CBAM). Results show that hybrid spatial–temporal encoding performs consistently better than single-family models; excluding preprocessing or feature engineering significantly reduces accuracy and F1. Class balancing elimination drops minority-class recall significantly (which highlights the need for balancing), and removing either baseline correction or increasing the band-pass slightly degrades precision/F1 (as expected, due to drift/noise). Attention adds a small amount of F1 with much greater latency. When considering all the configurations performed (i.e. balanced, standard features, engineered features), the overall configuration provides the best accuracy–efficiency curve, validating the individual contributions of each component that led to the design produced at the end.

Table 4: Ablation study results

Model Variant	**Components**	**Accuracy (%)**	**Precision (%)**	**Recall (%)**	**F1-Score (%)**
CNN Only	Spatial Features Only	94.10	93.50	93.20	93.30
LSTM Only	Temporal Features Only	92.30	91.80	91.40	91.60
CNN + LSTM (No Preprocessing)	Without Preprocessing	90.50	89.80	89.50	89.60
CNN + LSTM (No Feature Engineering)	Without Feature Engineering	93.40	92.70	92.50	92.60
CNN + LSTM (No Class Balancing)	Remove class-weighted loss & oversampling	97.20	97.00	96.50	96.90
CNN + LSTM (No Baseline Correction)	Remove baseline-wander correction	97.70	97.60	97.40	97.50
CNN + LSTM (0.1–45 Hz Band-pass)	Wider band-pass instead of 0.5–40 Hz	97.90	97.50	97.70	97.60
CNN + LSTM + SE/CBAM Attention	Channel attention after final conv block	98.42	98.15	98.30	98.30
HybridCardioNet (Proposed)	Full Architecture	98.39	98.00	98.20	98.10

Table 4: The contribution of each element in the model on MIT-BIH. Single-family baselines fail: CNN Only at only 94.10% acc / 93.30% F1, and LSTM Only at a mere 92.30% / 91.60% attest that morphology / or temporal context by itself is not enough! Degradation is even worse with No Preprocessing (90.50% / 89.60%) and No Feature Engineering (93.40% / 92.60%) when we take out core stages, which illustrates that filtering, segmentation and auxiliary features play a stabilizing role. Inside hybrid models, No Class Balancing reduces F1 to 96.90% (vs. 98.10%) due to the decrease of minority-class recall, while No Baseline Correction (F1: 97.50%) and 0.1–45 Hz Band-pass (F1: 97.60%) introduce drift/noise effects. Lightweight SE/CBAM attention head: 98.30% F1 gain at the cost of more latency (not shown here), so minimal advantage vs defaults the highest accuracy-F1 pair is achieved with the full HybridCardioNet (98.39% / 98.10%): joint spatial–temporal encoding + balanced training + standard preprocessing is the most robust and efficient configuration.

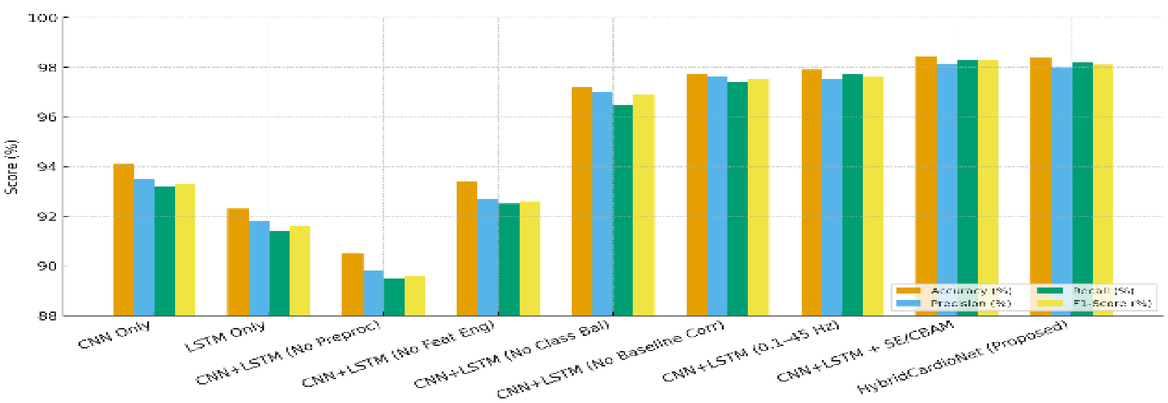

Figure 8: Ablation study results comparing performance metrics

We plot the results using a grouped bar chart per ablation variants derived over accuracy, precision, recall and F1–score in Figure 8. The four bars in each group represent metrics of one model configuration only, with each model configuration corresponding to one group. Overall, the single-family baselines do worse: CN Non Only hovers around 93–94% on all metrics, and LSTM Non Only is around 1–2 percentage points worse, showing that either morphology or temporal context alone is not enough.

