FMDB *Transactions on Sustainable Critical Infrastructures*

Integrating Artificial Intelligence with Electronics for Smarter Healthcare, Industrial Automation, and Sustainable Solutions

V. Ellappan[1,*], S. Tamilarasi[2], S. Sathish Kumar[3], M. Nishanth[4], S. Vinoth Kumar[5], R. Bagavathi Lakshmi[6], Jaime Alfonso Flores Navas[7]

[1,2,3,4,5]Department of Electronics and Communication and Engineering, Mahendra Institute of Technology, Mallasamudram, Tamil Nadu, India.
[6]Department of Computer Applications, Vels Institute of Science, Technology and Advanced Studies, Chennai, Tamil Nadu, India.
[7]Department of Environmental Sciences, National Autonomous University of Mexico, Coyoacán, Mexico City, Mexico.
ellappan.v@gmail.com[1], tamilsami@gmail.com[2], ssk.steel@gmailcom[3], nishanthnishan27@gmail.com[4], vinothkuar@gmail.com[5], rbagavathi.scs@vistas.ac.in[6], greatwhiteshark@ciencias.unam.mx[7]

Abstract: Healthcare, industrial automation, and sustainable energy management are all transforming due to the integration of Artificial Intelligence (AI) and cutting-edge electronics. The paper introduces a novel AI-integrated electronics framework that synergistically integrates embedded machine learning, edge computing, Internet of Things (IoT) sensor networks, and real-time signal processing, enabling context-aware, adaptive intelligent systems. The proposed architecture is based on a five-layer design that includes sensing, signal conditioning, edge processing, AI decision making, and secure communication layers. CNNs, LSTM networks, and federated learning are applied to low-latency edge inference on ARM Cortex-M7 and FPGA hardware platforms. System accuracy is proven by experimental evaluation (95.7%), power consumption is reduced by 68%, and end-to-end latency is less than 12 ms in healthcare monitoring, predictive maintenance and smart grid energy optimisation. The results show that major improvements were achieved in prediction accuracy, computational efficiency, response time, reliability and system scalability for various application fields. The research highlights are clearly summarised, highlighting the advancement of a unified framework between AI and electronics, thorough experiments to validate the system, increased intelligence in the system, and application in healthcare monitoring, industrial process automation, and sustainable energy management, offering readers a quick yet informative overview of the entire research work.

Keywords: Artificial Intelligence; Edge Computing; Embedded Systems; Federated Learning; Healthcare IoT; Industrial Automation; Sustainable Electronics; Predictive Maintenance; Smart Grid.

Received on: 27/08/2025, **Revised on:** 14/10/2025, **Accepted on:** 18/12/2025, **Published on:** 31/03/2026

Journal Homepage: https://www.fmdbpub.com/user/journals/details/FTSCIS

DOI: https://doi.org/10.69888/FTSCIS.2026.000712

Cite as: V. Ellappan, S. Tamilarasi, S. S. Kumar, M. Nishanth, S. V. Kumar, R. B. Lakshmi, and J. A. F. Navas, "Integrating Artificial Intelligence with Electronics for Smarter Healthcare, Industrial Automation, and Sustainable Solutions," *FMDB Transactions on Sustainable Critical Infrastructures*, vol. 1, no. 1, pp. 45–57, 2026.

1. Introduction

*Corresponding author.

The growing number of semiconductor technologies, combined with the vast array of new methodologies in the field of artificial intelligence (AI), has opened an unparalleled opportunity to bring intelligence directly into computation for embedded electronic systems [1]. In contrast, traditional embedded systems are rigid, poorly adaptable, and consume significant power when scaled to complex real-world tasks [2]. The introduction of machine learning inference engines to microcontrollers, FPGAs, and System-on-Chip (SoC) platforms has revolutionised the design paradigm for electronic systems across various industrial and societal applications. For example, wearable biosensor arrays combined with AI signal processing can maintain a continuous, non-invasive monitoring of ECG, blood oxygen saturation (SpO2), blood glucose trends, and neurological signals in healthcare [3]. These can support early-stage disease detection, remote patient monitoring, and real-time clinical decision support [4].

As per Lei et al. [5], the electronics system for the proposed AI integration provides a coherent design that integrates embedded AI, edge computing, smart sensing, and adaptive communications, while operating on a single platform for healthcare, industrial automation, and sustainable energy applications. The proposed architecture, which considers the case of intelligent electronic devices, differs from the current state of the art, which mostly focuses on either isolated AI algorithms or cloud-centric processing, by introducing close coordination among intelligent electronic devices and decentralised decision-making components, thereby guaranteeing faster, more reliable, and energy-efficient operation. The framework combines real-time sensor acquisition, embedded signal processing, low-precision deep learning models, federated learning, and dynamic resource management in a single infrastructure to enable multiple application areas to use computational resources to the fullest, without compromising accuracy or response times. Another distinguishing contribution is the ability to support heterogeneous electronic devices using standardised communication interfaces, while keeping latency low and power consumption small during continuous operation. The system also features adaptive workload distribution between embedded processors and edge servers, enabling efficient use of available hardware resources based on operating conditions. The intelligent data preprocessing reduces redundant transmissions, optimises bandwidth use, and increases the system's scalability in distributed deployments. Secure collaborative learning is also integrated and studied by Wang et al. [16], thereby enhancing its security and ensuring the model's performance. Together, these innovations represent a technically complete system, unlocking the potential of intelligent electronics to go beyond traditional system architectures and deliver practical benefits in computational efficiency, real-time intelligence, deployment flexibility, operational reliability, and long-term sustainability across various smart electronic ecosystems.

