AIJFR Advanced International Journal for Research (AIJFR)
E-ISSN: 3048-7641 ● Website: www.aijfr.com ● Email: editor@aijfr.com

IoT-Based Remote Patient Monitoring System for Continuous Chronic Disease Management: A Machine Learning Approach

Mahesh Krishna[1], Irimpan Arnold[2], Dr. Kannagi Anbazhagan[3]

[1,2,3] School of Computer Science and Information Technology, Jain University

Abstract

Diabetes management represents a critical challenge in modern healthcare, where continuous monitoring costs significantly exceed episodic treatment expenses. This research presents a systematic analysis of an IoT-based Remote Patient Monitoring (RPM) system for diabetic patients, integrating wearable sensors, cloud computing, and machine learning for real-time health tracking. Using an ESP32 microcontroller architecture with MAX30102 heart rate sensors, glucose monitors, and temperature sensors, a comprehensive data pipeline is implemented, including wireless transmission via MQTT protocol, cloud-based storage, and anomaly detection algorithms. The system addresses the 26.5% gap in continuous care delivery for middle-class diabetic populations, where traditional hospital-based monitoring proves economically prohibitive. Results demonstrate that real-time monitoring enables proactive intervention, reducing emergency hospitalisations by 40% while lowering long-term treatment costs by 30% compared to conventional periodic check-ups. Feature analysis reveals glucose level fluctuations, heart rate variability, and body temperature as primary health indicators requiring immediate intervention. An integrated alert framework connects predictive outputs to mobile applications and healthcare provider dashboards for immediate response. This research advances both theoretical understanding of remote healthcare dynamics and practical implementation strategies for medical institutions seeking cost-effective continuous monitoring capabilities.

Keywords: Remote Patient Monitoring, IoT Healthcare, Diabetes Management, Wearable Sensors, ESP32, Machine Learning, Cloud Computing, MQTT Protocol

1. Introduction

The healthcare industry faces unprecedented challenges in managing chronic diseases in an era of ageing populations and resource constraints. With diabetes prevalence exceeding 537 million adults globally and projected to reach 783 million by 2045, the strategic imperative has shifted from reactive treatment to proactive management [1]. Diabetes mellitus—characterised by persistent hyperglycaemia requiring continuous monitoring—directly impacts mortality rates, healthcare costs, and quality of life.

Industry research indicates that continuous monitoring costs 5–7 times less than emergency interventions, while hospital readmission rates for diabetic complications typically range from 18–35% depending on disease severity and monitoring frequency [2]. For a healthcare system serving 5 million diabetic patients experiencing a 2% monthly emergency hospitalisation rate, the annual cost impact exceeds $600 million in acute care expenses, excluding long-term complication management costs. This economic reality has catalysed investment in remote monitoring technologies capable of identifying health deterioration before emergency situations arise.

Traditional diabetes management approaches—including periodic blood glucose testing, scheduled clinic visits, and post-complication treatment—suffer from temporal lag, incomplete data capture, and limited preventive capability. These methods identify health risks only after complications have manifested, restricting intervention effectiveness. The convergence of Internet of Things (IoT) infrastructure, advanced wearable sensors, and wireless communication protocols has enabled the transformation from reactive to proactive healthcare paradigms.

This research addresses the gap between technological capability and clinical deployment by developing a comprehensive framework encompassing sensor integration, data transmission, cloud analytics, and alert system implementation. Four technology components representing distinct system layers are systematically integrated: wearable sensors (data acquisition), ESP32 microcontroller (edge processing), cloud platform (data storage and analysis), and mobile applications (user interface), with performance evaluated across reliability, latency, accuracy, and cost-effectiveness metrics within a healthcare context.

A. Research Objectives

This study pursues four primary objectives: (1) develop a robust hardware architecture integrating multiple physiological sensors with a low-power microcontroller for continuous data acquisition; (2) implement a wireless data transmission pipeline utilising MQTT protocol for real-time cloud synchronisation with minimal latency; (3) design a machine learning-based anomaly detection system identifying hypoglycaemic and hyperglycaemic events through pattern recognition; and (4) construct an actionable alert framework translating sensor data into immediate notifications for patients and healthcare providers.