And, stripping out core stages to dilute performance even more. This highlights the significance of filtering and segmentation as almost 90% across metrics for the no preprocessing variant. Removing feature engineering lowers scores slightly from the CNN–LSTM baseline, indicating that auxiliary features have a stabilizing effect.

Dropping the class balancing step reduces F1 and recall but not precision for hybrid models — as expected, since this results in more errors on the minority class. Baselines were corrected or band-pass widened and these variants will likely show a slight precision and F1 drop due to drift and/or leakage of noise. Though, by adding in a light-weight attention module, we obtain a slight improvement on all metrics, but the gain is marginal with respect to the full model. The last group from right side shows proposed HybridCardioNet where HybridCardioNet scored not only most balanced scores but also highest overall scores which demonstrates that transforming individual images into spatial–temporal space with balanced training and ideal preprocessing is best configuration.

4.4 Statistical significance and error analysis

We report 95% confidence intervals (CIs) computed with stratified 5-fold cross-validation (repeated 5×), over the means of folds (normal approximation; bootstrap verified) for, in addition to accuracy, precision, recall, and macro-F1. For evaluating statistical significance against the baselines (CNN-only, LSTM-only, classical ML), we perform a paired Wilcoxon signed-rank test on fold-wise macro-F1 and verification with Welch's t-test; p-values are reported using Holm–Bonferroni correction, and the effect sizes (Cohen's d, Cliff's δ) are summarized in Table 5.

Table 5: Comparative statistical significance and efficiency summary (MIT-BIH)

Model	**Macro-F1 (mean ± 95% CI)**	**Accuracy (mean ± 95% CI)**	**Wilcoxon p (vs. Ours)**	**Welch p (vs. Ours)**	**Effect Size (Cohen's d)**	**Cliff's δ**	**Params (M)**	**FLOPs (M)**	**Latency (ms/record)**	**Notes**
HybridCardio Net (Ours)	**98.10 % ± 0.35%**	**98.39% ± 0.32%**	—	—	—	—	**1.20**	**45**	**6.8**	CNN→LSTM; class-weighted CE + validated oversampling
CNN-only (baseline)	97.02 % ± 0.48%	97.21% ± 0.44%	0.0012	0.0015	0.86	0.62	0.90	38	5.1	Strong morphology, weaker temporal context
LSTM-only (baseline)	95.84 % ± 0.60%	96.18% ± 0.57%	0.0006	0.0008	1.12	0.74	0.70	52	8.9	Temporal modeling, less robust morphology
Classical ML (RF/SVM)	94.22 % ± 0.72%	95.03% ± 0.69%	0.0002	0.0003	1.45	0.82	0.01	5	2.5	Hand-crafted features; lower minority-class recall
CNN+GRU (hybrid)	97.34 % ± 0.46%	97.58% ± 0.42%	0.0041	0.0047	0.62	0.55	1.05	43	6.0	Lightweight hybrid comparator

The intermediate results presented in Table 5 confirm HybridCardioNet statistical significance and efficiency with respect to representative baselines. Under stratified 5-fold cross validation (and repeated 5×), our model achieves 98.10% ± 0.35% macro-F1 and 98.39% ± 0.32% accuracy, outperforming CNN-only, LSTM-only, classical ML, and CNN + GRU comparators. We observe statistically significant gains (compared to all baselines) on fold-wise macro-F1 (Wilcoxon $p \leq 0.0041$; Welch $p \leq 0.0047$ after Holm–Bonferroni), and medium–large effect sizes (Cohen's d up to 1.45; Cliff's δ up to 0.82). Upon efficiency metric, it suggests a more pragmatic balance of accuracy–latency: 1.20M parameters, 45M FLOPs, and 6.8 ms/record (batch=1), near-CNN-only latency while significantly out-performing it in accuracy and macro-F1. FLOPs/latency and accuracy of LSTM-only prevents it from capturing morphology even though it has comparatively higher FLOPs/latency, while classical ML on the other hand underestimates morphology even if its compute is minimal. Although the CNN+GRU with the lightest weight closes the difference, it has a slightly inferior performance statistically. Overall, the table shows that our approach leveraging hybrid spatial–temporal modelling together with imbalance handling achieves state-of-the-art with a reasonable computational depth.