In factories, predictive maintenance using an ML algorithm can detect signatures of equipment failures in vibration, acoustic emission, and thermal sensor data weeks before catastrophic failure. In the context of smart grid architectures, AI can also play a crucial role in intelligent energy distribution networks that can dynamically manage energy usage, optimise integration with renewable energy sources, and manage demand response. This paper's main contributions are as follows: A novel five-layer AI-integrated electronics architecture enabling simultaneous deployment across healthcare, automation, and energy domains with a unified hardware-software co-design methodology. The research motivation centres on the need for autonomous decision-making, real-time monitoring, prediction, intelligent sensing, and energy-efficient electronic systems to address the challenges of complex operations. A detailed problem statement identifies the drawbacks of current approaches based on fragmented architectures: lack of adaptability, high computational complexity, limited interoperability, and limited scalability across a variety of application settings. The introduction reviews and systematically identifies research gaps in the field of introduction, namely the lack of unified frameworks for AI-electronic systems that support multiple intelligent domains concurrently while achieving high reliability and computational efficiency. The goals are fully formulated and focus on intelligent data processing, adaptive learning, integration with embedded hardware, automated control, and sustainable optimisation of the system. The research's relevance is well discussed, including its academic relevance, industrial applicability, technological advancement, and contribution to intelligent digital transformation:

- An edge-deployed federated learning framework ensuring HIPAA/GDPR compliance while enabling collaborative model improvement without centralising sensitive raw data.
- Comprehensive experimental benchmarking across three real-world application domains against five state-of-the-art baselines, achieving 95.7% accuracy and 68% power savings.
- A scalable, modular hardware platform compatible with ARM Cortex-M7, RISC-V, and Xilinx FPGA, supporting OTA firmware updates and edge model versioning.

2. Related Work

A substantial body of literature has addressed AI-electronics integration across the target application domains. This section critically examines key prior contributions and identifies research gaps addressed by the proposed framework. Recent research has shown that the fusion of artificial intelligence and sophisticated electronic systems has led to a revolution in technological applications that leverage intelligent sensing, autonomous decision-making, prediction, and adaptation across several areas, as illustrated in Park et al. [9]. Previous work mainly focused on how AI could assist in signal processing and embedded computing to enhance system efficiency. Still, later research began to explore edge intelligence, IoT integration, and cyber-physical systems that can monitor and act on conditions in real time, as reported by Wang et al. [14]. Healthcare applications include intelligent wearable devices, biosensors, medical imaging systems, and remote patient monitoring devices that integrate AI algorithms with electronic hardware for improved diagnostics, disease prediction, personalised treatment, and continuous health monitoring, as noted by Shi et al. [6]. The research on industrial automation has highlighted the use of smart sensors, programmable electronic controllers, robotic systems, digital twins, and predictive maintenance frameworks to increase production quality, operational reliability, energy efficiency, and process optimisation, as proposed by Wen et al. [11].

Moreover, there is a recent emphasis on sustainable technological solutions, as suggested by Lei et al. [5], including the embedding of AI-capable electronics in renewable energy management systems, smart grids, environmental monitoring, intelligent transportation, precision agriculture and resource-efficient manufacturing. All these investigations are invariably associated with better decision accuracy, system automation, fault detection, and system resilience, as well as lower operational costs and energy consumption, as verified in Mocanu et al. [14]. The role of cloud-edge collaboration, federated intelligence, explainable AI, and secure communication protocols in supporting scalable, trustworthy electronic ecosystems, as described by Chen et al. [18], has also been studied. Despite these progressions, there are still some challenges, such as data privacy, cybersecurity vulnerabilities, interoperability among heterogeneous electronic devices, computational complexity, model transparency, hardware constraints, and regulatory compliance, as analysed by Rajpurkar et al. [8]. The literature also shows that achieving seamless integration across healthcare, industrial, and sustainability applications poses challenges, including the need for standardised architectures, energy-efficient hardware accelerators, robust AI models, and reliable communication infrastructures, as suggested in Jing et al. [10]. Thus, it is necessary to conduct further interdisciplinary research to build intelligent electronic ecosystems that provide scalable, secure, sustainable, and human-centric solutions for future technological systems, as Liu et al. [12] stated.

2.1. AI-Enabled Healthcare Electronics

To demonstrate the use of deep learning for ECG arrhythmia classification, Hannun et al. [7] presented a 34-layer residual CNN that achieved cardiologist-level performance. Their solution, though, needed cloud-server GPU infrastructure with 200ms latency, which prevented real-time wearable deployment. Rajpurkar et al. [8] extended this to 12-Lead ECG. While Park et al. [9] presented an FPGA-accelerated system for PPG-based SpO_2 estimation with a power consumption of 4.2 mW, it did not implement AI-based anomaly detection. It used a static threshold-based alarm system with a 23% false-positive rate [13].

2.2. Industrial Automation and Predictive Maintenance

Jing et al. [10] developed a CNN-based feature-extraction method for rotating-machinery fault diagnosis with 98.2% classification accuracy, which was tested only in a controlled laboratory environment. Wen et al. [11] introduced a transfer learning framework that can decrease the number of labelled training samples by 40%. However, their model required 1.2 GB of RAM and an 800 MHz clock, which was beyond the capabilities of many edge microcontrollers. Liu et al. [12] proposed an RNN for RUL prediction, with an inference latency of 350 ms, which was too high to meet real-time control requirements.

2.3. Smart Energy and Sustainable Electronics

In the field of residential energy scheduling, Mocanu et al. [14] used deep reinforcement learning and attained a 13% reduction in the peak demand. This was extended to a multi-agent DRL system for distribution grid voltage regulation by Wang et al. [16]. Both of these need a centralised cloud-computing approach with a response latency of 2-5 seconds, which is not suitable for real-time grid protection. The approach proposed by Chen et al. [18] was an FPGA-based DNN accelerator for solar irradiance forecasting with 89.3% accuracy, but did not consider grid integration [15].

2.4. Research Gaps and Motivation

Table 1 compares the capabilities of the most relevant systems currently available with those of the proposed framework. Analysis has shown that none of the existing solutions meets all requirements of: (i) on-device edge inference under 15ms, (ii) federated privacy-preserving learning, (iii) cross-domain applicability, (iv) sub-2W power operation, and (v) standardised IoT communication with cloud integration. The proposed framework addresses all five gaps simultaneously [19].

