

Wearable Health Monitoring System: A Multi-Sensor IoT Architecture with Edge

Intelligence and Cloud-Based Anomaly Detection

Neal Bheda, Dr.Asha Durafe, Neel Mahadik, Heet Thakkar, Jeel Mange, Dhruv Trivedi
Aayushi Bhanushali

Department of Electronics and Computer Science Shah & Anchor Kutchhi Engineering College, India.

Abstract- Wearable health monitoring systems (WHMS) represent a transformative paradigm for continuous, non-invasive physiological surveillance, enabling early detection of life-threatening conditions outside traditional clinical settings. This paper presents the design, implementation, and clinical evaluation of a compact five-modality WHMS that concurrently acquires electrocardiographic (ECG) waveforms, photoplethysmographic (PPG) oxygen saturation (SpO₂), infrared skin-surface temperature, barometric-pressure-derived blood pressure estimates, and six-axis inertial motion data. An on-board Raspberry Pi Zero 2W edge processor executes real-time QRS detection, four-class LSTM arrhythmia classification, adaptive-filter SpO₂ computation, and random-forest fall detection before forwarding compressed feature vectors via Bluetooth Low Energy 5.0 to a paired smartphone gateway and thence to a HIPAA-compliant cloud backend. Clinical validation across 22 volunteers against gold-standard reference instruments yielded ECG heart-rate MAE of 0.7 bpm, SpO₂ MAE of 1.1%, systolic blood-pressure MAE of 4.8 mmHg (meeting AAMI SP-10), arrhythmia classification macro-F1 of 95.4%, fall-detection sensitivity of 97.3%, and an operational runtime of 11.2 hours under a realistic mixed-activity duty-cycle profile. These results demonstrate that deeply integrated edge-cloud WHMS architectures are capable of meeting clinical accuracy standards within a consumer-wearable form factor.

Keywords: Wearable Sensors, Health Monitoring, Electrocardiography, SpO₂, Edge Computing, LSTM, BLE 5.0, Fall Detection, Internet of Things, Anomaly Detection, Remote Patient Monitoring, Federated Learning.

I. INTRODUCTION

Chronic non-communicable diseases — cardiovascular disorders, type-2 diabetes, chronic respiratory ailments, and neurological conditions — account for approximately 74% of annual global mortality, with the World Health Organisation projecting that the burden will intensify most sharply in low-to-middle-income nations over the coming decades [1]. The conventional paradigm of episodic clinical encounters is structurally ill-suited to managing conditions whose trajectories are continuous and whose critical events — arrhythmias, hypoglycaemic crises, apnoeic spells — arise unpredictably and transiently between appointments.

Wearable health monitoring systems (WHMS) offer a clinically compelling alternative: miniaturised, body-worn sensor arrays that acquire, condition, and transmit physiological signals in real time. Early

devices such as the Holter monitor (introduced in 1961) established the diagnostic value of extended ambulatory cardiac recording, but were too cumbersome for true continuous daily wear. The convergence of microelectromechanical systems (MEMS) fabrication, ultra-low-power mixed-signal integrated circuits, energy-efficient short-range wireless protocols, and cloud-hosted machine learning has now enabled wristband-, chest-patch-, and smart-textile-form-factor devices capable of multi-parameter physiological surveillance across full waking days on a single charge [2].

This paper presents the design, implementation, and experimental validation of a compact five-modality WHMS prototype that simultaneously acquires electrocardiographic (ECG) waveforms, photoplethysmographic (PPG) oxygen saturation (SpO₂), infrared skin-surface temperature, barometric-pressure-derived blood pressure estimates, and six-axis inertial motion data. An on-

board Raspberry Pi Zero 2W edge node performs real-time QRS detection, four-class LSTM arrhythmia classification, adaptive-filter SpO₂ computation, and random-forest fall detection before transmitting compressed feature vectors via Bluetooth Low Energy 5.0 to a paired smartphone gateway. Annotated records are forwarded to a HIPAA-compliant cloud backend where a personalised LSTM anomaly detector flags clinically significant deviations for asynchronous clinician review [3].

The paper is organised as follows. Section II surveys related literature across sensor modalities, edge processing, wireless transport, and cloud analytics. Section III details the system architecture and hardware design. Section IV describes the signal-processing and machine-learning pipeline. Section V reports experimental results. Section VI discusses findings comparatively. Section VII concludes and outlines future directions.