B. Key Contributions

This research contributes: systematic integration of wearable sensors under a unified IoT architecture; a comprehensive wireless communication framework balancing real-time performance with power efficiency; empirical validation using continuous monitoring scenarios simulating real-world diabetic patient conditions; and a practical alert system architecture bridging sensor data and clinical intervention. Unlike prior studies focusing exclusively on algorithmic optimisation, this work addresses end-to-end deployment considerations essential for healthcare value realisation.

2. Literature Review

A. Remote Patient Monitoring Evolution

Remote patient monitoring research evolved from simple periodic telemetry in the 1990s to sophisticated continuous monitoring architectures in contemporary applications. Early studies employed basic heart rate monitors for cardiac patients, enabling healthcare providers to track vital signs remotely [3]. However, limited sensor diversity and the lack of integrated platforms constrained comprehensive health assessment, motivating investigation of multi-parameter monitoring systems.

Kumar et al. [4] provided pioneering work in IoT-based wearable health monitoring, integrating ECG and SpO_2 sensors with cloud servers. Their findings demonstrated that real-time physiological monitoring significantly improved emergency response time, highlighting the value of continuous data streams for early intervention. Despite these promising results, the system focused exclusively on cardiac parameters without addressing metabolic conditions such as diabetes, and lacked machine learning integration for predictive analytics.

B. IoT Healthcare Architectures

Al-Hameed et al. [5] developed comprehensive wearable IoT devices for continuous monitoring of heart rate and body temperature, integrating mobile applications for patient-provider communication. Their system demonstrated user-friendly interfaces and improved data accessibility. However, limitations in advanced analytics, data security protocols, and scalability hindered adoption for long-term chronic disease management requiring robust privacy protections and continuous operation.

Li et al. [6] advanced the field by applying deep learning techniques including Convolutional Neural Networks (CNN) and Long Short-Term Memory (LSTM) models to physiological data streams. Their research achieved high accuracy in detecting cardiac anomalies and stress patterns through automated pattern recognition. However, computational complexity made real-time deployment challenging, particularly for resource-constrained wearable devices requiring low power consumption and extended battery life.

C. Diabetes-Specific Monitoring Systems

Gupta et al. [7] developed IoT-based glucose monitoring systems combining continuous glucose meters with smartphone applications for data visualisation. Their study demonstrated improved patient engagement and glycaemic control through real-time feedback. However, single-parameter focus on glucose levels without integrating complementary vital signs (heart rate, temperature, activity) limited comprehensive health assessment capabilities essential for detecting diabetes-related complications.

D. Research Gap

While extensive literature addresses individual system components, limited research integrates complete end-to-end architectures spanning hardware, communication, analytics, and intervention. Existing studies treat monitoring as terminal output rather than input for immediate clinical action. This research addresses this gap by developing a comprehensive RPM framework that integrates wearable sensors, wireless protocols, cloud analytics, and alert systems, enabling proactive diabetes management through real-time health tracking and automated intervention triggering.

3. Methodology

A. Research Framework

This research adopts the Cross-Industry Standard Process for Data Mining (CRISP-DM) framework [8], adapted for IoT healthcare applications, encompassing system design, hardware integration, data pipeline development, analytics implementation, and deployment conceptualisation. An experimental research design is employed within the engineering paradigm, utilising empirical testing and performance measurement for system validation.

B. System Architecture Overview

The proposed IoT-based Remote Patient Monitoring system comprises four integrated layers: (1) Sensor Layer—wearable devices capturing physiological parameters; (2) Edge Processing Layer—ESP32 microcontroller performing data aggregation and preliminary processing; (3) Communication Layer—wireless transmission via MQTT protocol to cloud infrastructure; and (4) Application Layer—cloud-based analytics, storage, and user interfaces for patients and providers.

The system architecture follows distributed processing principles where edge devices handle time-critical preprocessing while cloud platforms perform computationally intensive analytics. This hybrid approach optimises latency, power consumption, and scalability while maintaining real-time monitoring capabilities essential for diabetes management.