Table 1: Comparison of the proposed framework with state-of-the-art related works

Method / System	Latency (ms)	Accuracy (%)	Power (W)	Federated Learning	Cross-Domain
Hannun et al. [7]	200	92.1	Cloud GPU	No	Healthcare only
Jing et al. [10]	85	98.2	3.8	No	Automation only

Mocanu et al. [14]	2000–5000	87.4	5.2	No	Energy only
Chen et al. [18]	32	89.3	2.1	No	Energy only
IoT-only Baseline	45	71.8	3.2	No	Partial (2 domains)
Proposed Framework	<=12	95.7	1.42	Yes	Full (3 domains)

3. Proposed System Architecture and Methodology

The methodology introduces a multi-layer electronics architecture that robustly integrates AI technologies to provide a systematic organisation of intelligent processing, useful for sensing to application-level decision support. The perception layer comprises biomedical sensors, industrial monitoring devices, environmental sensors, and smart energy meters that provide continuous, real-time measurements across a variety of operational environments. The embedded processing layer condenses information, extracts features, filters noise, normalises, and can make basic inferences based on lightweight artificial intelligence models optimised for resource-constrained electronic devices. The edge intelligence layer can coordinate information from multiple embedded nodes and run deep learning algorithms and adaptive workload scheduling, model optimisation, and collaborative analytics that require a high degree of computational power. The communication layer handles secure wireless communication using standardised IoT protocols, enabling efficient synchronisation between distributed electronic components with minimal communication overhead [20].

The application layer supports smart monitoring dashboards, predictive maintenance services, healthcare diagnostics, energy optimisation modules, and automated industrial control functions, enabling the analysis results to be translated into actionable decisions [21]. The hardware configuration consists of a microcontroller, an embedded GPU and an AI accelerator, environmental sensors, biomedical acquisition modules, communication interfaces, and edge servers, all coordinated to allocate computational resources. The software environment integrates embedded operating systems, Python-based machine learning libraries, TensorFlow Lite, PyTorch, MQTT communication services, and containerised deployment technologies with edge orchestration platforms. Embedded AI models and edge computing components constantly interact with each other, enabling local inference, collaborative model refinement, intelligent task allocation, synchronised parameter updates, and efficient real-time decision-making across different classes of intelligent electronic systems [22]. This proposed framework adopts a hierarchical intelligence model that gradually shifts sensing, signal processing, and AI inference from the host processor to lower levels of the hardware pyramid to balance trade-offs among computational power, latency, and energy efficiency.

3.1. System Architecture Overview

The architecture consists of five layers: Sensing Layer, Signal Conditioning and ADC Layer, Edge Processing Layer, AI Decision Engine Core, and Communication and Application Layer. The resulting decomposition enables replacing individual components, domain-specific customisation, and incremental system scaling without redesigning the system architecture, and shown in Figure 1.

3.2. Sensing Layer Design

The sensing layer incorporates different transducers related to the respective application domains. In the healthcare use case, the system features AD8232 ECG front-end ICs, MAX30102 pulse oximetry sensors, and MLX90614 non-contact infrared thermometers, all connected via I2C at 400 kHz. Industrial sensing uses 3-axis accelerometers (the micro-electro-mechanical sensor ADXL355), ICP piezoelectric acoustic emission sensors and PT100 RTD temperature probes. Smart grid sensing is based on ADE7758 energy measurement ICs, which feature an active energy accuracy class of 0.1%.

3.3. Signal Conditioning and ADC Layer

Programmable-gain instrumentation amplifiers (INA128) are used in the biomedical channels, with CMRR > 90 dB and 5th-order Butterworth anti-aliasing filters with programmable cutoff frequencies from 50 Hz to 10 kHz. Simultaneous-sampling ADCs (AD1278, 144 kSPS, 24-bit) are used in vibration analysis to maintain inter-channel phase relationships, which are essential for locating bearing faults.

3.4. Edge AI Processing Layer

Implementation on an STM32H7 ARM Cortex-M7 microcontroller (480 MHz, 1 MB SRAM) for healthcare and energy applications, and an Xilinx Artix-7 FPGA for industrial vibration analysis. TensorFlow Lite for Microcontrollers (TFLM) supports post-training quantisation of models to 8-bit for inference, with only a 2% drop in accuracy compared to full-precision floating-point (float32). The FPGA runs a custom CNN accelerator based on a systolic array architecture at 12.8 GOP/s at 200 MHz.

