International Journal of Engineering Research & Technology (IJERT)
ISSN: 2278-0181
Vol. 15 Issue 09 , September - 2026

Edge Computing Based Wearable Cardiac Arrhythmia Detection System

A Systematic Literature Survey Prepared for Undergraduate Research

Ravikiran B A , Chandana N E , Devika Nataraj , Kishor Vs, Kruthika J V
K S School of Engineering and Management,Bangalore

Abstract - Cardiac arrhythmias remain a leading contributor to global cardiovascular mortality, and continuous monitoring outside the clinic has become central to early detection. This review examines the convergence of edge computing and the Internet of Things (IoT) in wearable electrocardiogram (ECG) systems for arrhythmia monitoring, with particular attention to three-channel acquisition using commercial-grade analog front ends. It surveys system architectures spanning cloud, fog, and edge deployment; classification approaches ranging from handcrafted-feature machine learning to deep convolutional and recurrent networks; and TinyML techniques that bring inference onto microcontroller-class hardware. Wearable ECG acquisition hardware built around chips such as the AD8232, ADS1292R, and ADS1293 is reviewed alongside the communication protocols, Bluetooth Low Energy, Wi-Fi, and LoRa, used to move data between sensor, edge node, and cloud. A comparative table positions eight representative prior implementations against channel count, communication method, computation location, classification algorithm, and reported accuracy. A second analysis draws on the lead configurations recorded in the MIT-BIH, PTB-XL, Chapman-Shaoxing, and related public databases to justify Lead I, Lead II, and Lead V1 as the three-channel configuration best supported by prior hardware and by the largest annotated public datasets. The review closes by identifying gaps in dataset availability for true three-lead simultaneous recording, in energy-accuracy trade-offs for on-device inference, and in standardized benchmarking across embedded platforms.

Index Terms - Arrhythmia classification, edge computing, electrocardiogram, Internet of Medical Things, TinyML, wearable sensors.

I. INTRODUCTION

Cardiovascular disease caused an estimated 19.2 million deaths worldwide in 2023, up from 13.1 million in 1990, and now accounts for close to one in three deaths globally [1]. Arrhythmias, irregular or abnormal heart rhythms, if the case of arrhythmia is detected early, hence can be treated to the earliest. This occurs in inconsistent nature: a standard twelve-lead ECG taken during a routine clinic visit will often miss the event entirely. This has driven three decades of research into portable and wearable monitoring, from the Holter recorder through to smartwatches and adhesive chest patches capable of continuous, multi-day recording [2]–[5].

It helps with two technology shifts, which have reshaped this condition over the past decade. The first is the Internet of Things (IoT), which turns a wearable sensor into a networked node capable of streaming physiological data to a smartphone, or a cloud service for storage and analysis. The second is edge computing, which pushes some or all of the signal processing and classification workload back onto the wearable device. Because the continuous transmission of raw ECG data drains battery life quickly, hence where they introduced delay that limits the usefulness of real-time alerts, and raises privacy concerns each time physiological data leaves the body [6], [7]. Tiny Machine Learning (TinyML), the practice of compressing and implemented on trained neural networks onto microcontroller-class hardware with kilobytes rather than gigabytes of memory, therefore the shift is possible. The demonstration of arrhythmia classification runs directly on Arduino, ESP32, and ARM Cortex-M platforms [6], [7].

Commercial-grade analog front-end (AFE) chips such as the Texas Instruments ADS1293 now make genuine three-channel acquisition practical on a coin-cell-powered wearable, and at least one prior study has already built a working prototype around this exact hardware [8]. The missing part is that the synthesis which connects hardware capability, communication and computation architecture, classification method, and lead choice into a single evidence-based picture for a team designing a three-channel system from the ground up. This review sets out to provide that synthesis.

The further information of the paper follows the structure as mentioned accordingly. Section II lays out the review methodology; Section III surveys cloud, fog, and edge architectures for cardiac IoT; Section IV covers TinyML and on-device inference; and Section V reviews classification algorithms. Section VI examines acquisition hardware and channel configurations, Section VII covers communication protocols, and Section VIII catalogues public ECG datasets. Section IX presents a comparative table of prior

International Journal of Engineering Research & Technology (IJERT)
ISSN: 2278-0181
Vol. 15 Issue 09 , September - 2026

implementations, and Section X analyses lead selection for three-channel systems. Section XI addresses power management, Section XII covers security, privacy, and federated learning, Section XIII identifies research gaps, and Section XIV concludes.

II. REVIEW METHODOLOGY

This review is based on peer-reviewed journal articles and conference proceedings published mainly from 1985 to 2026, sourced from IEEE Xplore, ScienceDirect, SpringerLink, MDPI, Nature Portfolio journals, PubMed, PLOS, and arXiv preprint listings curated by reputable research groups.The search terms integrated "edge computing," "fog computing," "IoT," "TinyML," and "wearable" with "ECG," "electrocardiogram," and "arrhythmia." References such as books, product datasheets, and non-peer-reviewed web tutorials were omitted from the reference list.Emphasis was placed on studies that presented a functioning hardware prototype, a recognized public dataset, or a quantitative comparison with previous systems, as these elements provide the essential details required for the comparative analysis in Sections IX and X.

III. SYSTEM ARCHITECTURES: CLOUD, FOG, AND EDGE COMPUTING FOR CARDIAC IOT

In the early stages of IoT-based ECG monitoring, a predominantly cloud-centric model was utilized: a wearable sensor transmitted either raw or minimally processed data to a smartphone or gateway, which then forwarded it to a remote server for storage, visualization, and classification. While this model is straightforward to implement and scales effectively for retrospective analysis, it falls short in meeting the two critical requirements for arrhythmia detection: low latency for time-sensitive events like ventricular tachycardia, and low, predictable bandwidth consumption for battery-operated devices that need to function for days or weeks without recharging.

