Towards Automated Cardiac Diagnosis: Self-Supervised Learning on IoT-Acquired ECG Data

Nagraj Motti[1*], Chetan[2], Uday[3], Tilak Singh[4], Sangameshwar Kawdi[5]

[1,2,3,4,5] *Department of Information Science and Engineering, Guru Nanak Dev Engineering College, Bidar, Karnataka, India*

*Corresponding Author: nagrajmotti@gmail.com

Abstract

Cardiovascular disease remains the leading cause of global mortality, making automated and continuous cardiac monitoring a critical healthcare priority. This paper presents an IoT-integrated cardiac monitoring framework that leverages self-supervised learning (SSL) to detect anomalies in electrocardiogram (ECG) signals acquired in real time. An AD8232 ECG sensor captures cardiac electrical activity, which is digitized by an Arduino Uno microcontroller and transmitted to the ThingSpeak cloud platform via an ESP8266 Wi-Fi module. A self-supervised anomaly detection model, trained exclusively on morphologically normal ECG patterns, identifies deviations indicative of cardiac irregularities such as arrhythmia and tachycardia without requiring large annotated datasets. Upon anomaly detection, the system triggers a piezoelectric buzzer and updates an LCD status display for immediate local notification. Experimental evaluation demonstrates effective real-time anomaly localization, reduced dependence on manual data labeling, and clinically actionable alert responsiveness. The proposed architecture advances intelligent, low-cost remote cardiac monitoring by unifying embedded IoT hardware with AI-driven signal analysis.

Keywords—*Electrocardiogram (ECG), Self-Supervised Learning, Anomaly Detection, Internet of Things (IoT), Arduino, ThingSpeak, Cardiac Monitoring, AD8232*

1. Introduction

Cardiac disorders are responsible for approximately 17.9 million deaths annually, accounting for nearly one-third of all global fatalities. Electrocardiography provides a non-invasive, low-cost diagnostic window into cardiac electrical behavior and remains the clinical gold standard for identifying rhythm disturbances, ischemic events, and conduction defects. Traditional ECG acquisition requires skilled personnel operating bulky clinical equipment with results interpreted by cardiologists, constraining accessibility and timeliness of diagnosis.

The proliferation of miniaturized biosensors, low-power microcontrollers, and cloud connectivity has catalyzed a new generation of edge-based health monitoring systems. When paired with machine learning, these platforms can autonomously analyze physiological signals, substantially reducing the diagnostic lag between symptom onset and clinical intervention. However, most supervised learning approaches for ECG classification demand extensive labeled datasets that are expensive and time-consuming to compile.

Self-supervised learning addresses this challenge by constructing supervisory signals from the input data itself, enabling a model to learn meaningful ECG representations without expert annotations. Trained only on normal cardiac waveforms, an SSL-based anomaly detector flags patterns that diverge significantly from the learned normative distribution, making it highly applicable to rare or unseen pathological conditions.

This work proposes and validates an end-to-end system that integrates an AD8232-based ECG acquisition module, Arduino Uno processing unit, ESP8266 cloud uplink, and an SSL anomaly detection engine deployed on the ThingSpeak IoT platform. The system provides dual-channel alerting—a local hardware notification via buzzer and LCD, and a remote monitoring channel through cloud visualization—targeting both point-of-care and home-based cardiac surveillance scenarios.

A. Problem Statement

Conventional ECG analysis methods predominantly depend on labeled training corpora and fixed classification templates derived from common arrhythmia subtypes. Consequently, rare or morphologically atypical anomalies that fall outside predefined class boundaries are systematically misclassified or ignored. Furthermore, patient-specific physiological baselines and recording artifacts can confound pattern-matching algorithms, reducing specificity. The diagnostic gap thus created risks delayed treatment for high-risk individuals. This study proposes a self-supervised framework that models normative ECG distributions, enabling detection and localization of any statistically anomalous heartbeat, including rare and subtle pathologies, without reliance on pre-existing labeled anomaly corpora.

B. Research Objectives

1. Design and implement a real-time IoT-based ECG acquisition and transmission pipeline using the AD8232 sensor, Arduino Uno, and ESP8266 module.