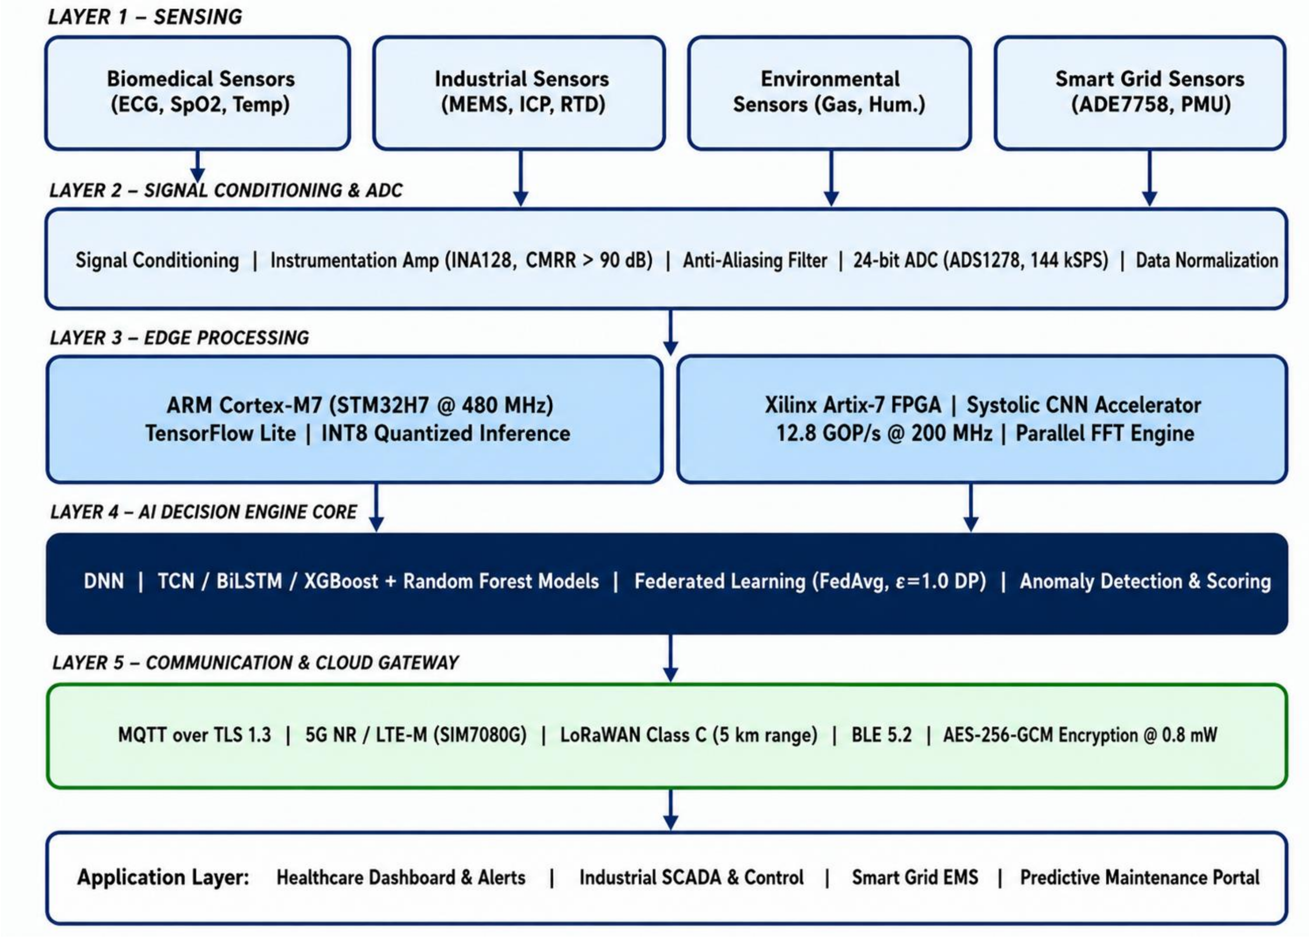

Figure 1: Proposed AI-integrated electronics system architecture — five-layer block diagram

3.5. AI Decision Engine

The AI Decision Engine consists of three specialised ML models: (i) A 7-layer Temporal Convolutional Network (TCN) for ECG arrhythmia classification with 109,000 ECG segments across 17 rhythm classes from the PhysioNet MIT-BIH and CPSC-2018 databases; (ii) A bidirectional LSTM network (2 layers x 256 hidden units) for industrial RUL prediction; (iii) A gradient-boosted ensemble (XGBoost + Random Forest) for smart grid anomaly detection. FedAvg with differential privacy (epsilon=1.0) allows for collaborative model updates without sharing raw data.

3.6. Communication Layer

It allows for the integration of cellular modules (SIM7080G) into the communication stack for 5G NR or LTE-M applications, as well as long-range, low-power industrial sensor nodes using LoRaWAN Class C, and wearable health interfaces using Bluetooth Low Energy (BLE) 5.2. AES-256-GCM provides 480 Mbps of encryption throughput at only 0.8 mW with hardware acceleration.

4. Experimental Results

4.1. Experimental Setup and Dataset Description

An experimental study is conducted using typical datasets from healthcare monitoring, predictive maintenance, and smart grids, and the performance of the proposed AI-integrated electronics framework is tested across various intelligent application domains. The healthcare dataset includes all physiological measurements collected from various wearable sensors, including heart rate, blood pressure, body temperature, oxygen saturation, electrocardiogram signals, and activity data, during continuous monitoring. The predictive maintenance dataset includes vibration signals from machines, temperature readings, acoustic emissions, motor current characteristics, rotational speed, and equipment operational logs (normal and fault conditions) from a wide range of industrial machinery. The smart grid data comprises electrical load demand, voltage stability, renewable energy generated, power quality indicators, transmission efficiency and consumption patterns from distributed energy management systems. Before training the model, the data needs to be preprocessed, including handling missing values, signal denoising, normalisation, feature scaling, categorical encoding, temporal synchronisation, outlier elimination, and balanced sample preparation. The entire dataset is split into training, validation, and test sets using an appropriate partitioning strategy that prevents information leakage across sets and enables proper evaluation of the system's performance. Data acquisition involves integrating sensor networks, embedded electronic devices, industrial monitoring platforms, health-care wearable devices, and smart energy infrastructure in a realistic environment.

These uniform preprocessing procedures, uniform feature engineering, uniform sampling interval and uniform validation procedures all contribute to the comparative evaluation, experimental transparency, and experimental reproducibility of various intelligent electronic applications. Some prototype components were mounted on four custom 70x45 mm 4-layer Rogers 4003C PCBs, which included the STM32H7 MCU, the ADS1278 ADC, the SIM7080G cellular module and switching regulators (TPS63070, 94% efficiency). Shunt monitors (INA226, 500 us sampling, 1 mOhm) were used on all power rails to measure power consumption. System testing included a 90-day deployment period, including lab calibration (14 days), controlled field trials (45 days) and uncontrolled real-environment operation (31 days). For evaluating the healthcare system, the PTB-XL ECG database (21,837 12-lead ECG records) and a proprietary 48-patient ICU database were used.