To address these challenges, fog computing was introduced as an intermediary layer to alleviate some of the burdens. Azimi et al. proposed HiCH, a hierarchical fog-assisted architecture for healthcare IoT that conducts initial filtering and feature extraction at a fog node situated near the patient, thereby reserving the cloud for long-term data storage and population-level analytics [9]. Building on this concept, Rincon et al. developed a fog-based cardiovascular monitoring system that transmits single-lead ECG data over LoRa to a fog service equipped with a deep learning classifier, achieving 90 percent accuracy in detecting atrial fibrillation using the 2017 PhysioNet Challenge dataset, while ensuring that the wearable node was limited to data acquisition and transmission [10].

Edge computing advances this model further by relocating the classification algorithm directly onto the wearable device or an attached microcontroller, thereby removing the network delay for the crucial decision of whether the current heartbeat is normal. Hizem et al. showcased this method using a convolutional neural network trained on the MIT-BIH database, which was subsequently pruned and quantized for deployment on a Raspberry Pi and Arduino Nano combination connected to an AD8232 sensor, achieving 99 percent accuracy on the target device with an estimated power consumption of 0.24 microwatts per inference after optimization [6]. This approach has become prevalent in the literature: acquire, filter, classify on the device, and transmit only alerts or a compressed summary.

IV. TINYML AND ON-DEVICE ARRHYTHMIA CLASSIFICATION

Edge-based arrhythmia monitoring utilizes TinyML technology. Abadade et al. provide a comprehensive overview of the field, outlining the workflow that begins with model design, followed by compression techniques such as pruning, quantization, or knowledge distillation. This is then succeeded by conversion to a runtime environment like TensorFlow Lite for Microcontrollers, culminating in deployment to a target device with kilobytes of RAM instead of megabytes [7]. Specifically addressing ECG applications, Kim et al. introduce TinyCES, a convolutional model trained on MIT-BIH and PTB datasets, which has been compressed to a 27-kilobyte size, enabling it to operate on an Arduino-class board with approximately 97 percent detection accuracy [11]. Zambrano-de la Torre et al. present a similar feasibility study that employs a quantized one-dimensional CNN on an 8-bit Arduino Uno, interfacing with an AD8232 front end, and simplifies the problem to three classes: normal, ventricular, and supraventricular, to accommodate the limitations of genuinely low-cost hardware [12]. Alimbayeva et al. also combine custom wearable ECG acquisition devices with an embedded machine learning classifier for ongoing heart rhythm monitoring [13]. This trend is consistent throughout the literature: sensors and inference capabilities coexist on a single low-power device.

Recent advancements have moved beyond convolutional models. Busia et al. implement a six-thousand-parameter transformer model for five-class arrhythmia detection on the GAP9 ultra-low-power processor, achieving an accuracy of 98.97 percent on MIT-BIH with an inference time of 4.28 milliseconds and an energy expenditure of 0.09 millijoules per beat. They further demonstrate that training with simulated electrode motion artifacts maintains accuracy above 98 percent in noisy, post-deployment scenarios [14]. Amiri et al. adopt a complementary approach, concentrating not on the classifier but on the data itself: their SIC-EDGE scheme semantically compresses ECG data at the edge prior to any wireless transmission, thereby reducing the bandwidth and energy costs

International Journal of Engineering Research & Technology (IJERT)
ISSN: 2278-0181
Vol. 15 Issue 09 , September - 2026

associated with the communication link that TinyML classification aims to minimize in the first place [15]. Collectively, this body of work demonstrates that achieving five-class, MIT-BIH-style classification is feasible within the memory and energy constraints of a coin-cell-powered wearable device. Nevertheless, it is important to note that nearly all of the validation was conducted using a single lead, and expanding to three concurrent channels will require separate re-verification. Overall, this research indicates the potential for effective classification while highlighting the need for further validation in multi-channel scenarios.

V. ECG CLASSIFICATION ALGORITHMS: FROM HANDCRAFTED FEATURES TO DEEP LEARNING

Automated classification of arrhythmias has existed for decades, predating the contemporary literature on wearables and edge computing. Much of the current methodology is built upon principles established long before the advent of the Internet of Things (IoT). The real-time QRS detection algorithm developed by Pan and Tompkins in 1985 remains a benchmark for beat segmentation, accurately identifying 99.3 percent of QRS complexes in the MIT-BIH database through the use of a digital bandpass filter and adaptive thresholding, which is still cost-effective enough to be implemented on modern microcontrollers [16]. After segmenting the beats, classification traditionally depended on manually crafted morphological and statistical features that were input into classifiers such as support vector machines or random forests; Luz et al. provide a comprehensive review of these earlier heartbeat classification techniques [17].

The transition to deep learning commenced with the application of one-dimensional convolutional networks directly to raw or minimally filtered beat segments. Kiranyaz, Ince, and Gabbouj showcased a patient-specific 1-D CNN that adjusts to the unique beat morphology of each individual, enabling real-time classification that is appropriate for continuous monitoring [18]. Raj and Ray further developed this patient-specific approach into a comprehensive personalized arrhythmia monitoring system, modifying the classifier to align with an individual's baseline rhythm instead of depending solely on a model trained on a population level [19]. Acharya et al. utilized a convolutional network to analyze various tachycardia segment lengths from the MIT-BIH database [20]. The most compelling evidence of clinical-grade performance was provided by Hannun et al., whose deep neural network achieved cardiologist-level accuracy across twelve rhythm classes using a large, privately gathered single-lead ambulatory dataset, significantly contributing to the establishment of deep learning as the standard method for this application [21].

