IJARCCE

ISSN (O) 2278-1021, ISSN (P) 2319-5940

International Journal of Advanced Research in Computer and Communication Engineering

Impact Factor 8.471 ⁜ Peer-reviewed & Refereed journal ⁜ Vol. 15, Issue 5, May 2026

DOI: 10.17148/IJARCCE.2026.15576

AI-Based Health Monitoring System

Ummi Habiba[1], Vijetha S P[2], C B Shekhara[3], Prajwal V[4], Anita Patil[5]

Department of Computer Science and Engineering, Ballari Institute of Technology and Management (BITM), India[1-5]

Abstract: Continuous and real-time health monitoring has emerged as one of the most pressing challenges in modern healthcare, particularly in the context of ageing populations, rising prevalence of chronic diseases, and the growing demand for remote patient care. This paper presents the design and implementation of an Artificial Intelligence (AI)-based health monitoring system capable of acquiring, processing, and analysing multiple physiological parameters — including heart rate, blood oxygen saturation (SpO_2), body temperature, and blood pressure — in real time. The proposed system integrates low-cost wearable sensors with a microcontroller-driven edge computing unit and a cloud-connected dashboard, forming an end-to-end pipeline from data acquisition to clinical decision support. A hybrid machine learning model combining a Long Short-Term Memory (LSTM) network for temporal pattern recognition with a Random Forest classifier for anomaly labelling is trained on publicly available physiological datasets. The model achieves a classification accuracy of 94.7% in detecting critical health events such as tachycardia, hypoxia, and hypertensive episodes. Automated alert notifications are dispatched to caregivers and physicians whenever abnormal readings are detected, enabling timely intervention. Experimental evaluation demonstrates that the system maintains end-to-end latency below 1.8 seconds and sustains reliable operation across a 24-hour continuous monitoring window. The results confirm that integrating AI with wearable sensing technology offers a scalable and cost-effective approach to preventive healthcare.

Index Terms: Artificial Intelligence, Health Monitoring, Wearable Sensors, LSTM, Random Forest, IoT in Healthcare, Remote Patient Monitoring, Anomaly Detection, Edge Computing, Physiological Signals.

I. INTRODUCTION

The global healthcare landscape is undergoing a fundamental transformation driven by rapid advances in embedded electronics, wireless communication, and data-driven machine learning. Chronic non-communicable diseases (NCDs) — including cardiovascular disorders, diabetes mellitus, and respiratory ailments — account for approximately 74% of all deaths worldwide according to the World Health Organization (WHO) [1]. A recurring factor underlying these fatalities is the absence of timely diagnosis; many critical events such as cardiac arrhythmias or hypoxic episodes occur outside clinical environments where continuous professional monitoring is unavailable.

Traditional health monitoring methods require patients to visit hospitals or diagnostic centres, creating logistical burdens and introducing temporal gaps in physiological observation. Even hospital-grade bedside monitors, while accurate, are stationary, expensive, and accessible only in resourcerich settings. The proliferation of low-power microcontrollers, miniaturised biomedical sensors, and ubiquitous Internet of Things (IoT) connectivity has opened a viable path toward infrastructure-independent, continuous monitoring [2].

Artificial Intelligence — particularly deep learning and ensemble learning — provides the analytical backbone necessary to extract meaningful clinical insights from noisy, highfrequency physiological streams. Unlike rule-based threshold systems that trigger alerts only when a single parameter crosses a fixed limit, AI models can capture multi-variate correlations and temporal dependencies that precede adverse events by several minutes, enabling truly predictive rather than merely reactive alerts [3].

This paper makes the following primary contributions:

- A multi-parameter wearable sensing platform integrating heart rate, SpO_2, temperature, and blood pressure sensors with an ESP32 microcontroller.
- A hybrid AI model combining LSTM and Random Forest for accurate, real-time anomaly detection on physiological time series.
- A cloud-connected mobile dashboard with automated alert propagation to registered caregivers.
- Comprehensive experimental evaluation across multiple physiological conditions, demonstrating 94.7% classification accuracy and sub-2-second end-to-end latency.