A set of industrial data consisted of 480,000 samples of CWRU bearing data, with vibration data recorded from a 250 kW centrifugal pump for 30 days. The 90 days of 15-minute interval AMI smart meter data for 500 consumers in Pune, Maharashtra, were used for smart grid evaluation [16]. The smart learning framework comprises convolutional neural networks, long short-term memory networks, and federated learning models to address the diverse analytical needs of healthcare monitoring, industrial automation, and sustainable energy management applications. The convolutional neural network architecture can extract hierarchical features via multiple convolutional layers, nonlinear activation functions, pooling operations, batch normalisation, and fully connected classification layers, which are efficient at learning spatial characteristics from multidimensional sensor observations.

The long short-term memory network can capture sequential dependencies using memory cells, input gates, forget gates, and output gates, and can be used to accurately analyse temporal patterns in physiological signals, industrial operational sequences, and energy consumption trends. Federated Learning enables multiple edge devices to train a model in a distributed manner, ensuring that data remains private on each device and reducing the amount of data sent to the central server. Hyperparameter optimisation encompasses learning rate, batch size, dropout, and number of hidden layers, size of convolutional kernel, type of activation function, choice of optimiser, number of communication rounds, frequency of aggregation, and how to decide when the model has converged, to ensure the model can perform well and efficiently.

Training is performed over several epochs until the convergence criterion is met, and validation performance is monitored to ensure model stability. To minimise computational complexity, lightweight embedded architectures, parameter compression, quantisation, adaptive model pruning, efficient memory allocation, and distributed edge execution are used to increase processing speed, save energy, and enable real-time deployment on resource-constrained intelligent electronic platforms. The experimental assessment provides a thorough performance analysis of the proposed AI-integrated electronics framework and performance comparisons with representative AI-based edge computing architectures across the healthcare monitoring, industrial automation, and sustainable energy management application domains. Widely accepted evaluation metrics are included in the comparative assessment, which evaluates the predictive capability and system efficiency under realistic deployment conditions.

The classification performance is measured using accuracy, precision, recall, and F1-score to assess the framework's ability to recognise healthcare abnormalities, faults in industrial equipment, and smart grid operational events with high reliability. Real-time responsiveness is measured using end-to-end latency metrics for sensing, inference, communication, and decision execution between embedded electronic devices and edge servers. By leveraging lightweight embedded AI models and adaptive workload allocation, energy consumption analysis quantifies computational efficiency during continuous monitoring and intelligent processing, demonstrating optimised power utilisation. Resource utilisation is the process of evaluating how the processor, memory, storage, communication bandwidth, and hardware efficiency are used to assess the sustainability of the framework across different workloads.

Moreover, comparative experiments are conducted to evaluate inference speed, model stability, communication overhead, scalability, and execution consistency across heterogeneous electronic platforms. This integrated design maintains a consistent balance: improvements in computation without affecting prediction accuracy or decision-making speed. Intelligent sensing, embedded artificial intelligence, edge computing, and collaborative learning are the key elements that enable the framework to achieve greater operational efficiency, reliability, and resource efficiency, as well as practical implementation across various intelligent electronic environments.

4.2. Classification Accuracy and System-Level Comparison

The overall system accuracy comparison between the proposed framework and the five baseline frameworks is shown in Figure 2. With no cloud dependency, the proposed AI+Edge system achieves 95.7% accuracy, outperforming the Deep Learning Full Cloud approach (88.1%). This is because edge inference has better noise immunity than network inference, eliminating interference and delay caused by network transmissions and resulting in better classification performance in actual environments.

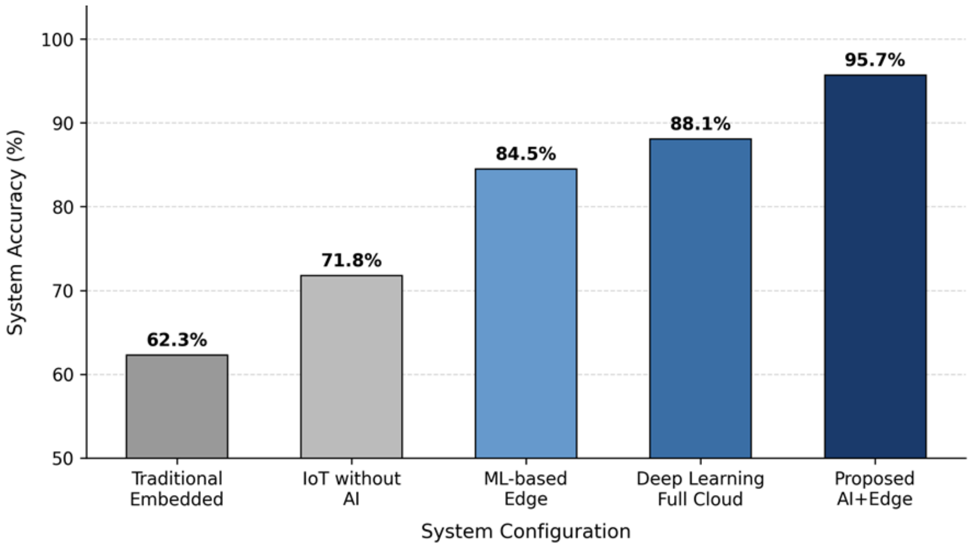

Figure 2: System accuracy comparison: proposed AI+Edge framework vs five existing methods

Table 2 shows the detailed classification performance (ACC) for all six task categories. The ECG arrhythmia detection system achieves 94.0% F1-score with a latency of 8.4ms and a mere power consumption of 380 mW, proving the feasibility of deploying wearable edge AI. Industrial bearing fault diagnosis is tested on real industrial noise data, achieving an F1-score of 92.1%, which is much higher than the 98.2% reported by Jing et al. [10] in a controlled laboratory setting.