C. Hardware Implementation

1. Sensor Selection and Integration

Physiological data acquisition utilises three primary sensors integrated on a wearable wristband form factor. The MAX30102 pulse oximetry sensor captures heart rate (30–220 BPM) and SpO_2 levels (70–100%) using photoplethysmography (PPG) technology with red and infrared LEDs. A DS18B20 digital temperature sensor measures skin temperature (±0.5°C accuracy) for fever detection and metabolic activity assessment. Non-invasive glucose monitoring employs optical spectroscopy techniques correlating blood glucose with tissue absorption characteristics (±15 mg/dL accuracy).

Sensor integration with the ESP32 microcontroller utilises I2C protocol for MAX30102 communication and one-wire protocol for DS18B20, minimising GPIO pin requirements while enabling multi-sensor coordination. Sampling rates are configured at 1 Hz for temperature (sufficient for slow-changing thermal dynamics), 10 Hz for glucose (capturing meal-response patterns), and 25 Hz for heart rate (enabling heart rate variability analysis).

2. Microcontroller Configuration

The ESP32-WROOM-32 microcontroller serves as the edge processing unit, selected for integrated Wi-Fi connectivity (802.11 b/g/n), dual-core Tensilica LX6 architecture (240 MHz clock), and low power consumption (80 mA active, 5 µA deep sleep). The dual-core architecture enables parallel sensor reading (Core 0) and wireless transmission (Core 1), preventing data acquisition interruption during network operations.

Power management implements duty cycling where the ESP32 enters light sleep between sensor readings, reducing average power consumption to 12 mA for 20+ hour battery life on a 1000 mAh LiPo battery. Wake-on-sensor interrupt enables immediate activation for critical glucose level changes without continuous active monitoring.

D. Data Transmission Pipeline

1. MQTT Protocol Implementation

Wireless data transmission utilises Message Queuing Telemetry Transport (MQTT) protocol, selected for lightweight overhead (2-byte minimum header), publish-subscribe architecture, and Quality of Service (QoS) levels appropriate for healthcare telemetry. The system implements QoS level 1 (at-leastonce delivery), balancing reliability with bandwidth efficiency and ensuring critical health data reaches the cloud platform without excessive retransmission overhead.

Topic hierarchy follows the structure patients/{patient_id}/{sensor_type}, enabling granular subscription management where healthcare providers monitor specific patients while analytics engines process aggregate data streams. Payload formatting uses JSON for human readability and parsing flexibility, with typical message size below 200 bytes, maintaining bandwidth efficiency on 2.4 GHz Wi-Fi networks.

2. Cloud Platform Integration

Cloud infrastructure utilises a HIPAA-compliant platform (AWS IoT Core) providing MQTT broker services, device authentication via X.509 certificates, and encrypted transmission (TLS 1.2). Device shadows maintain last-known sensor states, enabling mobile applications to retrieve current patient status without direct device communication—crucial for scenarios where wearable devices operate intermittently or in low-connectivity environments.

The data ingestion pipeline processes incoming MQTT messages through AWS Lambda functions performing validation, timestamp normalisation, and database insertion into DynamoDB time-series tables. Partitioning by patient ID and time range optimises query performance for historical trend analysis while maintaining sub-second latency for real-time dashboard updates.

E. Machine Learning Analytics

1. Anomaly Detection Algorithms

Hypoglycaemia and hyperglycaemia detection employs threshold-based rules combined with trend analysis. Static thresholds (glucose <70 mg/dL for hypoglycaemia; >180 mg/dL for hyperglycaemia) trigger immediate alerts. Dynamic trend analysis using moving average convergence (5-minute and 15-minute windows) detects rapid glucose changes exceeding 30 mg/dL per 15 minutes, indicating impending complications requiring preemptive intervention.

Heart rate anomaly detection utilises statistical process control, identifying values exceeding two standard deviations from personalised baselines established during a 14-day calibration period. This approach accounts for individual variability where normal ranges differ significantly across patients based on age, fitness level, and medications. Temperature monitoring flags readings above 37.5°C as potential infection indicators requiring clinical evaluation.

2. Predictive Pattern Recognition

Long Short-Term Memory (LSTM) neural networks trained on historical glucose patterns predict future levels 30–60 minutes ahead, enabling proactive intervention before dangerous thresholds are reached. The network architecture comprises a 64-unit LSTM layer, a 32-unit dense layer, and a single output neuron predicting continuous glucose value. Training data includes 6-hour sliding windows capturing meal patterns, medication timing, and activity levels.

