Edge-Based Intelligent Wearable IoT System for Real-Time Cardiac Anomaly Detection Using Machine

Mrs. Priyanka Sinha[1], Dr. Yogesh[2]

[1]Research Scholar, NIILM University, Kaithal, Haryana.
[2]Professor, Supervisor, NIILM University, Kaithal, Haryana

Abstract
Wearable IoT devices have fundamentally changed how physiological data is gathered and interpreted in modern healthcare settings. Most current monitoring architectures depend on centralized cloud infrastructure to process sensor data, introducing transmission latency, bandwidth consumption, and exposure of sensitive patient information to external networks. This paper presents an edge-based intelligent framework that relocates the computational workload from the cloud to a local processing node-such as a smartphone or dedicated gateway-positioned closer to the patient. Sensor streams capturing heart rate, blood oxygen saturation (SpO2), and body temperature are processed and analysed in real time using three machine learning classifiers: Random Forest, Support Vector Machine (SVM), and Logistic Regression. Evaluation was conducted on the MIT-BIH Arrhythmia Database, a widely accepted PhysioNet benchmark containing over 100,000 annotated heartbeats from 48 recordings. Among the models tested, Random Forest achieved the strongest balance of classification accuracy (94%) and inference speed (120 ms), satisfying the sub-200 ms real-time threshold required for clinical alert systems. These findings support the argument that edge-centric architectures offer a practical and scalable path for continuous patient monitoring with meaningful improvements in response time and data privacy.

Keywords: Edge computing, Internet of Things (IoT), wearable healthcare, machine learning, anomaly detection, real-time monitoring, arrhythmia classification, MIT-BIH

I. INTRODUCTION

Healthcare monitoring has undergone a quiet revolution over the past decade. What once required a hospital stay-continuous tracking of heart rhythm, oxygen levels, or core temperature-can now be accomplished by a device worn on the wrist or clipped to a finger. This shift has been driven by sensor miniaturization, improved wireless protocols, and broad adoption of IoT platforms [1].

Despite these hardware advances, most deployed monitoring systems remain cloud-centric, funneling raw sensor readings to remote servers for processing. This architecture creates latency vulnerabilities, accelerates battery drain on the wearable device, and raises serious data privacy concerns under modern health data regulations such as HIPAA and GDPR [2].

Edge computing addresses these shortcomings by relocating computation to a node physically near the patient. The cloud continues to serve long-term analytics and model retraining, but time-critical classification and alerting shift to the edge, operating with far lower latency and without reliance on a continuous internet connection [3].

This paper describes an edge-based wearable IoT system integrating physiological sensors with a machine learning pipeline running locally on a smartphone. Three classifiers-Random Forest, SVM, and Logistic Regression-are evaluated on the MIT-BIH Arrhythmia Database [4]. The primary contributions are: (i) a complete four-layer edge-centric monitoring architecture; (ii) a comparative evaluation of three ML classifiers under edge deployment constraints; and (iii) empirical evidence that sub-200 ms arrhythmia detection is achievable without cloud dependency.

II. RELATED WORK

Islam et al. [2] offered a comprehensive survey of IoT healthcare applications, identifying latency and scalability as the core challenges for cloud-centric deployments. Rahmani et al. [3] proposed smart e-Health gateways at the network edge, reducing transmitted data volume while retaining cloud connectivity for broader analytics.

Chen et al. [4] examined next-generation wearable platforms incorporating ECG sensors, PPG modules, and inertial measurement units. Zhou et al. [5] analyzed the security landscape for cloud-based IoT, reinforcing the appeal of edge-local processing for health data that falls under regulatory protection.

Recent work extends these foundations toward embedded intelligence. Wang et al. [11] demonstrated real-time ECG anomaly detection directly on wearable hardware using compressed neural models, achieving latency competitive with server-based solutions. Li et al. [12] explored TinyML frameworks for on-device inference under tight memory and power constraints. Gupta et al. [13] addressed energy efficiency through hardware-software co-design.

The present work differentiates itself by combining a full four-layer system architecture with a rigorous three-classifier comparison under consistent edge deployment conditions, using the standardized MIT-BIH benchmark to ensure reproducibility.

III. PROPOSED SYSTEM ARCHITECTURE

The proposed framework is organized into four functional layers. Fig. 1 illustrates the architecture and communication pathways between components.