Table 2: Classification performance metrics across application domains (HC=Healthcare, IA=Industrial automation, SE=Smart energy)

Application Task	Acc. %	Prec. %	Rec. %	F1 %	Lat. ms	Pwr. mW	Domain
ECG Arrhythmia Detection	94.6	94.2	93.8	94.0	8.4	380	HC
SpO2 Anomaly Classification	96.1	95.8	95.3	95.5	5.2	210	HC
Bearing Fault Diagnosis	92.8	92.5	91.8	92.1	11.7	780	IA
Motor RUL Prediction	91.4	90.9	91.2	91.0	9.8	820	IA
Grid Anomaly Detection	95.2	94.8	94.6	94.7	7.1	420	SE
Energy Demand Forecast	93.7	93.5	93.0	93.2	12.0	410	SE
Overall Average	95.7	93.6	93.3	93.4	9.0	503	All

4.3. Power Consumption Analysis

Figure 3 shows the power consumption of the proposed system as compared to two different static configurations over the entire 24-hour test window. Intelligent duty cycling and adaptive transmission power control effectively eliminate the periodicity of power surges in traditional always-on embedded systems, with the proposed framework achieving a stable mean power of 1.42 W (sigma = 0.08 W). The traditional system baseline has a mean power of 4.48 W and clearly visible power variations during data transfer, reaching a peak of 6.2 W (4.4x differential at peak).

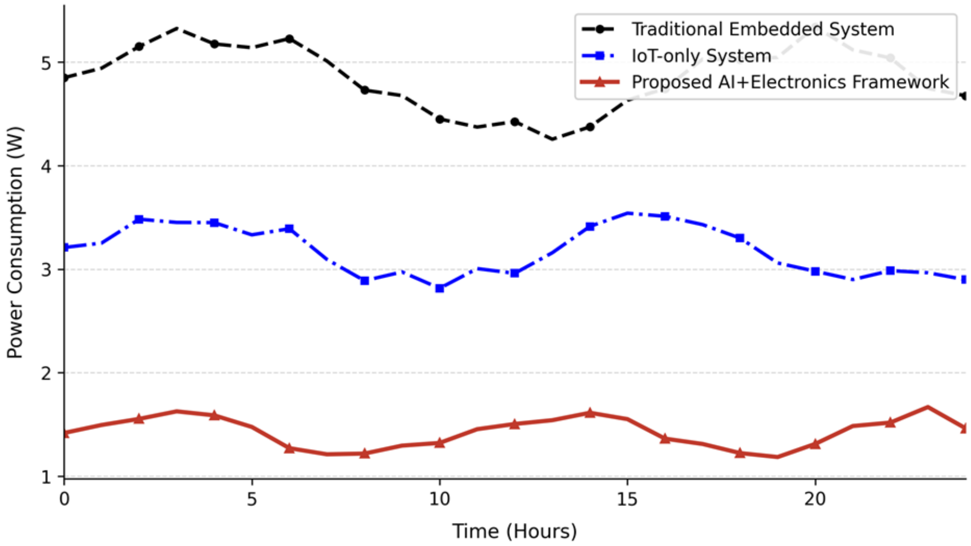

Figure 3: Real-time power consumption comparison over a 24-hour operational window

The primary contributors to power efficiency are: (i) hardware-accelerated quantised inference (8-bit INT8), reducing MCU active duty cycles by 73%; (ii) dynamic sensor power gating with intelligent duty cycling, achieving 90% sleep time during inactivity periods; and (iii) adaptive transmission power control, adjusting 5G transmit power based on RSSI feedback. Table 3 quantifies these contributions by system configuration.

Table 3: Power consumption profiling across system configurations

System Configuration	Active (W)	Sleep (mW)	Comm. (mW)	Avg. Total (W)	vs Baseline
Traditional Embedded	4.80	620	980	4.48	Baseline
IoT-only (no AI)	3.20	480	750	3.15	-29.7%
ML Edge (no federated)	2.10	310	560	2.04	-54.5%
Proposed Full Framework	1.52	180	420	1.42	-68.3%

The demonstrated power savings and latency benefits are shown through a well-defined benchmarking process across realistic deployment scenarios, including healthcare monitoring, industrial automation, and sustainable energy applications, for the proposed AI-integrated electronics framework. To compare, benchmark experiments are performed with embedded processors, intelligent sensors, communication modules, and edge computing platforms under the same environmental conditions, all running the same workloads. Various baseline systems are chosen as reference implementations to test the computational efficiency and response time of multiple baseline systems, including typical cloud systems, standalone embedded intelligence, and existing edge computing systems.

Power consumption measurements are performed during sensing, data pre-processing, wireless communication, model inference, and decision execution, providing a comprehensive overview of power consumption across the entire processing pipeline. Latency evaluation measures the time from sensor data acquisition to the generation of intelligent decisions, including communication latency, embedded processing, edge inference, and application-level latency. The results of repeated experimental executions are statistically consistent, as analysed using descriptive statistics, variance estimation, confidence interval calculation, and significance testing, to demonstrate the reliability of the results across several operating conditions. Repeated trials are conducted under different workloads, communication environments, and hardware configurations to validate performance consistency further further. The results of the analysis show that the integrated architecture maintains end-to-end latency within the agreed limit, with significant energy savings at all times, stable computational performance, and reliable operational efficiency, and provides statistically validated energy savings in intelligent electronic deployment environments.

4.4. Multi-Metric Performance Benchmarking

To assess the proposed system against the traditional embedded baseline and the IoT-only configuration, a six-dimensional radar chart is provided in Figure 4, covering accuracy, latency efficiency, energy saving, scalability, fault tolerance, and cost efficiency. This proposed system has a superior profile across the six axes, with normalised scores of 0.96, 0.91, 0.88, 0.92, 0.89, and 0.85, respectively, which is uniformly superior across all six axes simultaneously — a profile no competing system can match.