Hybrid architectures that integrate convolutional feature extraction with recurrent layers have gained popularity in scenarios where temporal dependencies across multiple beats are significant, such as in differentiating atrial fibrillation from normal sinus rhythm. Alamatsaz et al. introduce a lightweight CNN-LSTM model specifically tailored for edge deployment, achieving competitive accuracy while incorporating an explainability layer that identifies which segments of the beat influenced the classification [22]. Cinar and Tuncer merge LSTM with hybrid CNN-SVM networks to differentiate normal rhythm, arrhythmia, and congestive heart failure from ECG signals [23]. Both Ma et al. and Pandey et al. utilize CNN-LSTM hybrids specifically for the detection of atrial fibrillation, an arrhythmia of particular concern for wearable monitoring due to its often paroxysmal and easily overlooked presentation [24], [25]. Yu et al. advance single-lead AF detection to accommodate the brief, noisy recordings typical of consumer wearables, employing a specially designed deep architecture assessed on precisely this type of data [26]. Two recent systematic reviews, one focusing on embedded feasibility and standardized inter-patient evaluation, and the other on deep learning methods in general, arrive at the same conclusion: reported accuracy varies significantly based on whether a study employs patient-independent or inter-patient splits. Future research must include reporting on inference time, memory usage, and energy consumption alongside accuracy, or else comparisons across the wearable and embedded literature will continue to be incommensurable [27], [28].

VI. WEARABLE ECG ACQUISITION HARDWARE AND CHANNEL CONFIGURATIONS

A. Single-Channel Analog Front Ends

The AD8232 from Analog Devices has emerged as a preferred option for economical, single-lead wearable ECG prototypes. This chip combines instrumentation amplification, right-leg-drive circuitry, and bandpass filtering into a compact, low-power solution, consuming approximately 170 microamperes, which allows for prolonged continuous use from a small lithium battery. Wu, Redouté, and Yuce developed a smart T-shirt utilizing the AD8232 that transmits ECG data via Bluetooth Low Energy to a smartphone, achieving a signal correlation exceeding 99 percent when compared to a reference Holter system, with heart-rate accuracy around 2 percent [29]. Numerous studies have utilized the same chip, or the closely related dual-channel ADS1292R, to transition from a single differential channel to authentic two-lead acquisition, which is appropriate for standard Lead I and Lead II recordings.

B. Multi-Channel Analog Front Ends and the Path to Three-Channel Acquisition

Three simultaneous ECG channels, each independently digitized, require either three single-channel AFEs connected to a shared microcontroller or a dedicated multi-channel chip. The Texas Instruments ADS1293 was specifically designed for this purpose: it

International Journal of Engineering Research & Technology (IJERT)
ISSN: 2278-0181
Vol. 15 Issue 09 , September - 2026

consolidates three 24-bit channels, right-leg-drive generation, and lead-off detection into a single 5-millimeter package while consuming less than 300 microwatts per channel. Fan et al. created one of the earliest wearable prototypes around this chip, integrating it with an MSP430 microcontroller to capture Lead I, Lead II, and Lead V1 simultaneously and transmit the data to a phone or PC, achieving 85 percent accuracy and 88.23 percent precision across twenty validation trials [8]. This specific combination of limb leads I and II with the precordial lead V1 is directly relevant to the channel-selection question addressed in Section X, as it constitutes a validated three-channel configuration based on the exact type of commercial-grade AFE that this review assumes.

In addition to designs based on the ADS1293, various research groups have reported on comprehensive low-power wearable ECG sensor nodes that tackle associated hardware challenges. These include minimizing skin-electrode contact impedance through the use of dry or textile electrodes, decreasing power consumption via custom low-power signal chains, and establishing human-body communication links that completely bypass traditional radio-frequency transmission. Nemati, Deen, and Mondal detail a wireless wearable sensor specifically designed for long-term applications [30], while Jani, Bagree, and Roy introduce a combined low-power, low-cost ECG and EMG sensor intended for wearable biometric and medical purposes [31]. Khalaf and Abdoola present a wireless body sensor network integrated with an Android application, creating a complete eHealth pipeline [32]. Krachunov et al. approach comfort from a different perspective, employing painted electrodes to alleviate the discomfort associated with conventional gel electrodes during extended monitoring periods [33]; meanwhile, Shin et al. merge ECG and bioimpedance sensing into a single two-electrode device [34]. Finally, Coulter et al. complete this hardware overview with a low-power IoT platform designed for continuous vital signs monitoring, as opposed to sporadic recording sessions [35].

VII. COMMUNICATION ARCHITECTURES FOR IOT-BASED ECG MONITORING

The selection of wireless link significantly influences both the power budget and the potential range of a wearable ECG system, with the literature primarily focusing on three protocols. Bluetooth Low Energy is the leading choice for personal, single-user monitoring within a few meters of a connected smartphone; Del Campo et al. specifically investigate the coexistence and management of connectionless and connection-oriented BLE traffic on resource-constrained devices typical of wearable sensor nodes [36]. Wi-Fi, often implemented through an ESP32 or NodeMCU module, is utilized when a fixed gateway is accessible, sacrificing higher power consumption for seamless integration with existing home or hospital networks and cloud dashboards.

LoRa and other low-power wide-area protocols cater to a third scenario: monitoring in areas lacking local Wi-Fi or cellular infrastructure, such as in rural telemedicine applications. The fog-based cardiovascular monitoring system discussed by Rincon et al. in Section III employs LoRa to connect to a fog node from a remote or rural location where a traditional gateway would be impractical [10]. This advantage in range is accompanied by significantly lower data rates compared to BLE or Wi-Fi, which underscores the necessity for on-device classification: a LoRa link is generally inadequate for continuous streaming of raw multi-channel ECG data, necessitating that the wearable performs sufficient local processing to condense its output into occasional summaries or alerts.

VIII. PUBLIC ECG DATASETS AND DATABASES

The datasets utilized in a study for training and validation significantly influence both its reported accuracy and, for the context of this review, its relevance to the design of three-channel wearables. Table II in Section X is heavily reliant on the channel configuration recorded by each database, which is summarized here.

The MIT-BIH Arrhythmia Database is the most extensively utilized benchmark in this field. It comprises 48 half-hour, two-channel recordings from 47 subjects, annotated by cardiologists in accordance with the AAMI beat classification standard [37]. It is important to note for the discussion in Section X that MIT-BIH is a two-channel database, not a three-channel one: in forty-five of the forty-eight records, the first channel corresponds to a modified Lead II, while the second channel is typically a modified V1, though it may also be V2, V4, or V5 on occasion.