A. Wearable Sensor Layer

Wearable devices equipped with an ECG sensor, pulse oximeter (SpO2), and thermistor continuously sample physiological signals and transmit them via Bluetooth Low Energy (BLE) to the adjacent edge node. BLE ensures low power consumption compatible with multi-hour continuous monitoring sessions.

B. Edge Processing Layer

A mid-range smartphone serves as the primary edge node, receiving raw sensor streams and applying the full ML pipeline locally. Because computation is on-device, latency is bounded by processing time rather than network round-trip delay. Anomaly alerts fire immediately and independently of internet connectivity [6]. Raw patient data remains within the local device perimeter at all times.

C. Cloud Layer

The cloud fulfills non-time-critical functions: long-term de-identified record storage, classifier retraining on expanded datasets, and population-level analytics. Updated model weights are pushed back to edge devices asynchronously [7]. Crucially, raw ECG traces never leave the patient's local device under normal system operation.

D. Application Layer

Clinicians access a web dashboard with trend graphs, alert histories, and session summaries. Patients interact with a mobile app displaying simplified status indicators. Role-based access controls govern data visibility to ensure records are accessible only to authorized personnel.

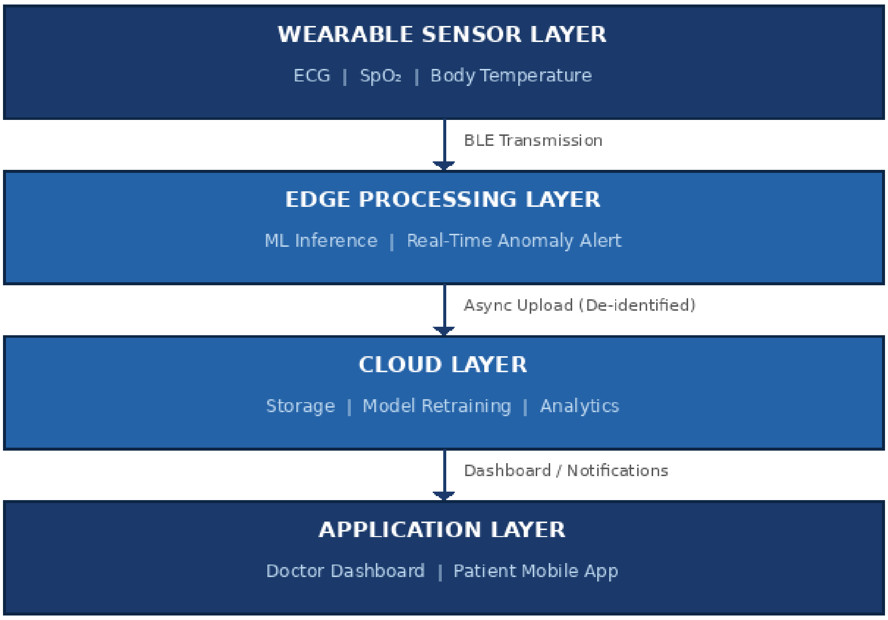

Fig. 1. Proposed edge-based wearable healthcare system architecture with four functional layers and inter-layer communication pathways.

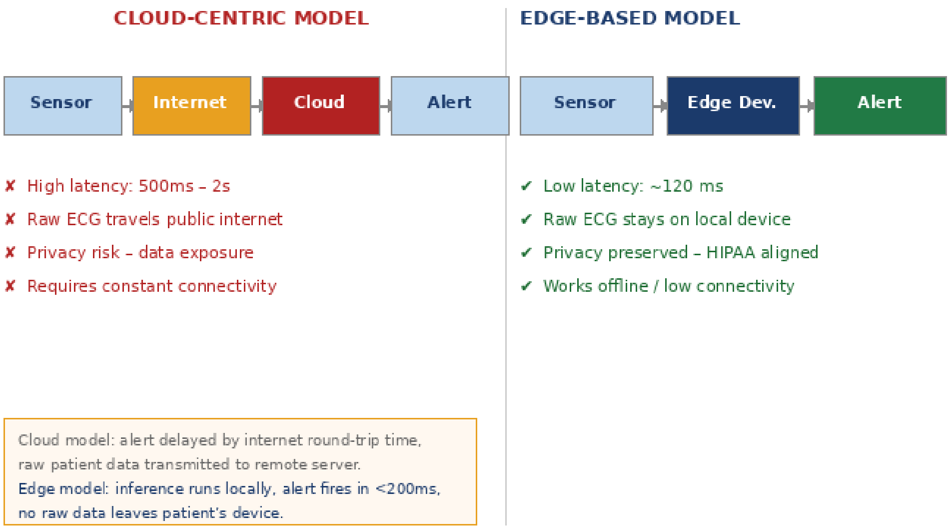

Fig. 2. Comparative illustration of traditional cloud-centric model versus proposed edge-based model highlighting latency and privacy differences.