II. LITERATURE REVIEW

A. Sensor Modalities and Signal Quality

Photoplethysmography has emerged as the dominant non-invasive modality for continuous SpO₂ and heart-rate estimation owing to the simplicity of its two-photodiode optical front end. Tamura et al. [4] surveyed wearable PPG sensors and identified motion artefact as the principal source of ambulatory measurement error, reporting that least-mean-squares adaptive filters driven by co-located accelerometer signals reduce SpO₂ bias from 4.2% to 0.9% under controlled treadmill exercise. ECG acquisition in wearable form factors introduces complementary challenges: textile dry electrodes exhibit contact impedances one to two orders of magnitude higher than gel Ag/AgCl electrodes, introducing baseline wander that standard high-pass filters suppress only partially. Gargiulo et al.

[5] demonstrated that impedance-adaptive notch filters attain SNRs within 1.8 dB of clinical electrodes without the skin preparation requirements that preclude self-application. Infrared thermopile temperature sensing gained renewed prominence during the COVID-19 pandemic. Haxha et al.

[6] showed that wrist-mounted pyrometric sensors correlate with tympanic core-body temperature at $r = 0.97$ after a per-user calibration offset is applied, providing a non-contact alternative to oral or axillary probes. Blood pressure estimation from photoplethysmographic pulse-transit time remains an active area; current consumer implementations achieve mean absolute errors of 4–8 mmHg for systolic pressure — within the ± 5 mmHg threshold of the AAMI SP-10 standard under controlled static conditions [14].

B. Edge Inference on Embedded Platforms

Transmitting raw physiological waveforms to the cloud incurs latency and energy overhead incompatible with real-time alerting. Lane et al. [7] profiled fifteen machine-learning workloads on ARM Cortex-M microcontrollers and found that INT8-quantised CNNs operating on spectral features consume 42% less energy than equivalent float32 models while sacrificing fewer than two percentage points of accuracy. Zhang et al. [8] implemented a binary neural network for R-peak detection that achieved 98.6% accuracy on the MIT-BIH database consuming only 0.8 mW on a 90 nm CMOS test chip — demonstrating milliwatt-class arrhythmia screening. More recently, TinyML frameworks (TensorFlow Lite Micro, Edge Impulse) have lowered the barrier to deploying recurrent architectures on embedded Cortex-A platforms without custom silicon [7].

C. Wireless Body-Area Networking

Bluetooth Low Energy has become the de facto BAN data-offload standard, owing to sub-milliwatt advertising power, ubiquitous smartphone support, and a well-defined Generic Attribute Profile for health device interoperability. Siekkinen et al. [9] compared BLE, ZigBee, and ANT+ for biometric streaming and concluded that BLE 5.0's 2 Mbps PHY and extended advertising PDU sustain twelve-hour uninterrupted ECG streaming from a 200 mAh cell. For infrastructure-less direct-to-cloud connectivity, NB-IoT offers wide-area reach but at a 23 dBm transmit power that imposes a substantially higher per-bit energy cost than BLE [10], making it unsuitable for wrist-worn form factors without auxiliary energy harvesting.

D. Cloud Analytics and Anomaly Detection

Cloud-hosted anomaly detectors close the diagnostic loop by correlating instantaneous readings against personalised physiological baselines and population reference norms. Clif-ford et al. [11] benchmarked seven classifiers on the PhysioNet/Computing in Cardiology 2017 atrial-fibrillation dataset and found that single-layer LSTMs outperformed hand-crafted feature classifiers by 7.3 percentage points in F1-score. Sub-sequent work has investigated privacy-preserving federated learning architectures that train personalised anomaly models across distributed patient cohorts without centralising raw physiological records [12], addressing a key barrier to deployment in jurisdictions with stringent data-sovereignty regulations.

E. Identified Gaps in Prior Work

Three systematic gaps emerge from the reviewed corpus. First, most prototypes evaluate sensors in isolation; end-to-end studies spanning sensor fusion, on-device inference, wireless transport, and cloud analytics within a single cohort are uncommon. Second, battery lifetime is frequently reported under continuous-stream conditions rather than realistic duty-cycled profiles, inflating claimed runtimes. Third, clinical validation against gold-standard reference instruments — a prerequisite for translational relevance — is omitted from the majority of engineering-focused publications. The present work directly addresses all three gaps.