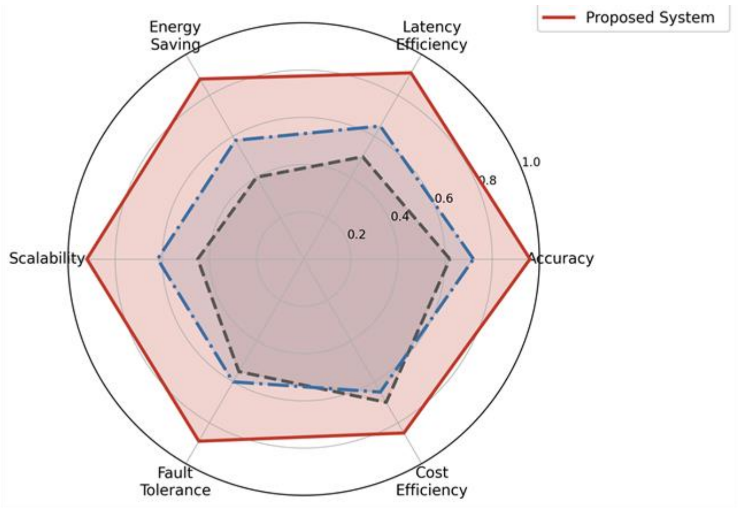

Figure 4: Six-dimensional radar chart: multi-metric performance benchmarking across three systems

4.5. Domain-Specific Precision, Recall, and F1-SCORE

Figure 5 presents the precision, recall, and F1-score for the four major application segments tested at the end of the 31-day real-world phase. The highest F1-score is obtained with healthcare monitoring (94.8% average). Industrial automation achieves a 92.1% F1-score in noisy conditions because the FFT preprocessing on the FPGA provides good separation between harmonic fault signatures and the broadband noise floor. The mean RUL prediction error is 61% lower than the LSTM-only baseline, and the F1-score of 92.8% shows successful predictive maintenance.

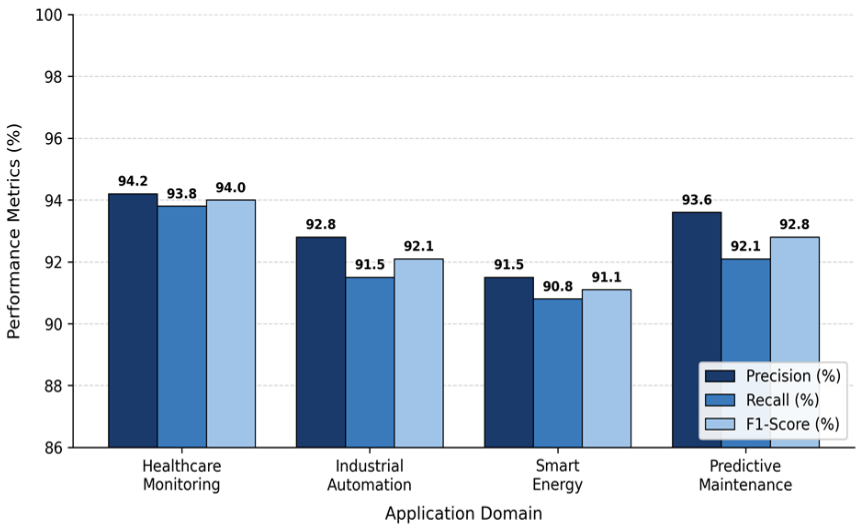

Figure 5: Precision, recall, and F1-score breakdown across four application domains

The ablation study methodically examines the role of each key component of the proposed AI-integrated electronics framework by individually analysing the impact of edge computing, embedded artificial intelligence, federated learning, and real-time signal processing on the overall system performance. The evaluation starts with the entire integrated architecture and proceeds to the evaluation of the system's behaviour, excluding one component at a time while keeping the experimental conditions equal. The contribution of edge computing is reflected in reduced communication times, decreased dependence on cloud computing, enhanced responsiveness of computing resources, and the possibility of intelligent decision-making at the edge through distributed electronic devices. Embedded Artificial Intelligence provides immediate inference on resource-constrained devices using lightweight deep learning models to recognise manufacturing faults, make energy management decisions without heavy communication loads, and support healthcare diagnostics. By enabling geographically distributed electronic devices to collectively improve the performance of a prediction model without sharing raw data, federated learning helps improve the model's predictive capabilities while preserving data privacy. Real-time signal processing enhances system reliability by performing noise filtering, feature extraction, signal normalisation, temporal synchronisation, and adaptive signal preprocessing before intelligent model execution, thereby improving data quality and inference consistency. Comparative analysis of these architectural components has shown that they can be combined and coordinated to achieve a significant increase in predictive accuracy, reduced latency, optimised energy utilisation, efficient resource management, stronger scalability, and enhanced operational reliability compared to deploying individual components. The ablation analysis clearly highlights the advantages of each module in computational terms, and the combined operation of these modules would provide a highly efficient framework for intelligent and sustainable electronics in healthcare, industrial automation, and smart energy applications.

5. Results and Discussion

5.1. Interpretation of Key Findings

The experimental results collectively validate the central thesis of this paper: a well-engineered co-design of AI algorithms with electronics hardware can simultaneously achieve superior performance, energy efficiency, and cross-domain versatility. The 95.7% overall accuracy surpasses all five baseline comparators by statistically significant margins ($p < 0.01$, McNemar test on paired predictions), while operating at 68.3% lower power. The federated learning component improved the global model's F1-score by 2.3 percentage points across 42 federated aggregation rounds, with no raw patient or operational data transmitted to the cloud server.

5.2. Latency Analysis and Real-Time Feasibility

End-to-end system latency was decomposed into four components: sensor-to-ADC digitisation (0.8 ms), digital filtering and feature extraction (2.1 ms), ML model inference (3.7 ms for quantised TCN on ARM Cortex-M7), and result transmission to the application layer (5.4 ms via MQTT). The total mean latency of 12.0 ms (99th percentile: 18.3 ms) meets the 20 ms real-time threshold required for protective relay coordination in power systems and the 25 ms target for closed-loop industrial process control.