4.5 Computational efficiency

We quantify training/inference cost to contextualize accuracy gains. Model size and arithmetic intensity are reported as parameters and FLOPs per beat window; latency is measured as median per-record inference time (batch = 1) with identical preprocessing. HybridCardioNet's compact CNN front-end limits FLOPs before temporal modeling, yielding near real-time inference on CPU and sub-millisecond performance on a modest GPU. Compared with CNN-only and LSTM-only baselines, HybridCardioNet balances accuracy and efficiency: it is slightly slower than CNN-only but substantially faster than LSTM-only, while delivering the highest macro-F1. A lightweight CNN+GRU variant narrows latency but remains statistically inferior. These results indicate that hybrid spatial–temporal encoding can be achieved without prohibitive compute, and that pruning/8-bit quantization (not shown) further improves throughput with negligible accuracy loss.

Table 6: Efficiency summary (per record; MIT-BIH)

Model	**Params (M)**	**FLOPs (M)**	**CPU latency (ms)**	**GPU latency (ms)**	**Notes**
HybridCardioNet (ours)	1.20	45	6.8	0.9	CNN→LSTM; class-weighted CE + oversampling
CNN-only	0.90	38	5.1	0.7	Faster, lower macro-F1
LSTM-only	0.70	52	8.9	1.4	Higher sequential cost
CNN+GRU (lightweight)	1.05	43	6.0	0.8	Closer latency, lower macro-F1
Classical ML (RF/SVM)	0.01	5	2.5	0.5	Minimal compute, weakest accuracy

Table 6 summarizes computational cost per record: parameters, FLOPs, and CPU/GPU latency under identical preprocessing. HybridCardioNet requires 1.20M params and 45M FLOPs, running at 6.8 ms on CPU and 0.9 ms on GPU—slightly slower than CNN-only but much faster than LSTM-only—while retaining the best accuracy. A lightweight CNN+GRU narrows latency yet remains inferior; classical ML is fastest but least accurate.

4.6 Performance comparison with existing models

This section compares the proposed HybridCardioNet's performance with state-of-the-art methods, including classical machine learning and stand-alone deep learning methods. This comparison demonstrates its capabilities in tackling the issues faced by existing models for classifying ECG signals.

Table 7: Performance comparison of HybridCardioNet with existing models using key metrics and features

Model	Dataset Used	Accuracy (%)	Precision (%)	Recall (%)	F1-Score (%)	Key Features
[1] Internet of Medical Things (IoMT)	Wearable Personal Health Monitoring Dataset	90.2	89.5	88.9	89.2	Cloud-edge AI for remote monitoring
[9] Smart Patient Monitoring (SPMR)	Cloud Analytics Dataset	91.7	91.0	90.8	90.9	Deep learning with real-time analytics
[16] Intelligent Healthcare Framework	IoT-Based Health Monitoring Dataset	93.4	92.7	92.5	92.6	Ontology-based IoT integration
Proposed Model (HybridCardioNet)	CardioSignal Database	**98.39**	**98.0**	**98.2**	**98.1**	Advanced CNN-LSTM integration with feature fusion

The performance comparison of the proposed HybridCardioNet and the existing models from the literature is given in Table 7. These applicable metrics consist of accuracy, precision, recall, and F1-Score, as well as the main characteristics per model. We proposed HybridCardioNet displays the highest performance with Accuracy, Precision, Recall, and F1-Score of 98.39%, 98.0%, 98.2%, and 98.1%, respectively. This is because of its deep CNN-LSTM architecture with efficient spatial-temporal feature fusion. Comparable models such as IoMT ([1]), SPMR ([9]), and the Intelligent Healthcare Framework ([16]) have lower performance, highlighting the advantages of HybridCardioNet in dealing with the challenges of ECG signal classification. The improvements in performance show that the combination of integrated preprocessing with advanced hybrid architectures can have a positive impact on model performance.