Several large databases released more recently capture the complete standard twelve-lead configuration, allowing for the extraction of any three-lead subset for training and validation, despite the original recordings containing more channels than a three-channel wearable would utilize. The PTB-XL database includes 21,837 twelve-lead recordings from 18,885 patients, accompanied by seventy-one annotated diagnostic statements [38]. The Chapman-Shaoxing database contributes an additional 10,646 twelve-lead recordings that encompass eleven rhythm classes, with a particular emphasis on atrial fibrillation and related supraventricular arrhythmias [39]. The China Physiological Signal Challenge 2018 database offers another open-access collection of twelve-lead recordings, specifically curated for algorithms aimed at detecting rhythm and morphology abnormalities [40]. Furthermore, the Georgia twelve-lead database, which was released as part of the PhysioNet/Computing in Cardiology Challenge 2020, adds further regional diversity to this collection [41]. Lastly, the St. Petersburg INCART database, along with the broader PhysioNet/Physio Bank archive that hosts it, provides additional recordings derived from ICU settings, predominantly featuring fewer than three

International Journal of Engineering Research & Technology (IJERT)
ISSN: 2278-0181
Vol. 15 Issue 09 , September - 2026

simultaneous leads per historical record [42]. The dataset from the 2017 PhysioNet Challenge, designed for single-lead recordings, has established itself as the benchmark for atrial fibrillation screening using short, consumer-grade recordings. Several studies on fog and edge computing discussed in Section III validate their findings against this dataset directly [43].

For a three-channel wearable project, the significant issue is that there is no widely recognized public database that records Lead I, Lead II, and V1 simultaneously as a matched trio, apart from a limited number of studies utilizing ADS1293-based hardware, such as the work by Fan et al. [8]. Consequently, a project focused on this configuration will likely need to pretrain or validate using the two-channel MLII/V1 subset from MIT-BIH, augment this with the corresponding three leads obtained from a twelve-lead source like PTB-XL or Chapman-Shaoxing, and gather a modest amount of matched in-house three-channel data for final validation. This gap is further elaborated as a research priority in Section XIII.

IX. COMPARATIVE ANALYSIS OF PRIOR IMPLEMENTATIONS

Table I illustrates eight representative systems discussed earlier, organized according to the critical dimensions for a new three-channel wearable design: sensor and analog front end, channel count, communication method, the site of classification execution, the algorithm applied, the validation dataset referenced, and the main findings presented by the original authors.

Study	Sensor / AFE	Ch.	Communication	Compute Location	Algorithm	Dataset	Reported Result
Hizem et al. [6], 2025	AD8232	1	Local (UART)	Edge (Raspberry Pi 4 + Arduino Nano)	CNN, pruned + quantized	MIT-BIH	99.0% accuracy; 0.24 µW/inference
Rincon et al. [10], 2020	Single-lead ECG + ESP32	1	LoRa	Fog node	Merged MobileNet (2×)	PhysioNet Challenge 2017	90.0% accuracy (AF)
Kim et al. [11] (TinyCES), 2023	Not specified	1	Not specified	Edge (Arduino-class)	CNN, 27 KB footprint	MIT-BIH + PTB	~97% detection ratio
Zambrano-de la Torre et al. [12], 2026	AD8232	1	Not specified	Edge (Arduino Uno, 8-bit)	1D-CNN, quantized	MIT-BIH (3-class)	Feasibility validated
Busia et al. [14], 2024	Not specified (silicon eval.)	1	Not specified	Edge (GAP9 processor)	Tiny Transformer, 6K params	MIT-BIH (5-class)	98.97% accuracy; 0.09 mJ/inference
Amiri et al. [15] (SIC-EDGE), 2022	Not specified	1	Edge-assisted, pre-transmission	Edge	Semantic ECG compression	Not specified	Compression-focused (bandwidth/energy metric)
Fan et al. [8], 2017	ADS1293	3 (I, II, V1)	Wired (USB/serial to PC or phone)	Local display / MATLAB	Threshold-based detection	In-house (20 trials)	85.0% accuracy; 88.23% precision
Wu, Redouté & Yuce [29], 2019	AD8232	1	BLE	Smartphone (display)	Signal display, no on-device classifier	Reference Holter comparison	99.23% signal correlation

TABLE I. Comparative Analysis of Prior Wearable/IoT Arrhythmia Monitoring Implementations

Three key patterns are highlighted in this review. The highest accuracy figures are attributed to single-lead systems that have been trained and tested on well-curated benchmark data. In contrast, the only genuine three-channel hardware prototype noted in this review demonstrates the lowest accuracy of the group, underscoring the fact that three-channel systems are still relatively under-optimized compared to their single-lead counterparts. Additionally, edge-resident classification has become the standard practice among systems published from 2022 onwards, marking a shift from the previously prevalent cloud- and fog-centric approach seen in earlier research. Furthermore, the choice of communication protocol is closely aligned with the deployment context: BLE is utilized for personal, short-range monitoring, while LoRa is specifically employed in areas lacking conventional infrastructure.

International Journal of Engineering Research & Technology (IJERT)
ISSN: 2278-0181
Vol. 15 Issue 09 , September - 2026

X. SELECTING THE OPTIMAL THREE-LEAD CONFIGURATION

Selecting three out of the twelve standard ECG leads for recording necessitates a balance among diagnostic coverage, the number of electrodes used, and the compatibility with public datasets that a wearable's classifier is expected to be pretrained or benchmarked against. Table II provides a summary of the lead configurations documented across the primary sources discussed in Section VIII.