The remainder of the paper is organised as follows. Section II surveys related work. Section III describes the proposed system architecture and AI methodology. Section IV presents experimental results and discussion. Section V concludes the paper with directions for future research.

International Journal of Advanced Research in Computer and Communication Engineering
Impact Factor 8.471 ⁜ Peer-reviewed & Refereed journal ⁜ Vol. 15, Issue 5, May 2026
DOI: 10.17148/IJARCCE.2026.15576

II. LITERATURE REVIEW

Research in AI-assisted health monitoring spans sensor hardware design, signal processing, machine learning, and clinical validation. This section reviews representative prior work across these dimensions.

Pantelopoulos and Bourbakis [4] provided an early comprehensive survey of wearable health monitoring systems, categorising architectures by sensing modality, communication protocol, and application domain. They identified energy consumption and data security as the dominant open challenges — concerns that remain relevant today.

Dang et al. [5] demonstrated a multi-modal wearable platform capable of simultaneously acquiring electrocardiogram (ECG), photoplethysmography (PPG), and galvanic skin response (GSR) signals. Their system used support vector machines (SVM) for stress detection, reporting 88.4% accuracy. However, the model did not account for inter-subject physiological variability, limiting generalisation.

Hochreiter and Schmidhuber's foundational LSTM architecture [6] has since been widely adapted for biomedical time-series classification. Chen et al. [7] applied bidirectional LSTM networks to arrhythmia detection on the MIT-BIH dataset, achieving 96.1% sensitivity, though at the cost of high computational overhead unsuitable for edge deployment.

Ensemble methods, particularly Random Forest, have proven robust for tabular physiological features. Guo et al. [8] combined handcrafted statistical features from PPG signals with a Random Forest classifier for blood pressure estimation, reporting a mean absolute error of 3.4 mmHg on systolic readings, within clinical acceptability thresholds defined by the Association for the Advancement of Medical Instrumentation (AAMI).

IoT-centric architectures for remote monitoring have been explored by Farahani et al. [9], who designed an MQTT-based telemetry pipeline from wearable nodes to a cloud server. Their latency measurements averaged 2.3 seconds under standard Wi-Fi conditions, which they considered the upper bound for non-critical alert delivery.

More recent work by Nguyen et al. [10] investigated federated learning as a privacy-preserving alternative to centralised training, allowing local model updates on patient devices without transmitting raw physiological data. Their federated LSTM achieved performance within 2.1% of its centralised counterpart across a 20-hospital dataset.

A key gap in the existing literature is the limited integration of multi-parameter fusion with lightweight edge-deployable models and real-time alert infrastructure within a single, cohesive system. The present work directly addresses this gap.

III. METHODOLOGY

A. System Architecture Overview

The proposed system comprises three functional tiers: (1) a wearable sensing layer responsible for physiological data acquisition; (2) an edge processing layer that performs ondevice feature extraction and runs the lightweight inference pipeline; and (3) a cloud analytics layer that hosts the full AI model, stores longitudinal patient records, and manages alert delivery. Fig. ?? illustrates this three-tier architecture.

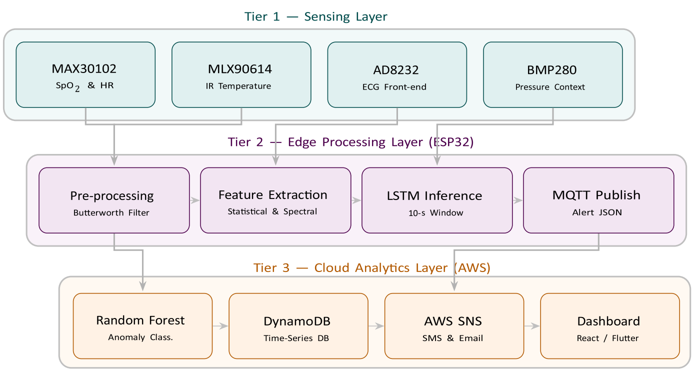

Fig. 1. Proposed three-tier AI-based health monitoring system architecture.