Feature engineering incorporates time-of-day encoding (sine/cosine transformation for circadian patterns), day-of-week indicators (capturing routine variations), and meal proximity flags (derived from glucose spike patterns). The model achieves a mean absolute error of 12.8 mg/dL on the test set, providing sufficient accuracy for early warning generation without excessive false alarms that would compromise user trust.

F. Alert System Architecture

1. Risk Stratification Framework

Alert generation follows a three-tier severity classification: Critical (immediate intervention required), Warning (monitoring intensification needed), and Informational (trend awareness). Critical alerts trigger for glucose below 60 or above 250 mg/dL, heart rate below 40 or above 130 BPM, or predicted hypoglycaemia within 30 minutes. Warning alerts activate for borderline values or sustained unfavourable trends. Informational notifications highlight daily summaries and achievement of health goals.

Severity levels determine notification routing. Critical alerts simultaneously notify the patient (mobile push notification with audible alarm), emergency contacts (SMS), and the on-call healthcare provider (pager integration). Warning alerts reach the patient and care team via email. Informational updates appear in the mobile application dashboard without interrupting user activities.

2. Alert Fatigue Mitigation

Excessive alerts reduce user responsiveness and system trust. The architecture implements intelligent throttling where repeated similar alerts within 30-minute windows consolidate into a single notification with updated severity. Machine learning-based alert refinement analyses historical false positive rates, adjusting thresholds dynamically to maintain a false positive rate below 5% while preserving sensitivity above 95% for true emergencies.

User feedback mechanisms enable patients to indicate false alarms, feeding supervised learning models that personalise alert thresholds to account for individual physiology variations. This adaptive approach balances standardised medical protocols with personalised care requirements, optimising alert relevance and actionability.

4. Results and Analysis

A. System Performance Evaluation

Table I presents comprehensive performance metrics across system components evaluated during a 30-day pilot deployment with 25 diabetic participants. Data transmission reliability achieved a 99.2% success rate with average latency of 1.8 seconds from sensor reading to cloud availability. Sensor accuracy validation against laboratory reference measurements demonstrated ±12 mg/dL glucose accuracy, ±3 BPM heart rate accuracy, and ±0.3°C temperature accuracy, meeting clinical acceptability thresholds for continuous monitoring applications.

B. Clinical Outcome Analysis

Pilot study results demonstrate significant improvements in diabetes management metrics compared to the conventional periodic monitoring baseline. Participants experienced a 42% reduction in hypoglycaemic events (2.3 vs. 4.0 episodes per month), a 38% reduction in hyperglycaemic events (3.1 vs. 5.0 episodes per month), and a 47% reduction in emergency department visits (0.3 vs. 0.6 visits per month). Average HbA1c levels decreased from 8.2% to 7.4% over the 90-day monitoring period, indicating improved long-term glycaemic control.

Patient engagement metrics revealed an 89% daily system utilisation rate with average wear time of 18.2 hours per day. User satisfaction surveys (n = 25) showed 92% agreement with the statement "the system helps me manage my diabetes better" and 88% willingness to continue using the system long-term. Reported barriers included skin irritation from the wearable band (16% of users) and occasional connectivity issues in rural areas with poor Wi-Fi coverage (12% of users).

C. Alert System Effectiveness

The alert system generated 847 total notifications during the pilot period. Table II summarises the classification and trigger conditions for each alert type. Healthcare provider review of Critical alerts confirmed 296 true positives (94.9% precision) with 16 false positives, primarily caused by sensor displacement or motion artefacts. Alert response time analysis showed patients acknowledged Critical alerts within an average of 4.2 minutes, with 87% taking appropriate corrective action—consuming carbohydrates for hypoglycaemia or administering insulin for hyperglycaemia.

precision, coupled with average patient response times of 4.2 minutes, validates the practical clinical utility of the alert framework.

The 42–47% reductions in glycaemic events and emergency visits align with the economic projections motivating this research, supporting the hypothesis that continuous monitoring enables preventive intervention before complications escalate. The LSTM predictive model's 12.8 mg/dL mean absolute error provides actionable early warnings while remaining within clinically acceptable bounds. Observed barriers skin irritation and rural connectivity identify priority areas for hardware refinement and network infrastructure investment in future iterations.