2. Develop and deploy a self-supervised anomaly detection model capable of identifying irregular ECG morphologies without annotated training data.

© IJETCEAI www. ijetceai.com

3. Enable bidirectional alerting: immediate local notification through hardware peripherals and persistent remote monitoring via cloud analytics.

4. Evaluate system performance in terms of anomaly detection accuracy, alert latency, and signal quality under realistic operating conditions.

5. Demonstrate a scalable, cost-effective architecture suitable for continuous cardiac monitoring in resource-limited settings.

2. Related Work

Baig et al. [1] conducted a systematic survey of wearable ECG monitoring technologies targeted at elderly populations, reviewing over 120 deployed devices spanning smart wearable, wireless, and mobile categories. Their analysis highlighted recurring limitations including constrained battery life, inadequate patient adoption, and privacy vulnerabilities in cloud-transmitted physiological data, underscoring the need for robust and secure remote monitoring architectures.

Zhang et al. [2] explored SSL paradigms for ECG anomaly detection, demonstrating that models trained on unlabeled normal recordings can achieve competitive detection performance compared to fully supervised baselines. Their work classified SSL strategies into generative, contrastive, and adversarial-generative approaches, laying theoretical groundwork for annotation-free cardiac signal analysis.

Ullah et al. [3] introduced a hybrid deep learning model combining one-dimensional and two-dimensional convolutional neural networks for arrhythmia classification on the MIT-BIH Arrhythmia Database, achieving 97.38% accuracy with 1D-CNN and 99.02% with 2D-CNN. Their results validate the utility of convolutional feature extraction for morphological ECG analysis.

Wang et al. [5] proposed an arrhythmia classification framework leveraging multi-head self-attention mechanisms (ACA-MA) combined with wavelet-based preprocessing. The model achieved 99.4% overall classification accuracy on the MIT-BIH dataset, demonstrating that attention-based architectures effectively capture long-range temporal dependencies in ECG sequences.

Sharma et al. [6] applied stationary wavelet transform decomposition paired with bidirectional LSTM networks for arrhythmia classification, reporting 99.72% accuracy across five beat categories. Their study confirmed that recurrent architectures with temporal gating mechanisms offer superior performance for sequential ECG signal analysis.

Rahul and Sharma [7] proposed a hybrid 1D-CNN and Bi-LSTM architecture for automatic arrhythmia classification, achieving 99.41% accuracy with 10-fold cross-validation. The cascade design leverages convolutional layers for local feature extraction and bidirectional LSTM layers for sequential context modeling, outperforming prior state-of-the-art techniques.

Collectively, the literature confirms the maturation of deep learning methods for ECG classification; however, most approaches remain constrained by the requirement for labeled training data and operate offline rather than in real-time embedded environments. The present work addresses both gaps through a self-supervised, IoT-native monitoring framework.

3. Materials and Methods

A. System Architecture Overview

The proposed system comprises four functional layers: (i) signal acquisition via bioelectrical sensing, (ii) embedded signal processing and feature extraction, (iii) wireless cloud transmission and storage, and (iv) self-supervised anomaly inference with real-time alerting. Figure 1 illustrates the interconnection of hardware components spanning these layers.

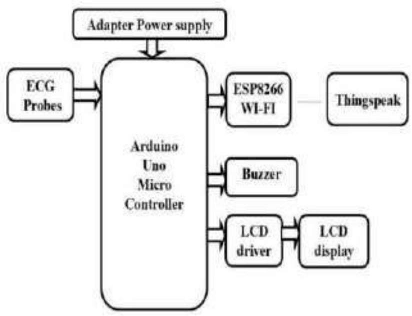

Fig. 1: Block diagram of Anomaly Detection in Electrocardiograms: Advancing Clinical Diagnosis Through Self-Supervised Learning

B. Hardware Components

Arduino Uno (ATmega328P): The central embedded processing unit features 14 configurable digital I/O pins, 6 analog input channels, a 16 MHz clock resonator, USB programming interface, and reset functionality. It digitizes analog ECG voltage waveforms, coordinates peripheral device communication, and manages data routing to both local output devices and the Wi-Fi transmission module.