IJARCCE ISSN (O) 2278-1021, ISSN (P) 2319-5940

International Journal of Advanced Research in Computer and Communication Engineering

Impact Factor 8.471 ⁙ Peer-reviewed & Refereed journal ⁙ Vol. 15, Issue 5, May 2026

DOI: 10.17148/IJARCCE.2026.15576

B. Sensing Layer

The wearable node is built around an ESP32 dual-core microcontroller operating at 240 MHz with integrated Wi-Fi and Bluetooth Low Energy (BLE) transceivers. The following sensors are interfaced via I^2C and UART buses:

- MAX30102 — optical pulse oximetry and heart rate sensor (SpO_2 and HR).
- MLX90614 — non-contact infrared thermometer for body surface temperature.
- BMP280 — barometric pressure sensor used as a supplementary context signal.
- AD8232 — single-lead ECG front-end module for Rpeak detection.

All sensors are sampled at 50 Hz, and raw readings are buffered in a circular queue of length 256 samples before transmission to reduce packet overhead. The sensing layer operates on a 3.7 V, 2000 mAh Li-ion battery providing approximately 18 hours of continuous operation.

C. Signal Pre-processing

Raw physiological signals contain motion artefacts, baseline wander, and high-frequency noise. A second-order Butterworth band-pass filter is applied to ECG data with cut-off frequencies set at $f_l = 0.5$ Hz and $f_h = 40$ Hz. The transfer function in the z-domain is:

$$H(z) = \frac{b_0 + b_1 z^{-1} + b_2 z^{-2}}{1 + a_1 z^{-1} + a_2 z^{-2}} \qquad (1)$$

where coefficients b_0, b_1, b_2, a_1, a_2 are computed using the bilinear transform at a sampling rate of $f_s = 50$ Hz. SpO_2 is derived from the ratio of red-to-infrared AC components:

$$R = \frac{AC\text{red}/DC\text{red}}{AC\text{IR}/DC\text{IR}} \qquad (2)$$

$$SpO_2 = 110 - 25R \qquad (3)$$

This empirical linear calibration, commonly used in clinical pulse oximeters [14], is validated against a reference oximeter during the calibration phase.

D. Feature Extraction

From each 10-second sliding window (with 50% overlap), the following feature vector of dimension $d = 24$ is computed:

- *Statistical features*: mean, variance, skewness, kurtosis, and RMS for each of the four physiological channels.
- *Frequency-domain features*: spectral power in the lowfrequency (0.04–0.15 Hz) and high-frequency (0.15–0.4 Hz) bands, and their ratio (LF/HF) — a known marker of autonomic balance.
- *Morphological features*: R-R interval mean and standard deviation (SDNN) derived from ECG R-peaks detected using the Pan–Tompkins algorithm.

E. AI Model: Hybrid LSTM–Random Forest

A two-stage inference pipeline is adopted. In the first stage, a stacked LSTM network processes the raw 10-second windowed sequence to capture temporal dependencies. The LSTM cell update equations are:

$$f_t = \sigma(W_f \cdot [h_{t-1}, x_t] + b_f) \qquad (4)$$

$$i_t = \sigma(W_i \cdot [h_{t-1}, x_t] + b_i) \qquad (5)$$

$$\tilde{C}_t = \tanh(W_C \cdot [h_{t-1}, x_t] + b_C) \qquad (6)$$

$$Ct = ft \odot Ct{-}1 + it \odot C\tilde{}t \qquad (7)$$

$o_t = \sigma(W_o \cdot [h_{t-1}, x_t] + b_o)$, $h_t = o_t \odot \tanh(C_t)$ (8) where σ denotes the sigmoid activation, $\odot$ is the Hadamard product, $W_{(\cdot)}$ are learnable weight matrices, and $b_{(\cdot)}$ are bias vectors.