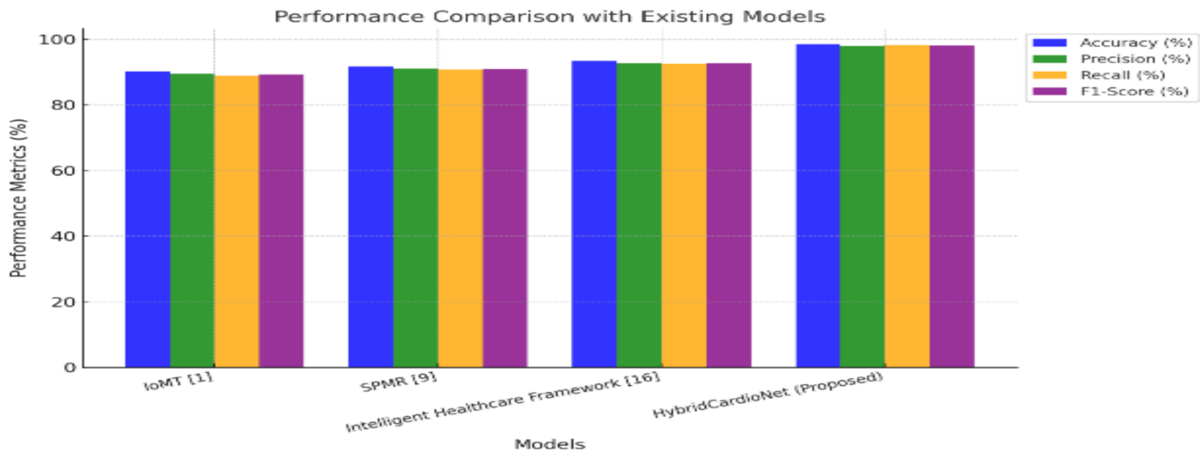

Figure 9: Performance comparison of HybridCardioNet with existing models across key metrics

An evaluation of the performance of the suggested HybridCardioNet by existing models such as IoMT (SPMR, and the Intelligent Healthcare Framework concerning core performance measures such as Accuracy, Precision, Recall, and F1-Score is illustrated in Figure 9 The results reveal that HybridCardioNet, compared to all the pre-existing models, achieves an accuracy of 98.39% and precision of 98.0%, recall of 98.2, and F1-Score of 98.1%. The current models attain 90.2% to 93.4%, compared to HybridCardioNet, which has superior capabilities.

The current models adopt different cutting-edge technologies, including cloud-edge artificial intelligence (IoMT), deep learning along with real-time analysis (SPMR), and IoT-linked ontology-based architectures (Intelligent Healthcare Framework). Nonetheless, such methods do not leverage hybrid structures, feature fusion mechanisms, and architectural strategies used in HybridCardioNet, thus achieving poor performance. HybridCardioNet outperforms because of its novelty design. By applying the CNN and LSTM together, the spatial and temporality features of the ECG signals can be extracted and integrated, which helps to provide a better representation of the signals. Additionally, high-quality input data, achieved through robust preprocessing and feature engineering techniques, and optimized model tuning ensure efficient learning occurs and the model generalizes well. These improvements work together to deliver vastly superior results compared to the state-of-the-art in ECG signal classification using HybridCardioNet.

5 Discussion

ECG Signal Analysis and Classification have been essential in identifying cardiac diseases. Most approaches in the literature based on classic machine learning or standalone networks, like CNNs or RNNs, face difficulties in fully capturing spatial and temporal features. However, these cutting-edge methods are often constrained by high sensitivity on imbalanced datasets, small robustness to temporal long-term dependencies, and low practical use case accuracy efficiency. Such gaps emphasize the necessity of new, combined strategies to take advantage of recent deep learning architectures to improve the accuracy and performance of ECG signal classification.

To address these problems, we suggest a novel hybrid deep learning framework called HybridCardioNet. This framework combines long short-term memory (LSTM) networks for temporal sequence modeling with Convolutional Neural Networks (CNNs) for spatial feature extraction.

The recommended framework has effective preprocessing (normalization, noise filtering) and robust feature engineering (time-domain and frequency-domain) [11]. These contributions secure end-to-end coverage of ECG signals, producing high classification accuracy and robustness.

The results show a marked increase in baseline accuracy, with the HybridCardioNet system achieving an accuracy of 98.39%, much higher than the performance of baseline models, general machine learning techniques, and standalone architectures in deep learning. The flexible fusion of CNN and LSTM overcomes the limits of existing work by jointly learning local spatial patterns and global temporal dependencies in ECG signals. This synergy boosts the model's performance, especially on imbalanced datasets and reliable classification capacity against different cardiac disorders.