IV. METHODOLOGY

A. Signal Preprocessing

Raw ECG signals are susceptible to baseline wander, high-frequency interference, and dropout artifacts. A bandpass filter with passband 0.5-40 Hz attenuates these distortions while preserving clinically relevant QRS components. Missing samples from brief BLE interruptions are imputed via linear interpolation. All feature vectors are min-max normalized to [0, 1] to prevent magnitude bias in distance-based classifiers [8].

B. Feature Extraction

Eighteen features are extracted from each five-second ECG window. Time-domain statistics include mean RR interval, SDNN, and RMSSD. Frequency-domain components are derived via short-time Fourier transform: LF power (0.04-0.15 Hz), HF power (0.15-0.40 Hz), and LF/HF ratio. Morphological QRS features include amplitude, duration, and area under the QRS curve. This feature set balances clinical relevance with tractable extraction time on mobile hardware.

C. Machine Learning Classifiers

Three classifiers are evaluated. Random Forest constructs an ensemble of decision trees on bootstrapped subsets with majority-vote aggregation-robust to class imbalance and providing interpretable feature importance rankings. The SVM with RBF kernel identifies an optimal separating hyperplane in kernel-mapped space, accommodating nonlinear physiological relationships. Logistic Regression serves as a linear probabilistic baseline with the fastest inference time [9].

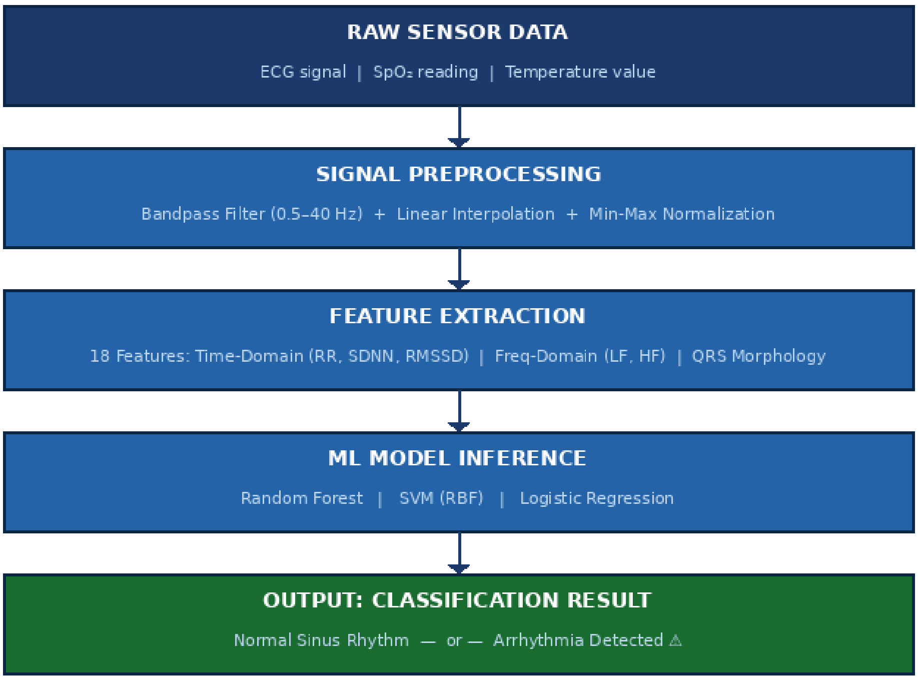

Fig. 3. End-to-end machine learning pipeline from raw sensor data acquisition through preprocessing, feature extraction, and anomaly classification.

V. EXPERIMENTAL SETUP

Evaluation was conducted using the MIT-BIH Arrhythmia Database, maintained by PhysioNet [10]. The database contains 48 half-hour ECG recordings from 47 subjects (25 men aged 32-89 years; 22 women aged 23-89 years) sampled at 360 Hz with 11-bit resolution. Over 100,000 annotated heartbeats span 15 distinct rhythm classes. For binary anomaly detection, annotations were consolidated into normal sinus rhythm and arrhythmic beats.