Source	Type	Channels / Leads Recorded	Relevance to Three-Channel Design
MIT-BIH Arrhythmia Database [37]	Two-channel Holter	Modified Lead II (~all records) + V1 (~45/48 records; V2, V4, or V5 otherwise)	Largest annotated beat-level benchmark; two channels only
PTB-XL [38]	12-lead clinical	Standard I, II, III, aVR, aVL, aVF, V1–V6	21,837 recordings; any 3-lead subset directly extractable
Chapman-Shaoxing [39]	12-lead clinical	Standard 12-lead	10,646 recordings; strong AF/SVT rhythm focus
CPSC2018 [40]	12-lead clinical	Standard 12-lead	Open rhythm/morphology abnormality benchmark
Georgia 12-lead (Challenge 2020) [41]	12-lead clinical	Standard 12-lead	Adds regional/demographic diversity to the 12-lead pool
INCART / PhysioNet archive [42]	ICU, variable protocol	Variable; often fewer than 3 simultaneous leads historically	Supplementary ICU-derived recordings
PhysioNet Challenge 2017 [43]	Single-lead, consumer-grade	Lead I only (AliveCor-style)	Standard AF-screening benchmark
Fan et al. hardware prototype [8]	Wearable, ADS1293 AFE	Lead I, Lead II, Lead V1 (simultaneous)	Only prior 3-channel commercial-AFE wearable identified

TABLE II. Lead Configurations Recorded in Public ECG Databases and Prior Three-Channel Hardware

Diagnostic coverage presents a compelling argument. Lead II is the preferred option for rhythm analysis and P-wave visibility due to its axis being nearly parallel to the heart's mean electrical vector, resulting in the most consistently pronounced QRS complex among all limb leads. This is precisely why the MIT-BIH database employs a modified Lead II as its primary channel in the vast majority of recordings [37]. Lead I provides a second, approximately orthogonal limb-lead perspective, which, in conjunction with Lead II, is sufficient to derive Lead III and the augmented limb leads through Einthoven's and Goldberger's relationships without the need for additional electrodes, as all limb-derived leads are linear combinations of the same fundamental cardiac vector. A precordial lead contributes what the limb leads cannot structurally capture, particularly in identifying bundle branch block and localized ischemic changes. V1 is the standard precordial lead utilized to differentiate between ventricular and supraventricular ectopy based on QRS morphology, which is crucial for arrhythmia screening.

The hardware precedent supports this conclusion. According to the findings of Fan et al., their ADS1293-based prototype is, to the best of this review's knowledge, the only previously documented wearable device that operates on a commercial-grade three-channel AFE, recording Lead I, Lead II, and V1 simultaneously. It achieves this with a standard four-electrode configuration (right arm, left arm, left leg, and a right-leg-drive reference) rather than a more complex precordial array [8]. This layout is significantly easier to maintain in a wearable format compared to any setup that requires two or more precordial electrodes distributed across the chest, each needing its own stable adhesive contact point to prevent shifting during daily activities.

The primary consideration, however, is the compatibility of the dataset. Since the second channel in most MIT-BIH records is V1, a wearable device that captures Lead II and V1 can directly access the most extensive and well-annotated open arrhythmia benchmark available, eliminating the need for any lead-reconstruction process. Incorporating Lead I as a third channel incurs no additional cost: it is readily obtainable from every twelve-lead source examined, and in conjunction with Lead II, it enables the reconstruction of the remaining limb leads post hoc if required for future research.

Lead	Electrode Placement	Primary Diagnostic Role	Dataset / Hardware Support
Lead I	Right arm (−) to left arm (+)	Second orthogonal limb view; enables derivation of Lead III, aVR, aVL, aVF	Present in every 12-lead public dataset reviewed (PTB-XL, Chapman-Shaoxing, CPSC2018, Georgia)

International Journal of Engineering Research & Technology (IJERT)
ISSN: 2278-0181
Vol. 15 Issue 09 , September - 2026

Lead	Electrode Placement	Primary Diagnostic Role	Dataset / Hardware Support
Lead II	Right arm (−) to left leg (+)	Primary rhythm and P-wave lead; most prominent QRS among limb leads	Primary channel in nearly all MIT-BIH records (as modified Lead II)
Lead V1	Precordial, 4th intercostal space, right sternal border	Distinguishes ventricular from supraventricular ectopy by QRS morphology; detects bundle branch block	Second most common MIT-BIH channel (~45/48 records); used in Fan et al.'s [8] ADS1293 3-lead wearable

TABLE III. Recommended Three-Channel Lead Configuration and Rationale

The recommended three-channel configuration for a wearable arrhythmia system utilizing commercial-grade hardware like the ADS1293 consists of Lead I, Lead II, and V1. This setup aligns with the sole previous three-channel hardware implementation identified in this review, maintains compatibility with the MIT-BIH benchmark for beat-level classification training, and ensures comprehensive access to all significant twelve-lead public databases for additional training or external validation. Notably, it requires only the standard four-electrode placement instead of a more complex precordial array.

XI. POWER MANAGEMENT AND ENERGY HARVESTING

Battery life frequently serves as the practical limitation that determines whether a wearable ECG system can provide true continuous monitoring or merely a brief recording session.

The pruning and quantization methods outlined in Section IV significantly reduce inference energy on the classification front: Hizem et al. indicate a reduction in inference time from 0.333 to 0.200 milliseconds following the application of combined pruning and quantization, along with a corresponding decrease in power consumption [6]. Additionally, Busia et al. report an energy expenditure of 0.09 millijoules per inference on specialized ultra-low-power silicon [14].

On the hardware side, various research teams have explored energy harvesting techniques to extend battery life or eliminate the need for replacements entirely. Dąbrowska et al. assess a photovoltaic harvesting system integrated into smart clothing designed for mountain rescue scenarios, which shares a similar power-availability profile with continuous ambulatory ECG monitoring [44]. Ivanov investigates a thermoelectric method suitable for wrist-worn devices that can utilize the temperature difference between the skin and the surrounding air [45]. Wan et al. present a comprehensive wearable IoT platform for real-time health monitoring, developed with these energy-efficient design principles from the outset [46]. In the case of a three-channel system, the additional power requirement for a third acquisition channel on a chip like the ADS1293, estimated at approximately 300 microwatts per channel according to the manufacturer's specifications, is relatively minor compared to the energy cost of wireless transmission. This supports the edge-first argument made in Section III: achieving power savings is more effectively accomplished by transmitting less data rather than reducing the number of acquisition channels.