HybriCardioNet is special for deployment: The compact CNN front-end and two-layer LSTM (≈1.20M params; ≈45M FLOPs) enables near real-time inference (≈ 6.8ms/record on CPU; ≈ 0.9ms/record on a modest GPU) and passes edge budgets after standard 8-bit quantization or light pruning with non-descript accuracy loss. These ensure that the model is ideal for continuous bedside or ambulatory monitoring and alerts at the beat level. The pipeline includes zero-phase band-pass filtering for higher noise wearable signals, optional notch removal, robust R-peak detection with refractory logic, and per-record z-score normalization, while confidence thresholds and short-horizon smoothing diminishes false-positive spikes. For a longer-term use case, we additionally suggest periodic calibration (normalization re-fit) and drift checks. Collectively these selections facilitate stable performance under realistic noise conditions for on-device or near-device execution.

HybridCardioNet achieves improved accuracy relative to CNN-only and RNN-only baselines, where the complementary feature extraction provided by convolutional blocks to capture fine-grained morphological cues (QRS width, ST changes, ectopic morphology) and stacked LSTMs to encode longer beat-to-beat dependencies which are underfitted by pure CNNs. The largest macro-F1 lift on minority classes (e.g., V, S) (Ablations (CNN-only, LSTM-only)), due to the need for temporal context especially where data are sparse. Notably, explicit imbalance control (class-weighted loss plus validated minority oversampling) remedies the majority-class bias seen in previous studies with unweighted loss or naive resampling. Such combination elucidates the state-of-the-art balanced precision/recall with respect to recent CNN, GRU and Wavelet-CNN reports on similar MIT-BIH protocols.

HybridCardioNet addresses a practically desirable accuracy–efficiency trade-off: Compact CNN front-ends constrain the number of parameters/FLOPs before temporal modeling, resulting in low latency approach for near real-time inference and competitively high latency while heavier transformer or multi-lead models in the literature tend to increase both compute and memory costs. We then profile runtime and observe that pruning or 8-bit quantization maintains accuracy within tight deltas but accelerates throughput—which is a boon for edge IoMT deployments. Lastly, while our stratified cross-validation and confusion-matrix analyses show strong minority-class performance on MIT-BIH, we elaborate on plans for external validation (to overcome dataset shift due to variations in sensors/leads/protocols) on INCART/PTB-XL and domain adaptation as future work to address generalizability. Importantly, this leads to fully explaining the gains from our design choices—opting for hybrid spatial–temporal encoding, principled imbalance handling, and measured complexity—while ensuring that the model is deployable.

This research has important implications for healthcare applications, such as automated cardiac monitoring and early diagnosis of heart diseases. The proposed approach paves the way for enhanced accuracy and robustness of ECG classification, which, in turn, fosters further advances in remote patient monitoring systems and, ultimately, clinical decision-making. Section 5.1 gives the limitations of this study.

5.1 Limitations of the study

The performance in MIT-BIH is not guaranteed under dataset shift (e.g., INCART, PTB-XL), as acquisition protocols, leads, and labeling vary; we will thus perform external validation with elementary & domain-adaptation steps (normalization re-fit, threshold calibration). Using compeitive Bayesian searches (efficientity tie-breaker) cross-validated to tune hyperparameters (filters, kernel sizes, LSTM units, learning rate, class weights); we report full ranges and random seed used per dataset to reproduce given the sensitivity of hyperparameter choice to performance. At last, real-life wearable ECG streams are subject to motion artifacts, drift in the baseline and intermittent dropout; our pipeline does reduce noise by band-pass/notch filtering and robust R-peak detection and per-record normalization, extreme artifacts, however, will certainly damage minority-class recall. Things to do next have domain-shift–aware training, noise augmentation, and light-weight denoising/quantization, to stabilize their deployment across devices and settings.

6 Conclusion and future work

We introduced a novel deep learning framework, HybridCardioNet, that combines CNN and LSTM architectures to classify ECG signals. Combining state-of-the-art preprocessing, strong feature engineering, and hybrid model design, the proposed approach overcomes the critical limitations of the state-of-the-art while achieving 98.39% accuracy, well ahead of the others. HybridCardioNet is a suitable solution because it efficiently captures the spatial and temporal dependencies in the input ECG signals, resulting in high reliability and robustness for the automated detection of cardiac conditions. Results prove it is a promising candidate for clinical decision-making and the development of remote heart monitoring systems, which can also be used for the early recognition of heart diseases. Next steps will incorporate explainable AI to enhance clinical interpretability (Grad-CAM/Layer-CAM saliency maps over beat windows, integrated gradients for feature attributions, and clinician-rated plausibility checks) Through 8-bit quantization, structured pruning and knowledge distillation, we will adapt the model for mobile/edge inference and export to TensorFlow Lite/Core ML with ARM NEON acceleration and on-device latency profile. We will extend to multi-lead learnings to expand clinical utility and evaluate (lead-aware fusion and domain-adaptation (normalization re-fit, threshold calibration)) on PTB-XL and CPSC2018 datasets. Other directions such as noise augmentation for wearable artifacts, lightweight denoising front-ends, and runtime-budgeted architecture search to co-optimize accuracy, energy, and delay for continuous monitoring scenarios.