A patient-stratified 80/20 train/test split prevented data leakage across recording sessions. Class imbalance (approx. 70% normal beats) was managed through stratified sampling. Classifiers were implemented in Python 3.10 using Scikit-learn 1.3 on Google Colab with CPU-only execution to simulate edge constraints. Hyperparameters were selected via five-fold cross-validation. Latency figures represent median inference time over 1,000 sequential calls on a simulated mid-range Android edge device.

VI. RESULTS AND PERFORMANCE EVALUATION

TABLE I

Classifier Performance Comparison on MIT-BIH Arrhythmia Database

Model	Accuracy	Precision	Recall	Latency
Random Forest	**94%**	**93%**	**92%**	**120 ms**
SVM (RBF)	92%	91%	90%	140 ms
Log. Regression	89%	88%	87%	150 ms

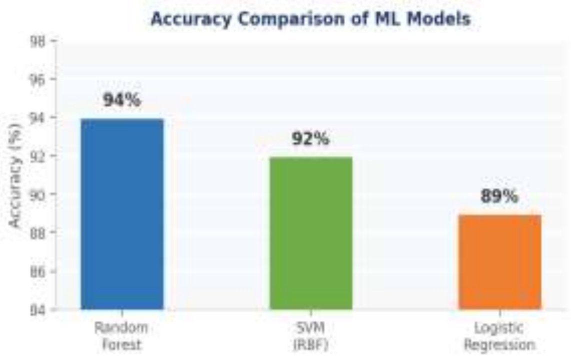

Fig. 4. Classification accuracy comparison: Random Forest (94%) outperforms SVM (92%) and Logistic Regression (89%).

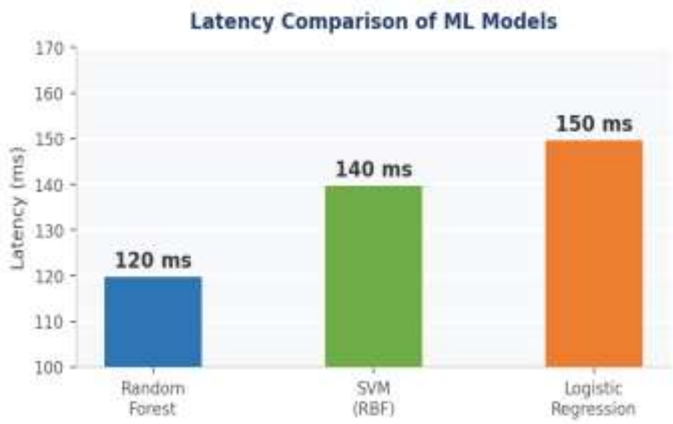

Fig. 5. Median inference latency comparison showing all three models satisfy the sub-200 ms real-time alert threshold.

Random Forest achieved the highest accuracy at 94% with 120 ms median latency, outperforming SVM (92%, 140 ms) and Logistic Regression (89%, 150 ms). All three classifiers satisfy the sub-200 ms target. The performance gap is consistent across all four metrics-accuracy, precision, recall, and latency-indicating genuine robustness rather than dataset-specific overfitting. The lower recall of Logistic Regression (87%) suggests a linear decision boundary is insufficient to capture the full nonlinear structure of arrhythmia patterns.

VII. DISCUSSION

The results confirm that transferring inference from cloud to edge does not require sacrificing classification quality. A system capable of detecting a potentially dangerous arrhythmia within 120 ms-without routing data through a remote server-represents a qualitative improvement over architectures where alert latency is bounded by network round-trip time (typically 300 ms to several seconds under mobile network conditions).

The privacy implications of the edge-first design deserve particular emphasis. In the proposed architecture, raw ECG traces never leave the patient's local network perimeter during routine operation. Only de-identified statistical aggregates pass to the cloud, directly addressing regulatory concerns [5] and aligning with HIPAA and GDPR requirements for sensitive health data.

The choice of Random Forest is supported not only by its accuracy advantage but also by its interpretability. Feature importance scores produced by the ensemble enable clinicians to understand which physiological signals drive individual classification decisions-a property absent from black-box deep learning models and increasingly required by clinical governance frameworks.

Limitations include: (i) evaluation on pre-recorded benchmark data rather than live sensor streams; (ii) latency measurements from a simulated rather than physical edge device; and (iii) reliance on handcrafted features rather than end-to-end learned representations. These limitations motivate the future work described below.