III. SYSTEM ARCHITECTURE

The proposed WHMS is organised into three tiers: (i) a wearable sensing-and-edge-computing node, (ii) a smartphone gateway, and (iii) a cloud analytics backend.

A. Wearable Sensing Node

The sensing node is realised on a four-layer flexible PCB measuring 55 mm × 38 mm × 4 mm and weighing 18 g inclusive of a 350 mAh lithium-polymer cell. Processing is performed by a Raspberry Pi Zero 2W (Broadcom BCM2710A1, quad-core Cortex-A53 @ 600 MHz, 512 MB LPDDR2) operating at a reduced 600 MHz clock. Five sensor

modules are integrated; their specifications are detailed in Table I.

TABLE I
SENSOR MODULE SPECIFICATIONS AND
INTERFACE SUMMARY

SENSOR MODULE	PARAMETER	PROTOCOL	ACCURACY	POWER (mA)
AD8232 (ECG)	HEART RATE	SPI	±1.5%	0.17
MAX30102 (PPG)	SPO	I ² C	±2%	1.20
MLX90614 (IR)	TEMPERATURE	I ² C	±0.5°C	2.00
BMP390 (BARO)	BP ESTIMATION	SPI	±1 MMHG	0.17
MPU-6050 (IMU)	ACTIVITY/FALL	I ² C	±0.1 G	3.90

The AD8232 acquires cardiac signals at 500 SPS with 12-bit resolution. The MAX30102 samples at 100 Hz with an on-chip 18-bit ADC. The MLX90614 provides readings every two seconds. The BMP390 contributes data for pulse-transit-time blood-pressure estimation, and the MPU-6050 records motion data at 50 Hz.

B. Power Management Subsystem

A Texas Instruments BQ25895 USB-C power-path IC manages single-cell charging. Four operating modes are defined: Full-Active, Selective-Active, Standby, and Deep-Sleep. Transitions are governed by activity: high-activity triggers Full-Active, while sedentary periods trigger Selective-Active or Standby. Targeting a weighted-average current of 28 mA yields a projected 12-hour runtime.

C. Smartphone Gateway Application

A Flutter application maintains a persistent BLE GATT connection. On receipt of encrypted 20-byte feature vectors (AES-128-GCM), the gateway renders real-time charts and evaluates a rule engine (e.g., SpO₂ < 90% triggers an alert). Feature vectors are batched and forwarded over TLS 1.3 to the cloud every 60 s.

D. Cloud Analytics Backend

The cloud backend is deployed on a containerised microservices stack (Docker/Kubernetes) in an AWS HIPAA-eligible environment. Apache Kafka routes records to (a) Apache Flink for LSTM anomaly detection and (b) an S3-backed time-series store in Parquet format. Clinician dashboards are built in Grafana.

IV. SIGNAL PROCESSING AND ML PIPELINE

A. ECG Conditioning and QRS Detection

Incoming 500 Hz samples pass through a second-order Butterworth band-pass filter (0.5–40 Hz). QRS complexes are located using a Modified Pan-Tompkins algorithm. Heart-rate variability indices (SDNN, RMSSD, pNN50) are computed over five-minute windows. A 64-point FFT extracts eight frequency-domain features, forming a 13-element input vector to a four-class LSTM (Normal, AFib, PVC, BBB) trained on 109,000 beats from the MIT-BIH database [13].

B. SpO Computation with Artefact Rejection

SpO is estimated via the ratio-of-ratios method. A normalised least-mean-squares adaptive filter uses IMU accelerometer as a reference to cancel motion-induced fluctuations. This reduced MAE from 3.1% to 0.9% under treadmill walking.

C. Fall Detection

A 1 s sliding window yields a 36-element feature vector. A random forest of 40 decision trees, trained on the SisFall dataset, achieves 97.3% sensitivity. Edge inference completes in 11 ms.

D. Cloud LSTM Anomaly Detector

A two-layer LSTM (128 → 64 units) ingests 72-hour sequences. The model outputs an anomaly score; deviations exceeding two standard deviations trigger clinician notification. Fine-tuning uses federated gradient aggregation [12].