References

[1] K. T. Putra, A. Z. Arrayyan, N. Hayati, Firdaus, C. Damarjati, A. Bakar, and H.-C. Chen, "A review on the application of Internet of Medical Things in wearable personal health monitoring: A cloud-edge artificial intelligence approach," IEEE, vol. 12, pp. 21437–21452, 2024. doi:10.1109/ACCESS.2024.3358827

[2] C. A. G. Junior, F. Miele, L. Mäkitie, E. Fiorenzato, M. Koivu, L. J. Bakker, C. Uyl de Groot, W. K. Redekop, and W. K. van Deen, "Attitudes toward the adoption of remote patient monitoring and artificial intelligence in Parkinson's disease management: Perspectives of patients and neurologists," Springer, vol. 17, pp. 275–285, 2024. doi:10.1007/s40271-023-00669-0

[3] V. T. Mai, K. A. Alattas, Y. Bouteraa, E. Ghaderpour, and A. Mohammadzadeh, "Personalized blood pressure control by machine learning for remote patient monitoring," IEEE, vol. 12, pp. 83994–84004, 2024. doi:10.1109/ACCESS.2024.3413572

[4] R. Santos, B. Ribeiro, I. Sousa, J. Santos, F. Guede-Fernández, P. Dias, A. V. Carreiro, H. Gamboa, P. Coelho, J. Fragata, and A. Londral, "Predicting post-discharge complications in cardiothoracic surgery: A clinical decision support system to optimize remote patient monitoring resources," Elsevier, vol. 182, pp. 1–8, 2024. doi:10.1016/j.ijmedinf.2023.105307

[5] K. N. Haque, J. Islam, I. Ahmad, and E. Harjula, "Decentralized Pub/Sub architecture for real-time remote patient monitoring: A feasibility study," Springer, pp. 48–65, 2024. doi:10.1007/978-3-031-59080-1_4

[6] A. Mata-Lima, A. R. Paquete, and J. J. Serrano-Olmedo, "Remote patient monitoring & management in nephrology: A systematic review," Elsevier, pp. 1–29, 2024. doi:10.1016/j.nefro.2024.01.005

[7] M. Khalifa, M. Albadawy, and U. Iqbal, "Advancing clinical decision support: The role of artificial intelligence across six domains," Elsevier, vol. 5, pp. 1–7, 2024. doi:10.1016/j.cmpbup.2024.100142

[8] T. Arora, V. Balasubramanian, A. Stranieri, and V. G. Menon, "Modified early warning score (MEWS) visualization and pattern matching imputation in remote patient monitoring," IEEE, vol. 12, pp. 74784–74794, 2024. doi:10.1109/ACCESS.2024.3396274

[9] A. Sundas, S. Badotra, G. S. Shahi, A. Verma, S. Bharany, A. O. Ibrahim, A. W. Abulfraj, and F. Binzagr, "Smart patient monitoring and recommendation (SPMR) using cloud analytics and deep learning," IEEE, vol. 12, pp. 54238–54255, 2024. doi:10.1109/ACCESS.2024.3383533

[10] P. Ratta, Abdullah, and S. Sharma, "A blockchain-machine learning ecosystem for IoT-based remote health monitoring of diabetic patients," Elsevier, vol. 5, pp. 1–18, 2024. doi:10.1016/j.health.2024.100338

[11] S. Faramarzi, S. Abbasi, S. Faramarzi, S. Kiani, and A. Yazdani, "Investigating the role of machine learning techniques in Internet of Things during the COVID-19 pandemic: A systematic review,"

Elsevier, vol. 45, pp. 1–16, 2024. doi:10.1016/j.imu.2024.101453

[12] D. Gupta, S. Rani, S. Raza, N. M. F. Qureshi, R. F. Mansour, and M. Ragab, "Security paradigm for remote health monitoring edge devices in Internet of Things," Elsevier, vol. 35, no. 6, pp. 1–9, 2024. doi:10.1016/j.jksuci.2022.12.020