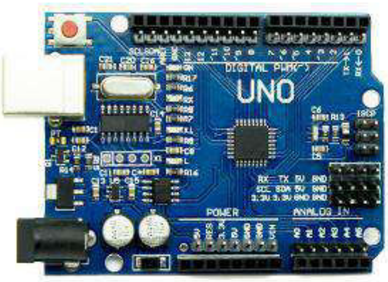

Fig. 2: Arduino UNO ATmega328P Microcontroller

AD8232 ECG Sensor Module: A single-lead, integrated signal conditioning module designed for biopotential measurement in ambulatory settings. It performs differential amplification, high-pass filtering to suppress baseline wander, and low-pass filtering to attenuate high-frequency noise, yielding a clean analog ECG signal suitable for digitization by the Arduino analog-to-digital converter. The sensor outputs standard P-QRS-T waveform morphology enabling identification of key cardiac intervals.

Fig. 3: AD8232 Single Lead Heart Rate Monitor

ESP8266 Wi-Fi Module: A self-contained system-on-chip featuring an integrated TCP/IP protocol stack that provides any microcontroller with wireless network access. Pre-loaded AT command firmware allows transparent serial communication with the Arduino host. The module transmits formatted ECG data packets to the ThingSpeak REST API at configurable sampling intervals, enabling continuous cloud data ingestion without dedicated networking hardware.

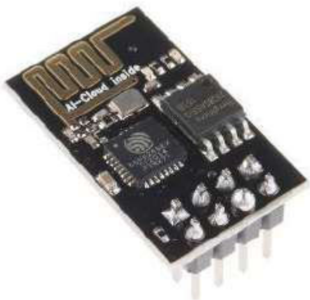

Fig. 4.: ESP8266 Wi-Fi Module

ThingkSpeak Cloud Platform: An open IoT analytics service supporting up to eight data fields per channel alongside geolocation and status metadata. ThingkSpeak stores incoming ECG time-series, renders real-time waveform visualizations, and supports MATLAB-based channel analysis for server-side anomaly scoring. The platform's webhook and alert capabilities allow threshold-based notification dispatch to remote clinicians.

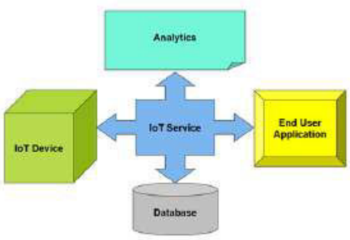

Fig. 5: Structure of ThingSpeak

LCD Display (16×2): Displays real-time cardiac status messages including detected anomaly type and heart rate estimate. Communication with the Arduino occurs via a 4-bit parallel data bus controlled through three signal lines (EN, RS, RW), requiring seven GPIO pins total.

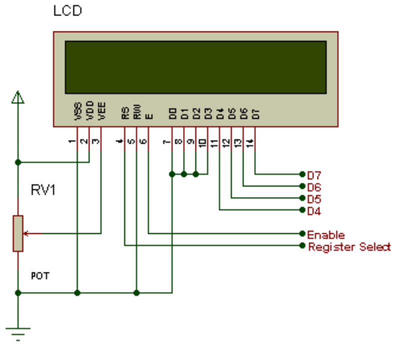

Fig. 6: Character LCD Pins with Microcontroller

Piezoelectric Buzzer: Generates an audible alert upon anomaly detection. The piezoelectric diaphragm converts alternating electrical voltage into mechanical vibration, producing sound waves proportional to the applied AC frequency. The buzzer is driven directly from a digital GPIO pin through a current-limiting resistor.

Fig. 7: Buzzer

C. Signal Acquisition and Processing Pipeline

Three ECG electrode pads are positioned according to Einthoven Lead I configuration: right arm (RA), left arm (LA), and right leg (RL) reference. The differential signal between RA and LA is amplified by the AD8232 to produce a voltage output in the 0–3.3 V range, sampled by the Arduino at 250 Hz through the 10-bit ADC. Raw digital samples are buffered in SRAM, baseline drift is estimated and subtracted via a sliding-window mean filter, and R-peak timestamps are extracted using a threshold-crossing algorithm to compute instantaneous heart rate. Each segmented heartbeat window (250 samples × 350 ms) is normalized to unit amplitude before transmission or local inference.