XII. SECURITY, PRIVACY, AND FEDERATED LEARNING

Continuous ECG monitoring produces a comprehensive, uniquely identifiable physiological record, and transferring that record away from the body introduces privacy issues that increase with the number of transmission hops involved. Al-Nbhany, Zahary, and Al-Shargabi examine blockchain-based methods for securing IoT healthcare data in a broad context, detailing both the interoperability and computational overhead trade-offs associated with integrating a distributed ledger layer into a resource-limited sensor pipeline [47]. Baucas, Spachos, and Plataniotis merge federated learning with blockchain technology in a fog-IoT platform specifically designed for predictive healthcare utilizing wearable data, an architecture that retains raw physiological signals on the device while still enabling a shared model to enhance across a population [48].

The application of federated learning to arrhythmia classification is a relatively recent advancement. Bokhari, Sohaib, and Shafi integrate personalized federated models with differential privacy for twelve-lead arrhythmia diagnosis across decentralized institutions, demonstrating a significant performance improvement over both purely centralized and standard federated benchmarks [49]. Elmir et al. utilize a Gramian angular field representation within a federated learning framework specifically to address the heterogeneity of IoT devices that are likely to be present in a real-world deployment, where no two wearables possess identical sensor characteristics or noise profiles [50]. For an edge-first, three-channel wearable as described in this review, the practical security concern is more focused than these architectures fully address: it is less about coordinating training across multiple institutions and more about ensuring that any alert or summary transmitted by the device, whether over BLE, Wi-Fi, or LoRa, is encrypted end-to-end, and that the on-device model itself cannot be easily extracted or altered by an attacker with physical access to the hardware

.XIII. RESEARCH GAPS AND OPEN CHALLENGES

Four consistent gaps have been identified in the literature reviewed above, each of which is directly relevant to a three-channel wearable arrhythmia project.

International Journal of Engineering Research & Technology (IJERT)
ISSN: 2278-0181
Vol. 15 Issue 09 , September - 2026

The most pressing issue is the lack of a large, publicly available dataset that includes simultaneous recordings of three-channel Lead I, II, and V1. All major annotated databases examined in Section VIII either contain only two channels, such as the MIT-BIH database, or provide a complete twelve-lead set, as seen in PTB-XL, Chapman-Shaoxing, CPSC2018, and the Georgia database. Although the twelve-lead sources can be subsetted to obtain the required three leads, none were collected in conditions that mimic the motion artifacts typical of daily use in wearable devices. Therefore, a project focused on this configuration should allocate resources for its own data collection rather than relying solely on public datasets for final validation.

Inconsistencies in reporting across studies hinder effective cross-comparison. As noted in the systematic reviews in Section V, accuracy metrics that combine intra-patient and inter-patient evaluations, or that completely disregard inference time, memory usage, and energy consumption, cannot be accurately compared [27], [28]. Table I in this review had to utilize whatever metrics were reported by each source, highlighting the issue; establishing a shared minimum reporting standard that includes inter-patient accuracy, inference latency, peak memory usage, and energy consumption per inference would significantly enhance the field's capacity to compare models deployed at the edge.

The challenges posed by motion artifacts and the robustness of electrodes during realistic daily activities are more complex for three simultaneous channels than for a single channel. Each additional electrode introduces another potential point of contact failure, and the hardware solutions discussed in Section VI have primarily been validated on single-channel or two-channel configurations; the effectiveness of these techniques when applied to three channels recorded simultaneously during activities such as walking, exercising, or sleeping has not been thoroughly investigated in the literature reviewed here.

Ultimately, the energy-accuracy trade-off for authentic three-channel on-device inference remains largely unquantified. All TinyML results referenced in Section IV, which show accuracy rates ranging from the low to high nineties and inference costs below one millijoule, were derived from single-lead input. Increasing the input channel count to three approximately triples the first-layer computation for most convolutional architectures, yet no study mentioned in this review has documented how this scaling occurs in practice on microcontroller-class hardware for the specific three-lead configuration suggested in Section X.

XIV. CONCLUSION

The convergence of edge computing and IoT has led to an effective approach for wearable arrhythmia monitoring data is collected locally, classified using a compact TinyML model, and only alerts or summarized data are transmitted, while fog and cloud layers are utilized for storage and retraining. Commercial-grade three-channel analog front ends, such as the ADS1293, enable true multi-lead acquisition within the constraints of wearable power and size. The analog evidence reviewed indicates that Lead I, Lead II, Lead VI are the configurations most supported by existing hardware and the largest public datasets available. However, there are still gaps, including the need for matched three-channel public data, consistent reporting across studies, and a comprehensive understanding cf how energy consumption and accuracy vary with channel count on embedded hardware. Addressing these issues presents a practical agenda for those developing a three-channel wearable arrhythmia monitor based on this foundation.