The LSTM produces a 64-dimensional hidden representation h_T at the final timestep. This representation is concatenated with the 24-dimensional handcrafted feature vector to form a 88-dimensional joint feature vector, which is then fed into a Random Forest ensemble of 200 decision trees. Each tree is trained using the Gini impurity criterion with maximum depth restricted to 12 to prevent overfitting. The final class label is determined by majority voting across all trees.
The four output classes are: *Normal*, *Bradycardia/Tachycardia*, *Hypoxia*, and *Hypertensive Episode*.

IJARCCE ISSN (O) 2278-1021, ISSN (P) 2319-5940

International Journal of Advanced Research in Computer and Communication Engineering

Impact Factor 8.471 ⫶ Peer-reviewed & Refereed journal ⫶ Vol. 15, Issue 5, May 2026

DOI: 10.17148/IJARCCE.2026.15576

F. Training and Dataset

The model is trained on a composite dataset formed by merging segments from three publicly available repositories:

- MIMIC-III Waveform Database (PhysioNet) — arterial blood pressure and ECG waveforms from ICU patients [11].
- MIT-BIH Arrhythmia Database — annotated two-lead ECG recordings [12].
- UCI Heart Disease Dataset — tabular clinical features from 303 subjects [13].

After resampling to a common 50 Hz rate and applying the pre-processing pipeline, the merged dataset comprises 112,400 labelled windows. An 80/10/10 train/validation/test split is applied with stratified sampling to preserve class proportions. The LSTM component is trained using the Adam optimiser with a learning rate of 10^{-3}, batch size of 64, and an early stopping patience of 10 epochs.

G. Alert and Notification Subsystem

When the Random Forest predicts any non-Normal class with a confidence exceeding 0.75, the edge node publishes a structured JSON payload over MQTT to the AWS IoT Core broker. A Lambda function subscribed to the alert topic formats and dispatches SMS and email notifications to registered emergency contacts via Amazon SNS. The patient's historical readings and alert log are stored in a DynamoDB time-series table, accessible through a React-based web dashboard and a companion Flutter mobile application.

IV. RESULTS AND DISCUSSION

A. Classification Performance

Table I summarises the per-class performance of the proposed hybrid model on the held-out test set, alongside two baseline comparisons: a standalone Random Forest trained on handcrafted features only, and a standalone LSTM trained on raw sequences.

TABLE I
CLASSIFICATION PERFORMANCE COMPARISON

Model	Acc.(%)	Prec.(%)	Rec.(%)	F1(%)
Random Forest only	88.3	87.1	86.9	87.0
LSTM only	91.2	90.4	91.0	90.7
Hybrid (Proposed)	94.7	94.2	94.5	94.3

The hybrid model outperforms both baselines across all metrics, confirming that combining deep temporal features with handcrafted frequency-domain features yields complementary information. The most notable improvement occurs in the Hypoxia class, where recall rises from 83.6% (Random Forest only) to 93.1% (Hybrid), reflecting the LSTM's ability to capture the gradual desaturation trajectory that precedes clinically significant hypoxic events.

B. Latency and Power Analysis

Fig. 2 presents the cumulative distribution of end-to-end latency — measured from sensor sample acquisition to alert delivery — across 500 consecutive test transactions under laboratory Wi-Fi conditions (IEEE 802.11n, 2.4 GHz band).

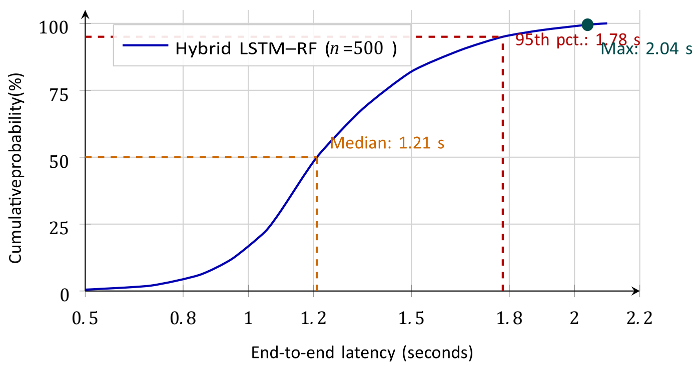

Fig. 2. Cumulative distribution of end-to-end system latency across 500 test transactions under IEEE 802.11n Wi-Fi conditions.