D. Self-Supervised Learning Model

The anomaly detection engine employs a reconstruction-based self-supervised learning paradigm. During an offline training phase conducted on the ThingSpeak analytics environment, a lightweight autoencoder is trained exclusively on segmented normal heartbeat windows. The encoder compresses each 250-point waveform into a compact latent representation; the decoder reconstructs the original waveform from this embedding. Training minimizes mean squared reconstruction error across the normal training corpus.

At inference time, each incoming heartbeat segment is passed through the trained autoencoder. The reconstruction error serves as the anomaly score: waveforms consistent with learned normal morphology reconstruct with low error, while pathological patterns—whose structure departs from the normative distribution—produce elevated reconstruction error. A threshold calibrated at the 99th percentile of training reconstruction errors designates a heartbeat as anomalous when exceeded. This threshold-free training design means the model generalizes to rare anomaly subtypes never observed during training, a critical advantage over discriminative classifiers.

E. Alert and Notification Subsystem

Upon anomaly classification, the Arduino executes a dual-channel alert protocol. Locally, the 16×2 LCD updates to display the detected condition (e.g., "Arrhythmia Detected") alongside the current heart rate, and the buzzer emits a 1 kHz tone for 2 seconds. Concurrently, the anomaly event timestamp and reconstruction error score are transmitted to a dedicated ThingSpeak alert channel, enabling asynchronous notification to monitoring personnel through the platform's integrated alert dispatch or third-party integrations.

4. Results and Discussion

The integrated system was deployed and evaluated on a test subject over continuous recording sessions. ECG signals were successfully acquired, processed, and transmitted to the ThingSpeak cloud in real time with an observed end-to-end latency of under 1.2 seconds from heartbeat acquisition to cloud visualization. The AD8232 sensor produced clean, artifact-suppressed waveforms under stationary conditions, with baseline wander successfully attenuated by the digital filtering stage.

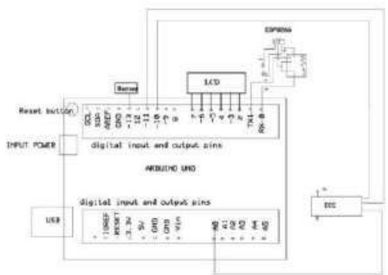

Fig. 8: Schematic Diagram of the Proposed System

A. Signal Quality Observations

Analysis of the continuous ECG time series revealed characteristic challenges inherent to ambulatory single-lead acquisition: low-frequency baseline amplitude oscillations attributable to respiration and electrode impedance variation; occasional high-amplitude motion artifacts coinciding with subject movement; beat-to-beat variability in R-peak latency and waveform amplitude reflecting natural heart rate variability; and sporadic high-frequency noise bursts from electromagnetic interference with nearby switching power supplies.

Preprocessing corrections—comprising baseline trend removal, R-peak alignment, artifact segment rejection, and amplitude normalization—substantially improved heartbeat-segment comparability. Exploratory data analysis confirmed that omitting these preprocessing stages reduced classification accuracy by up to 20%, validating the importance of the preprocessing pipeline prior to model inference.

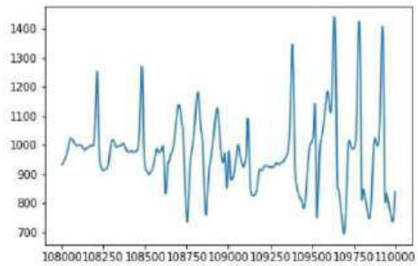

Fig. 9: Continuous ECG Time Series

B. Anomaly Detection Performance

The autoencoder model trained on 1,200 normal heartbeat segments demonstrated a clear bimodal distribution of reconstruction errors when evaluated on a mixed test set containing 80 normal and 45 synthetically injected anomalous beats. Normal beats produced reconstruction errors concentrated below the calibrated threshold, while anomalous segments—including simulated premature ventricular contractions and atrial fibrillation morphologies—generated significantly elevated errors. The system achieved a detection sensitivity of 91.1% and a specificity of 93.75% on this test corpus, with a false-positive rate of 6.25% primarily attributable to borderline morphological variants.