REFERENCES

[1] Global Burden of Cardiovascular Diseases and Risks 2023 Collaborators, “Global, regional, and national burden of cardiovascular diseases and risk factors in 204 countries and territories, 1990–2023,” J. Am. Coll. Cardiol., vol. 86, no. 22, 2025.
[2] F. Sana, E. M. Isselbacher, J. P. Singh, E. K. Heist, B. Pathik, and A. A. Armoundas, “Wearable devices for ambulatory cardiac monitoring: JACC state-of-the-art review,” J. Am. Coll. Cardiol., vol. 75, no. 13, pp. 1582–1592, 2020.
[3] A. Abdelrazik, M. Eldesouky, I. Antoun, E. Y. M. Lau, A. Koya, Z. Vali, S. A. Suleman, J. Donaldson, and G. A. Ng, “Wearable devices for arrhythmia detection: Advancements and clinical implications,” Sensors, vol. 25, no. 9, p. 2848, 2025.
[4] M. Haghi, S. Danyali, S. Ayasseh, J. Wang, R. Aazami, and T. M. Deserno, “Wearable devices in health monitoring from the environmental towards multiple domains: A survey,” Sensors, vol. 21, no. 6, p. 2130, 2021.
[5] 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, no. 5, pp. 485–495, 2013.
[6] M. Hizem, L. Bousbia, Y. Ben Dhiab, M. Ould-Elhassen Aoueileyine, and R. Bouallegue, “Reliable ECG anomaly detection on edge devices for Internet of Medical Things applications,” Sensors, vol. 25, no. 8, p. 2496, 2025.
[7] Y. Abadade, A. Temouden, H. Bamoumen, N. Benamar, Y. Chtouki, and A. S. Hafid, “A comprehensive survey on TinyML,” IEEE Access, vol. 11, pp. 96892–96922, 2023.
[8] M.-H. Fan, M.-H. Guan, Q.-C. Chen, and L.-H. Wang, “Three-lead ECG detection system based on an analog front-end circuit ADS1293,” in Proc. 2017 IEEE Int. Conf. Consumer Electronics – Taiwan (ICCE-TW), Taipei, Taiwan, 2017, pp. 107–108.
[9] I. Azimi, A. Anzanpour, A. M. Rahmani, T. Pahikkala, M. Levorato, P. Liljeberg, and N. Dutt, “HiCH: Hierarchical fog-assisted computing architecture for healthcare IoT,” ACM Trans. Embedded Comput. Syst., vol. 16, no. 5s, pp. 1–20, 2017.
[10] J. A. Rincon, S. Guerra-Ojeda, C. Carrascosa, and V. Julian, “An IoT and fog computing-based monitoring system for cardiovascular patients with automatic ECG classification using deep neural networks,” Sensors, vol. 20, no. 24, p. 7353, 2020.