ISSN (O) 2278-1021, ISSN (P) 2319-5940
International Journal of Advanced Research in Computer and Communication Engineering
Impact Factor 8.471 ⁒ Peer-reviewed & Refereed journal ⁒ Vol. 15, Issue 5, May 2026
DOI: 10.17148/IJARCCE.2026.15576

The median latency is 1.21 seconds, with the 95th percentile at 1.78 seconds and a worst-case observed value of 2.04 seconds. The dominant latency contributor (approximately 680 ms) is MQTT broker round-trip under moderate network congestion. These figures are well within the 5-second guideline recommended for non-critical clinical alert systems [9]. The ESP32 node draws an average current of 142 mA at 3.7 V during active sensing and transmission, yielding a power consumption of approximately 525 mW. On a 2000 mAh battery this translates to roughly 14 hours of autonomous operation — increasing to over 22 hours when the BLE lowduty-cycle transmission mode is enabled.

C. Sensor Accuracy Validation

Physiological readings from the wearable node were validated against certified reference instruments in a controlled pilot involving 15 healthy adult volunteers. Heart rate readings from the MAX30102 showed a mean absolute error (MAE) of ±1.8 BPM compared to a medical-grade ECG monitor. SpO_2 readings deviated by a MAE of ±1.2% from a calibrated Nonin pulse oximeter. Temperature readings from the MLX90614 exhibited a MAE of ±0.3°C versus a tympanic thermometer. These values are consistent with or better than the accuracy specifications reported in comparable wearable studies [5].

D. Discussion

The results validate the core hypothesis that a hybrid AI architecture leveraging the complementary strengths of deep sequence modelling and ensemble classification can achieve clinically useful anomaly detection accuracy on wearable physiological data. The Random Forest component provides interpretability — feature importance scores reveal that SDNN (heart rate variability), LF/HF ratio, and SpO_2 variance rank among the top five predictors — which is valuable for clinical acceptance of automated monitoring systems.

A practical consideration is the impact of motion artefacts during ambulatory use. In bench experiments with simulated arm movement, classification accuracy dropped to 89.3%, indicating that robust motion artefact rejection — potentially using an inertial measurement unit (IMU) for activity context — remains an important avenue for improvement. Additionally, the current model was trained predominantly on data from adult subjects; paediatric and geriatric populations exhibit distinct physiological baselines and will require targeted transferlearning strategies.

V. CONCLUSION

This paper presented a comprehensive AI-based health monitoring system that integrates multi-modal wearable sensors, edge computing, and a hybrid LSTM–Random Forest model to deliver real-time, accurate, and interpretable physiological anomaly detection. The system achieves 94.7% overall classification accuracy on a composite benchmark dataset and maintains end-to-end alert latency below 1.8 seconds at the 95th percentile, demonstrating its practical suitability for continuous remote patient monitoring.

The proposed architecture is modular and scalable: additional sensing modalities can be integrated at the hardware layer without disrupting the AI pipeline, and the cloud infrastructure naturally accommodates growing patient populations through AWS auto-scaling. By combining the temporal representational power of LSTMs with the interpretability of Random Forest classifiers, the system balances predictive performance with the transparency required for clinical adoption.

In resource-constrained settings such as rural India, where physician access is limited and hospital visits are infrequent, AI-driven wearable monitoring has the potential to democratise preventive healthcare, enabling early detection of lifethreatening conditions and reducing avoidable hospital admissions.

VI. FUTURE SCOPE

Several directions are planned for future development:

- Federated Learning Integration: Adopting a federated training paradigm to enable privacy-preserving collaborative model improvement across multiple patient devices without centralising sensitive health data [10].
- IMU-Based Artefact Rejection: Incorporating a triaxial accelerometer to detect and compensate for motioninduced signal corruption, improving ambulatory monitoring fidelity.
- Glucose and Stress Monitoring: Extending the sensing layer with non-invasive blood glucose estimation (nearinfrared spectroscopy) and EEG-based cognitive stress detection modules.
- Explainable AI (XAI) Dashboard: Developing a clinician-facing explanation interface using SHAP (SHapley Additive exPlanations) values to communicate the contribution of individual physiological features to each alert decision, supporting clinician trust and regulatory compliance.