V. EXPERIMENTAL RESULTS

A. Sensor Accuracy Against Clinical References
The prototype was evaluated against 22 volunteers. ECG heart rate yielded MAE 0.7 bpm ($r = 0.998$). SpO yielded MAE 1.1% against a Nellcor N-595. Temperature correlated at $r = 0.96$. Systolic blood pressure achieved MAE 4.8 mmHg, satisfying AAMI SP-10 [14].

B. Arrhythmia Classification Performance

The four-class LSTM achieved a macro-F1 score of 95.4% on the MIT-BIH test split, exceeding the 92.1% reported by Clifford et al. [11]. INT8 quantisation reduced model size to 0.54 MB.

C. Fall Detection Validation

In a separate 10-volunteer session, sensitivity was 96.8% and specificity 95.2%. Mean latency from fall to alert was 186 ms.

E. Power and Communication Performance

Measured average current was 31.2 mA, yielding a runtime of 11.2 h. BLE 5.0 packet loss was 0.4% at 15 m. Cloud refresh latency averaged 2.1 s.

F. Comparative Analysis

Table II compares this work against prior studies. This prototype is unique in integrating five modalities with edge and cloud intelligence in a sub-20 g form factor.

TABLE II
PERFORMANCE COMPARISON WITH
REPRESENTATIVE PRIOR WHMS STUDIES

Study	Sensors	Edge ML	Cloud AI	Battery	MAE(HR)
Jamura [4]	PPG	No	No	72 h	1.8 bpm
Zhang [8]	ECG	Yes	No	18 h	0.9 bpm
Siekkinen [9]	ECG+SpO	No	No	24 h	1.5 bpm
Roh [3]	Multi	No	Yes	16 h	1.2 bpm
Proposed	5-modal	Yes	Yes	11.2 h	0.7 bpm

VI. DISCUSSION

The results demonstrate clinical accuracy across five modalities. Limitations include a modest 22-volunteer cohort and an 11.2-hour battery life, which falls short of a full waking day for high-activity users. Federated personalisation [12] improved anomaly precision by 11.4%.

VII. CONCLUSION AND FUTURE WORK

This paper presented a five-modality WHMS with edge-cloud intelligence. Validation confirmed accuracy standards are met. Future work includes developing a custom low-power ASIC to extend battery life to 24+ hours, integrating glucose monitoring, and conducting a randomized trial in a cardiology department.

Acknowledgment

The authors thank the Department of Electronics and Computer Science, Shah & Anchor Kutchhi Engineering College, for support and the volunteer participants.

REFERENCES

1. World Health Organisation, "Non-communicable diseases," WHO Fact Sheets, Sep. 2023.
2. S. Patel et al., "A review of wearable sensors and systems," *J. NeuroEng. Rehabil.*, vol. 9, p. 21, 2012.
3. C. Roh et al., "Wearable sensor-based remote health monitoring system," *IEEE Access*, vol. 7, 2019.
4. Y. Tamura et al., "Wearable photoplethysmographic sensors," *Electronics*, vol. 3, 2014.
5. G. Gargiulo et al., "A mobile EEG system with dry electrodes," *Proc. IEEE BioCAS*, 2008.
6. S. Haxha et al., "Wearable non-contact PPG and temperature sensor," *IEEE Sens. J.*, vol. 20, 2020.
7. N. D. Lane et al., "Deep learning on wearables," *Proc. Int. Workshop IoT Design*, 2016.
8. Q. Zhang et al., "HeartID: ECG-based biometric identification," *IEEE Access*, vol. 5, 2017.
9. M. Siekkinen et al., "How low energy is Bluetooth low energy?" *Proc. IEEE WCNC*, 2012.
10. M. Lauridsen et al., "Coverage comparison of GPRS, NB-IoT," *Proc. IEEE VTC Spring*, 2017.
11. G. D. Clifford et al., "AF classification from a single lead ECG," *Proc. Comput. Cardiol.*, 2017.
12. T. Li et al., "Federated learning," *IEEE Signal Process. Mag.*, vol. 37, 2020.
13. G. B. Moody and R. G. Mark, "The impact of the MIT-BIH arrhythmia database," *IEEE Eng. Med. Biol. Mag.*, vol. 20, 2001.
14. AAMI Standard SP-10, *Electronic or Automated Sphygmomanometers*, 2002.