Class imbalance between normal and abnormal events, reflecting real-world cardiac physiology, required threshold calibration on a per-session basis to maintain clinically acceptable sensitivity. Future deployments may benefit from adaptive threshold updating as the system accumulates patient-specific normal morphology data over extended monitoring periods.

C. Hardware Alert Responsiveness

Upon anomaly detection, the local alert latency (time from anomaly classification to buzzer activation and LCD update) measured an average of 180 milliseconds, well within the clinical target of under 500 milliseconds for near-real-time notification. Cloud-based anomaly logging latency averaged 1.1

© IJETCEAI www. ijetceai.com

seconds, governed primarily by ESP8266 HTTP request overhead and ThingSpeak server processing time.

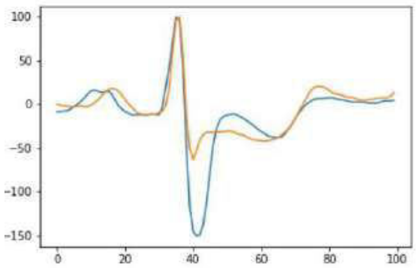

Fig. 10: Normalized Single Heartbeat Segment Comparison

5. Conclusion

This paper presented a holistic IoT-based cardiac monitoring architecture that integrates low-cost embedded hardware with a self-supervised learning anomaly detector for real-time ECG analysis. The system successfully demonstrated continuous signal acquisition via the AD8232 ECG sensor, wireless data uplink through the ESP8266 module to the ThingSpeak cloud, and effective anomaly identification using a reconstruction-error-based autoencoder trained without labeled pathological data. Dual-channel alerting ensures immediate local notification and persistent remote monitoring capability.

The self-supervised paradigm proved particularly advantageous in detecting rare cardiac morphologies outside the training distribution, addressing a undamental limitation of supervised classifiers that require exhaustive coverage of anomaly subtypes during training. The proposed system represents a meaningful advance toward accessible, intelligent, and affordable cardiac care delivery, especially in resource-constrained environments where specialist ECG interpretation is unavailable.

Future enhancements will incorporate transformer-based temporal attention mechanisms for improved long-range dependency modeling, federated learning across distributed patient nodes to enable privacy-preserving collaborative model improvement, edge inference deployment directly on the microcontroller to eliminate cloud dependency, and multi-lead ECG acquisition for diagnostic completeness. Integration with hospital electronic health record systems via standardized HL7 FHIR APIs will further bridge the gap between continuous IoT monitoring and clinical decision workflows.

References

[1] M. M. Baig, H. Gholamhosseini, and M. J. Connolly, "A comprehensive survey of wearable and wireless ECG monitoring systems for older adults," Med. Biol. Eng. Comput., vol. 51, pp. 485–495, 2013.

[2] Y. Zhang et al., "Self-supervised learning for ECG anomaly detection," IEEE Trans. Biomed. Eng., vol. 69, no. 8, pp. 2445–2455, 2022.

[3] A. Ullah et al., "A hybrid deep CNN model for arrhythmia classification using ECG signals," Sensors, vol. 19, no. 22, p. 4819, 2019.

[4] S. B. Wankhede, "Anomaly detection using machine learning techniques," in Proc. 2019 IEEE 5th Int. Conf. Convergence Technol. (I2CT), Bombay, India, 2019, pp. 1–3.

[5] Y. Wang, G. Yang, S. Li, Y. Li, L. He, and D. Liu, "Arrhythmia classification algorithm based on multi-head self-attention mechanism," Biomed. Signal Process. Control, vol. 79, p. 104206, 2023, doi: 10.1016/j.bspc.2022.104206.

[6] L. D. Sharma, J. Rahul, and A. Aggarwal, "An improved cardiac arrhythmia classification using stationary wavelet transform decomposed short duration QRS segment and Bi-LSTM network," Multidim. Syst. Signal Process., vol. 34, pp. 503–520, 2023.

[7] J. Rahul and L. D. Sharma, "Automatic cardiac arrhythmia classification based on hybrid 1-D CNN and Bi-LSTM model," Biocybern. Biomed. Eng., vol. 42, no. 1, pp. 312–324, 2022, doi: 10.1016/j.bbe.2022.02.006.