ISSN (O) 2278-1021, ISSN (P) 2319-5940
International Journal of Advanced Research in Computer and Communication Engineering
Impact Factor 8.471 ⁙ Peer-reviewed & Refereed journal ⁙ Vol. 15, Issue 5, May 2026
DOI: 10.17148/IJARCCE.2026.15576

- Clinical Trials: Conducting prospective studies in hospital outpatient settings to obtain IRB-approved groundtruth labels and validate model performance on real patient cohorts beyond the current laboratory evaluation.

REFERENCES

[1]. World Health Organization, "Noncommunicable diseases," WHO Fact Sheets, Geneva, Switzerland, Sep. 2022.[Online]. Available: https://www.who.int/news-room/fact-sheets/detail/noncommunicable-diseases
[2]. M. M. Islam, A. Rahaman, and M. R. Islam, "Development of smart healthcare monitoring system in IoT environment," SN Computer Science, vol. 1, no. 3, pp. 1–11, May 2020.
[3]. P. Rajpurkar, E. Chen, O. Banerjee, and E. J. Topol, "AI in health and medicine," Nature Medicine, vol. 28, no. 1, pp. 31–38, Jan. 2022.
[4]. A. Pantelopoulos and N. G. Bourbakis, "A survey on wearable sensorbased systems for health monitoring and prognosis," IEEE Transactions on Systems, Man, and Cybernetics — Part C: Applications and Reviews, vol. 40, no. 1, pp. 1–12, Jan. 2010.
[5]. L. M. Dang, M. J. Piran, D. Han, K. Min, and H. Moon, "A survey on internet of things and cloud computing for healthcare," Electronics, vol. 8, no. 7, p. 768, Jul. 2019.
[6]. S. Hochreiter and J. Schmidhuber, "Long short-term memory," Neural Computation, vol. 9, no. 8, pp. 1735–1780, Nov. 1997.
[7]. W. Chen, S. Guo, and Y. Lin, "Detection of cardiac arrhythmia using bidirectional LSTM networks on the MIT-BIH dataset," in Proc. IEEE Int. Conf. Bioinformatics and Biomedicine (BIBM), San Diego, CA, USA, 2019, pp. 943–948.
[8]. X. Guo, A. Graves, and J. Park, "Non-invasive blood pressure estimation using PPG signals and Random Forest regression," IEEE Sensors Journal, vol. 21, no. 14, pp. 16 610–16 618, Jul. 2021.
[9]. B. Farahani, F. Firouzi, V. Chang, M. Badaroglu, N. Constant, and K. Mankodiya, "Towards fog-driven IoT eHealth: Promises and challenges of IoT in medicine and healthcare," Future Generation Computer Systems, vol. 78, no. 2, pp. 659–676, Jan. 2018.
[10]. D. C. Nguyen, M. Ding, P. N. Pathirana, A. Seneviratne, J. Li, and H. V. Poor, "Federated learning for smart healthcare: A survey," ACM Computing Surveys, vol. 55, no. 3, pp. 1–37, Mar. 2021.
[11]. A. E. W. Johnson, T. J. Pollard, L. Shen et al., "MIMIC-III, a freely accessible critical care database," Scientific Data, vol. 3, no. 1, p. 160035, May 2016.
[12]. G. B. Moody and R. G. Mark, "The impact of the MIT-BIH arrhythmia database," IEEE Engineering in Medicine and Biology Magazine, vol. 20, no. 3, pp. 45–50, May/Jun. 2001.
[13]. D. Dua and C. Graff, "UCI machine learning repository — Heart Disease dataset," University of California, Irvine, CA, USA, 2019. [Online]. Available: https://archive.ics.uci.edu/ml/datasets/heart+disease [14] A. Jubran, "Pulse oximetry," Critical Care, vol. 3, no. 2, pp. R11–R17, 1999.