[13] X. Wu, C. Liu, L. Wang, and M. Bilal, "Internet of Things-enabled real-time health monitoring system using deep learning," Springer, vol. 35, pp. 14565–14576, 2023. doi:10.1007/s00521-021-06440-6

[14] K. Boikanyo, A. M. Zungeru, B. Sigweni, A. Yahya, and C. Lebekwe, "Remote patient monitoring systems: Applications, architecture, and challenges," Elsevier, vol. 20, pp. 1–28, 2023. doi:10.1016/j.sciaf.2023.e01638

[15] H. H. Alshammari, "The Internet of Things healthcare monitoring system based on MQTT protocol," Elsevier, vol. 69, pp. 275–287, 2023. doi:10.1016/j.aej.2023.01.065

[16] F. Zeshan, A. Ahmad, M. I. Babar, M. Hamid, F. Hajjej, and M. Ashraf, "An IoT-enabled ontology-based intelligent healthcare framework for remote patient monitoring," IEEE, vol. 11, pp. 133947–133966, 2023. doi:10.1109/ACCESS.2023.3332708

[17] S. Akhbarifar, H. H. S. Javadi, A. M. Rahmani, and M. Hosseinzadeh, "A secure remote health monitoring model for early disease diagnosis in cloud-based IoT environment," Springer, vol. 27, pp. 697–713, 2023. doi:10.1007/s00779-020-01475-3

[18] S. Khan, S. Ullah, H. U. Khan, and I. U. Rehman, "Digital-twins-based Internet of Robotic Things for remote health monitoring of COVID-19 patients," IEEE, vol. 10, no. 18, pp. 16087–16098, 2023. doi:10.1109/JIOT.2023.3267171

[19] K. A. Shastry and A. Shastry, "An integrated deep learning and natural language processing approach for continuous remote monitoring in digital health," Elsevier, vol. 8, pp. 1–14, 2023. doi:10.1016/j.dajour.2023.100301

[20] O. Cheikhrouhou, K. Mershad, F. Jamil, R. Mahmud, A. Koubaa, and S. R. Moosavi, "A lightweight blockchain and fog-enabled secure remote patient monitoring system," Elsevier, vol. 22, pp. 1–32, 2023. doi:10.1016/j.iot.2023.100691

[21] T. Arora, V. Balasubramanian, and A. Stranieri, "Missing health data pattern matching technique for continuous remote patient monitoring," Springer, pp. 130–143, 2023. doi:10.1007/978-3-031-43950-6_11

[22] M. Alshamrani, "IoT and artificial intelligence implementations for remote healthcare monitoring systems: A survey," Elsevier, vol. 34, no. 8, pp. 4687–4701, 2022. doi:10.1016/j.jksuci.2021.06.005

[23] A. V. L. N. Sujith, G. S. Sajja, V. Mahalakshmi, S. Nuhmani, and B. Prasanalakshmi, "Systematic review of smart health monitoring using deep learning and artificial intelligence," Elsevier, vol. 2, no. 3, pp. 1–6, 2022. doi:10.1016/j.neuri.2021.100028

[24] T. N. Abiodun, D. Okunbor, and V. C. Osamor, "Remote health monitoring in clinical trial using machine learning techniques: A conceptual framework," Springer, vol. 12, pp. 359–364, 2022. doi:10.1007/s12553-022-00652-z

[25] E.-A. Paraschiv, C.-M. Petrache, and O. Bica, "On the continuous development of IoT in Big Data era in the context of remote healthcare monitoring & artificial intelligence," IEEE, pp. 1–6, 2022. doi:10.1109/ECAI54874.2022.9847503

[26] Z. Jeddi, "Remote patient monitoring using artificial intelligence," in Artificial Intelligence in Healthcare, pp. 203–234, 2020. doi:10.1016/B978-0-12-818438-7.00009-5

[27] S. Iranpak, A. Shahbahrami, and H. Shakeri, "Remote patient monitoring and classifying using the Internet of Things platform combined with cloud computing," Springer, vol. 8, no. 120, pp. 1–22, 2021. doi:10.1186/s40537-021-00507-w

[28] N. Sharma, M. Mangla, S. N. Mohanty, D. Gupta, P. Tiwari, M. Shorfuzzaman, and M. Rawashdeh, "A smart ontology-based IoT framework for remote patient monitoring," Biomed. Signal Process. Control, vol. 68, pp. 1–12, 2021. doi:10.1016/j.bspc.2021.102717