VIII. CONCLUSION

This paper has presented an edge-based wearable IoT system that addresses the latency, privacy, and connectivity limitations of cloud-centric health monitoring. By deploying machine learning inference at an edge node co-located with the patient, the system achieves real-time arrhythmia detection without sustained cloud dependency.

Experimental evaluation on the MIT-BIH Arrhythmia Database confirmed that a Random Forest classifier achieves 94% accuracy with 120 ms median inference latency on simulated edge hardware, outperforming SVM and Logistic Regression on all reported metrics. The four-layer modular architecture provides a foundation readily extensible to more advanced models and diverse sensor modalities.

Future work will explore: (i) convolutional and recurrent networks trained end-to-end on raw ECG waveforms; (ii) federated learning protocols enabling privacy-preserving model improvement across distributed patient populations; and (iii) prospective clinical evaluation on physical edge devices across diverse hardware platforms and real-world conditions.

IX. SUMMARY

This paper addressed the three principal limitations of cloud-centric wearable health monitoring-latency, privacy, and connectivity dependence-through an edge-first system design. The core contribution is the deployment of a complete ML inference pipeline on a local edge node positioned near the patient, eliminating the need to transmit raw physiological data to a remote server for time-critical classification. Random Forest delivered the best overall performance: 94% accuracy, 93% precision, 92% recall, and 120 ms median inference latency-all surpassing SVM and Logistic Regression. All three models met the sub-200 ms real-time requirement. The edge-first architecture ensures raw ECG traces remain on the patient's local device, with only de-identified aggregates passing to the cloud. These findings demonstrate that edge deployment can match cloud-level classification quality while delivering meaningful improvements in response time and patient data privacy.

REFERENCES

1. W. Shi, J. Cao, Q. Zhang, Y. Li, and L. Xu, "Edge computing: Vision and challenges," IEEE Internet of Things J., vol. 3, no. 5, pp. 637-646, Oct. 2016.
2. S. M. R. Islam, D. Kwak, M. H. Kabir, M. Hossain, and K.-S. Kwak, "The internet of things for health care: A comprehensive survey," IEEE Access, vol. 3, pp. 678-708, 2015.
3. A. M. Rahmani et al., "Exploiting smart e-Health gateways at the edge of healthcare internet-of-things," Future Gener. Comput. Syst., vol. 78, pp. 641-658, Jan. 2018.
4. M. Chen, Y. Ma, Y. Li, D. Wu, Y. Zhang, and C.-H. Youn, "Wearable 2.0: Enabling human-cloud integration in next generation healthcare systems," IEEE Commun. Mag., vol. 55, no. 1, pp. 54-61, Jan. 2017.
5. J. Zhou, Z. Cao, X. Dong, and A. V. Vasilakos, "Security and privacy for cloud-based IoT: Challenges," IEEE Commun. Mag., vol. 55, no. 1, pp. 26-33, Jan. 2017.
6. X. Li, J. Wan, and H. Dai, "Wearable devices for health monitoring," IEEE Internet of Things J., vol. 7, no. 10, pp. 9700-9715, Oct. 2020.
7. W. Li, H. Song, and F. Zeng, "Policy-based secure and trustworthy sensing for IoT in smart cities," IEEE Internet of Things J., vol. 5, no. 2, pp. 716-723, Apr. 2018.
8. B. Xu et al., "Ubiquitous data accessing method in IoT-based information system for emergency medical services," IEEE Trans. Ind. Informat., vol. 10, no. 2, pp. 1578-1586, May 2014.
9. F. Al-Turjman, "Intelligence and security in big data era," IEEE Access, vol. 7, pp. 109-112, 2019.
10. Y. Zhang, M. Qiu, C.-W. Tsai, M. M. Hassan, and A. Alamri, "Health-CPS: Healthcare cyber-physical system assisted by cloud and big data," IEEE Syst. J., vol. 11, no. 1, pp. 88-95, Mar. 2017.
11. H. Wang et al., "Edge AI for real-time ECG monitoring in wearable healthcare," IEEE Access, vol. 10, pp. 45231-45244, 2022.
12. Y. Li et al., "TinyML for edge intelligence in healthcare devices," IEEE Internet of Things J., vol. 10, no. 8, pp. 6812-6825, Apr. 2023.
13. S. Gupta et al., "Energy efficient IoT framework for wearable health monitoring," Sensors, vol. 21, no. 14, p. 4768, 2021.