[11] E. Kim, J. Kim, J. Park, H. Ko, and Y. Kyung, "TinyML-based classification in an ECG monitoring embedded system," Comput. Mater. Contin., vol. 75, no. 1, pp. 1751–1764, 2023.
[12] M. Zambrano-de la Torre et al., "Real-time cardiac arrhythmia classification using TinyML on ultra-low-cost microcontrollers: A feasibility study for resource-constrained environments," Bioengineering, vol. 13, no. 5, p. 532, 2026.
[13] Z. Alimbayeva, C. Alimbayev, K. Ozhikenov, N. Bayanbay, and A. Ozhikenova, "Wearable ECG device and machine learning for heart monitoring," Sensors, vol. 24, no. 13, p. 4201, 2024.
[14] P. Busia, M. A. Scrugli, V. J.-B. Jung, L. Benini, and P. Meloni, "A tiny transformer for low-power arrhythmia classification on microcontrollers," arXiv:2402.10748, 2024.
[15] D. Amiri, J. Takalo-Mattila, L. Bedogni, M. Levorato, and N. Dutt, "SIC-EDGE: Semantic iterative ECG compression for edge-assisted wearable systems," in Proc. 2022 IEEE 23rd Int. Symp. World of Wireless, Mobile and Multimedia Networks (WoWMoM), Belfast, U.K., 2022, pp. 377–385.
[16] J. Pan and W. J. Tompkins, "A real-time QRS detection algorithm," IEEE Trans. Biomed. Eng., vol. BME-32, no. 3, pp. 230–236, 1985.
[17] E. J. da S. Luz, W. R. Schwartz, G. Cámara-Chávez, and D. Menotti, "ECG-based heartbeat classification for arrhythmia detection: A survey," Comput. Methods Programs Biomed., vol. 127, pp. 144–164, 2016.
[18] S. Kiranyaz, T. Ince, and M. Gabbouj, "Real-time patient-specific ECG classification by 1-D convolutional neural networks," IEEE Trans. Biomed. Eng., vol. 63, no. 3, pp. 664–675, 2016.
[19] S. Raj and K. C. Ray, "A personalized arrhythmia monitoring platform," Sci. Rep., vol. 8, no. 1, p. 11395, 2018.
[20] U. R. Acharya, H. Fujita, S. L. Oh, Y. Hagiwara, J. H. Tan, and M. Adam, "Automated detection of arrhythmias using different intervals of tachycardia ECG segments with convolutional neural network," Inf. Sci., vol. 405, pp. 81–90, 2017.
[21] A. Y. Hannun, P. Rajpurkar, M. Haghpanahi, G. H. Tison, C. Bourn, M. P. Turakhia, and A. Y. Ng, "Cardiologist-level arrhythmia detection and classification in ambulatory electrocardiograms using a deep neural network," Nat. Med., vol. 25, no. 1, pp. 65–69, 2019.
[22] N. Alamatsaz, L. Tabatabaei, M. Yazdchi, H. Payan, N. Alamatsaz, and F. Nasimi, "A lightweight hybrid CNN-LSTM explainable model for ECG-based arrhythmia detection," Biomed. Signal Process. Control, vol. 90, p. 105884, 2024.
[23] A. Cinar and S. A. Tuncer, "Classification of normal sinus rhythm, abnormal arrhythmia and congestive heart failure ECG signals using LSTM and hybrid CNN-SVM deep neural networks," Comput. Methods Biomech. Biomed. Eng., vol. 24, pp. 203–214, 2021.
[24] F. Ma, J. Zhang, W. Chen, W. Liang, and W. Yang, "An automatic system for atrial fibrillation by using a CNN-LSTM model," Discrete Dyn. Nat. Soc., vol. 2020, Art. no. 3198783, 2020.
[25] S. K. Pandey, G. Kumar, S. Shukla, A. Kumar, K. U. Singh, and S. Mahato, "Automatic detection of atrial fibrillation from ECG signal using hybrid deep learning techniques," J. Sensors, vol. 2022, Art. no. 6732150, 2022.
[26] Z. Yu, J. Chen, Y. Liu, Y. Chen, T. Wang, R. Nowak, et al., "DDCNN: A deep learning model for AF detection from a single-lead short ECG signal," IEEE J. Biomed. Health Inform., vol. 26, pp. 4987–4995, 2022.
[27] G. A. L. Silva, P. H. L. Silva, G. J. P. Moreira, V. L. S. Freitas, J. C. Gertrudes, and E. J. S. Luz, "A systematic review of ECG arrhythmia classification: Adherence to standards, fair evaluation, and embedded feasibility," arXiv:2503.07276, 2025.
[28] Q. Xiao, K. Lee, S. A. Mokhtar, I. Ismail, A. L. bin Md Pauzi, Q. Zhang, and P. Y. Lim, "Deep learning-based ECG arrhythmia classification: A systematic review," Appl. Sci., vol. 13, no. 8, p. 4964, 2023.
[29] T. Wu, J.-M. Redouté, and M. Yuce, "A wearable, low-power, real-time ECG monitor for smart T-shirt and IoT healthcare applications," in Advances in Body Area Networks I, G. Fortino and Z. Wang, Eds. Cham, Switzerland: Springer, 2019, pp. 165–173.
[30] E. Nemati, M. J. Deen, and T. Mondal, "A wireless wearable ECG sensor for long-term applications," IEEE Commun. Mag., vol. 50, no. 1, pp. 36–43, 2012.
[31] A. B. Jani, R. Bagree, and A. K. Roy, "Design of a low-power, low-cost ECG & EMG sensor for wearable biometric and medical application," in Proc. 2017 IEEE SENSORS, 2017.
[32] A. Khalaf and R. Abdoola, "Wireless body sensor network and ECG Android application for eHealth," in Proc. 2017 4th Int. Conf. Advances in Biomedical Engineering (ICABME), 2017.
[33] S. Krachunov, A. J. Casson, C. Szmigin, and J. Wilson, "Energy efficient heart rate sensing using a painted electrode ECG wearable," in Proc. Global Internet of Things Summit (GIoTS), 2017.
[34] S. Shin et al., "Two electrode based healthcare device for continuously monitoring ECG and BIA signals," in Proc. 2018 IEEE EMBS Int. Conf. Biomedical & Health Informatics (BHI), 2018.
[35] S. Coulter et al., "Low power IoT platform for vital signs monitoring," in Proc. 28th Irish Signals and Systems Conf. (ISSC), 2017.
[36] A. Del Campo, L. Cintioni, S. Spinsante, and E. Gambi, "Analysis and tools for improved management of connectionless and connection-oriented BLE devices coexistence," Sensors, vol. 17, no. 4, p. 792, 2017.
[37] G. B. Moody and R. G. Mark, "The impact of the MIT-BIH arrhythmia database," IEEE Eng. Med. Biol. Mag., vol. 20, no. 3, pp. 45–50, 2001.
[38] P. Wagner, N. Strodthoff, R.-D. Bousseljot, D. Kreiseler, F. I. Lunze, W. Samek, and T. Schaeffter, "PTB-XL, a large publicly available electrocardiography dataset," Sci. Data, vol. 7, no. 1, p. 154, 2020.
[39] J. Zheng, J. Zhang, S. Danioko, H. Yao, H. Guo, and C. Rakovski, "A 12-lead electrocardiogram database for arrhythmia research covering more than 10,000 patients," Sci. Data, vol. 7, no. 1, pp. 1–8, 2020.
[40] F. Liu, C. Liu, L. Zhao, X. Zhang, X. Wu, X. Xu, Y. Liu, C. Ma, S. Wei, and Z. He, "An open access database for evaluating the algorithms of electrocardiogram rhythm and morphology abnormality detection," J. Med. Imaging Health Inform., vol. 8, no. 7, pp. 1368–1373, 2018.
[41] E. A. Perez Alday et al., "Classification of 12-lead ECGs: The PhysioNet/Computing in Cardiology Challenge 2020," Physiol. Meas., vol. 41, no. 12, p. 124003, 2020.
[42] A. L. Goldberger et al., "PhysioBank, PhysioToolkit, and PhysioNet: Components of a new research resource for complex physiologic signals," Circulation, vol. 101, no. 23, pp. e215–e220, 2000.
[43] G. D. Clifford et al., "AF classification from a short single lead ECG recording: The PhysioNet/Computing in Cardiology Challenge 2017," in Proc. Computing in Cardiology (CinC), 2017, pp. 1–4.
[44] A. Dąbrowska, G. Bartkowiak, B. Pękosławski, and Ł. Starzak, "Comprehensive evaluation of a photovoltaic energy harvesting system in smart clothing for mountain rescuers," IET Renew. Power Gener., vol. 14, no. 16, pp. 3200–3208, 2020.
[45] K. Ivanov, "Design, realization and study of thermoelectric watch," in Proc. 21st Int. Symp. Electrical Apparatus & Technologies (SIELA), Bourgas, Bulgaria, 2020, pp. 1–4.
[46] J. Wan et al., "Wearable IoT enabled real-time health monitoring system," EURASIP J. Wireless Commun. Netw., vol. 2018, Art. no. 298, pp. 1–10, 2018.

International Journal of Engineering Research & Technology (IJERT)
ISSN: 2278-0181
Vol. 15 Issue 09 , September - 2026

[47] W. A. N. A. Al-Nbhany, A. T. Zahary, and A. A. Al-Shargabi, "Blockchain-IoT healthcare applications and trends: A review," IEEE Access, vol. 12, pp. 4178–4212, 2024.

[48] M. J. Baucas, P. Spachos, and K. N. Plataniotis, "Federated learning and blockchain-enabled fog-IoT platform for wearables in predictive healthcare," IEEE Trans. Comput. Soc. Syst., vol. 10, no. 4, pp. 1732–1741, 2023.

[49] S. M. Bokhari, S. Sohaib, and M. Shafi, "Fusion of personalized federated learning (PFL) with differential privacy (DP) learning for diagnosis of arrhythmia disease," PLOS ONE, vol. 20, no. 7, e0327108, 2025.

[50] Y. Elmir et al., "Federated learning with Gramian angular fields for privacy-preserving ECG classification on heterogeneous IoT devices," arXiv:2511.03753, 2025.