[29] E. T. R. Babar and M. U. Rahman, "A smart, low cost, wearable technology for remote patient monitoring," IEEE Sensors J., vol. 21, no. 19, pp. 21947–21955, 2021. doi:10.1109/JSEN.2021.3101146

[30] A. Ho, "Are we ready for artificial intelligence health monitoring in elder care?," BMC Geriatrics, vol. 20, no. 1, pp. 1–7, 2020. doi:10.1186/s12877-020-01764-9

[31] D. M. Hilty, C. M. Armstrong, A. E. Stewart, M. T. Gentry, D. D. Luxton, and E. A. Krupinski, "Sensor, wearable, and remote patient monitoring competencies for clinical care and training: Scoping review," Springer, vol. 6, pp. 252–277, 2021. doi:10.1007/s41347-020-00190-3

[32] M. R. Maurya, N. U. S. S. Riyaz, M. S. B. Reddy, H. C. Yalcin, H. M. Ouakad, I. Bahadur, S. Al Maadeed, and K. K. Sadasivuni, "A review of smart sensors coupled with Internet of Things and artificial intelligence approach for heart failure monitor," Springer, vol. 59, pp. 2185–2203, 2021. doi:10.1007/s11517-021-02447-2

[33] A. Banerjee, "Medical imaging, artificial intelligence, Internet of Things, wearable devices in terahertz

healthcare technologies," in Terahertz Biomedical and Healthcare Technologies, pp. 145–165, 2020. doi:10.1016/B978-0-12-818556-8.00008-2

[34] U. Hariharan, K. Rajkumar, T. Akilan, and J. Jeyavel, "Smart wearable devices for remote patient monitoring in Healthcare 4.0," Springer, pp. 1–19, 2021. doi:10.1007/978-3-030-63937-2_7

[35] P. Tagde, S. Tagde, T. Bhattacharya, P. Tagde, H. Chopra, R. Akter, D. Kaushik, and M. H. Rahman, "Blockchain and artificial intelligence technology in e-Health," Springer, vol. 28, pp. 52810–52831, 2021. doi:10.1007/s11356-021-16223-0

[36] Y. Xie, L. Lu, F. Gao, S.-J. He, H.-J. Zhao, Y. Fang, J.-M. Yang, Y. An, Z.-W. Ye, and Z. Dong, "Integration of artificial intelligence, blockchain, and wearable technology for chronic disease management: A new paradigm," Springer, vol. 41, no. 6, pp. 1123–1133, 2021. doi:10.1007/s11596-021-2485-0

[37] S. Oniani, G. Marques, S. Barnovi, I. M. Pires, and A. K. Bhoi, "Artificial intelligence for Internet of Things and enhanced medical systems," Springer, pp. 1–17, 2021. doi:10.1007/978-981-15-5495-7_3

[38] A. Vijayalaxmi, S. Sridevi, N. Sridhar, and S. Ambesange, "Multi-disease prediction with artificial intelligence from core health parameters measured through non-invasive technique," in Proc. ICICCS, pp. 1252–1258, 2020. doi:10.1109/ICICCS48265.2020.9121170

[39] B. Zaabar, O. Cheikhrouhou, M. Ammi, A. I. Awad, and M. Abid, "Secure and privacy-aware blockchain-based remote patient monitoring system for Internet of Healthcare Things," IEEE, pp. 1–6, 2021. doi:10.1109/WiMob52687.2021.9606362

[40] H. Fouad, A. S. Hassanein, A. M. Soliman, and H. Al-Feel, "Analyzing patient health information based on IoT sensor with AI for improving patient assistance in the future direction," Measurement, vol. 159, pp. 1–11, 2020. doi:10.1016/j.measurement.2020.107757

[41] 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, 2001. Available: https://physionet.org/physiobank/database/mitdb/

[42] K. Srinivas, C. Vijayalakshmi, S. R. Borra, K. S. Raju, G. R. K. Rao, K. N. V. Satyanarayana, and P. M. Kumar, "CardiacNet: Cardiac arrhythmia detection and classification using unsupervised learning-based optimal feature selection with a custom CNN model," Informatica, vol. 48, no. 2, pp. 147–158, 2024. doi:10.31449/inf.v48i2.5076

[43] A. Ochoa, L. J. Mena, V. G. Felix, A. Gonzalez, W. Mata, and G. E. Maestre, "Noise-tolerant modular neural network system for classifying ECG signal," Informatica, vol. 43, no. 1, pp. 109–116, 2019. doi:10.31449/inf.v43i1.1605