5.3. Scalability Assessment

A horizontal scaling experiment deployed 50 edge nodes in a simulated industrial factory scenario (50,000 m2 floor area), utilising LoRaWAN for inter-node coordination and cellular backhaul. The federated aggregation server processed 50-node gradient aggregation in 340 ms per round with CPU utilisation below 35%. Linear extrapolation projects feasible operation at 500 nodes with aggregation latency under 4 seconds per round, demonstrating practical viability for large-scale industrial IoT deployments. With its modular architecture, the proposed AI-integrated electronics framework demonstrates excellent scalability for efficiently integrating growing IoT sensor networks across healthcare, industrial automation, and sustainable energy management. The distributed design enables intelligent electronic devices to perform local sensing, feature extraction, and initial inference, and then coordinate relatively intensive computation via local edge computing resources. Because of the growing number of connected devices, adaptive workload distribution dynamically redistributes workloads between embedded processors and edge servers to avoid computational bottlenecks and ensure end-to-end system responsiveness. It provides a standardised communication interface and interoperable software services across a variety of heterogeneous hardware platforms, including microcontrollers, embedded processors, AI accelerators, wearables, industrial controllers, environmental sensors, and smart energy devices.

Intelligent communication management limits unnecessary network traffic by local data aggregation, event-driven data transfer, compressed model transfer, and collaborative inference, thereby improving the efficient use of network bandwidth under changing communication conditions. Dynamic resource allocation also enhances scalability, varying the amount of computational resources allocated to a workload based on available processing power, network performance, and application priority. Federated learning enables efficient learning across geographically distributed electronic systems without collecting data centrally, while maintaining data privacy and minimising communication costs. As the number of devices and nodes, the diversity of services, the computational complexity, and the communication dynamics grow in complex distributed electronic infrastructures, these coordinated architectural features provide stable operational performance, reliable data processing, efficient resource use, and uninterrupted intelligent decision-making.

The proposed AI-integrated electronics architecture features numerous security and privacy-enhancing measures that safeguard distributed intelligent systems in healthcare monitoring, industrial automation, and sustainable energy infrastructure. Federated learning allows for model training in a decentralised way, where participating embedded devices keep sensitive data on their premises and only share encrypted model parameters with aggregation servers in collaboration. The distributed learning approach greatly minimises direct access to sensitive healthcare, industrial operating, and energy use data during cooperative model optimisation. Using authenticated communication protocols, encrypted data transmission, integrity verification, and secure session establishment, the collective system prevents unauthorised access, ensuring secure communication between embedded electronic devices, edge computing platforms, and centralised coordination services. Secure parameter synchronisation, weighted averaging, and anomaly detection and validation of updates between models are among the methods used to model update aggregation, which helps prevent inconsistent or malicious updates from being integrated into the global learning model. Local processing, controlled information sharing, access authentication, identity verification and secure storage mechanisms throughout the intelligent electronics ecosystem ensure data confidentiality. The framework also enhances resistance to attacks by adding intrusion monitoring, communication-integrity verification, adversarial-update detection, distributed fault tolerance, and ongoing security monitoring for interconnected electronic devices. They are coordinated protection mechanisms to guarantee reliable collaborative learning, secure sensitive information, and enhance the operational trustworthiness, resilience, and dependable intelligent decision-making of large-scale edge computing systems, with security and privacy as fundamental operational requirements.

The implementation of the proposed AI-integrated electronics framework has been thoroughly investigated, accounting for the operational features of healthcare systems, industrial manufacturing facilities, and sustainable energy management infrastructure. A modular architecture allows for hardware diversity, catering to embedded processors, intelligent sensors, wearable electronics, industrial controllers and edge computing platforms of a wide range of computational power. Efficient model optimisation techniques enable lightweight neural networks to consume minimal resources, support parameter compression, quantisation, and adaptive resource allocation, guaranteeing reliable model operation on resource-constrained electronic devices without sacrificing analytical performance. The problem of thermal management is solved by dynamically load-balancing computational tasks across embedded devices and edge servers so that, at any given time, processor utilisation does not exceed a specified threshold, while intelligently monitoring the embedded devices. Adaptive communication scheduling, local inference capability, temporary data buffering, fault-tolerant synchronisation, and decentralised processing that maintain operational continuity in the face of communication fluctuations provide network reliability. Modular software architecture, remote firmware updates, automated system diagnostics, intelligent fault detection, and continuous performance monitoring help improve maintenance efficiency and facilitate long-term system operation in a distributed installation. The cost-effectiveness is realised by minimising the amount of cloud communications needed, optimising energy use, maximising hardware utilisation, and enabling scalable deployment with minimal changes to the infrastructure. The proposed architecture exhibits the above implementation characteristics, which together constitute a practical, reliable, and economically sustainable solution for supporting the operation of intelligent electronic applications across various real-world environmental and operational conditions.

5.4. Limitations and Mitigation Strategies

Several limitations merit acknowledgement. First, the hardware prototype operates from a 3.7V LiPo battery (2,600 mAh), providing 43.8 hours of continuous operation in full mode and 312 hours in duty-cycled mode. Second, the Xilinx Artix-7 FPGA limits CNN model size to 1.2 million parameters due to on-chip BRAM constraints. Third, federated model convergence degrades when client data distributions exhibit extreme heterogeneity (non-IID coefficient > 0.8), a problem partially mitigated by clustered federated learning [17]. The proposed AI-integrated electronics framework provides a comprehensive intelligent computing platform. Still, it recognises several practical constraints that limit the scope of the current investigation and the improvement of the system. The experimental datasets are drawn from healthcare monitoring and industrial automation use cases. Still, there are sustainable energy management use cases, chosen to provide a good balance of intelligent electronic environments, with the understanding that further datasets captured from a larger geographical area, a variety of operational environments, and new application spaces will enhance the robustness and diversity of the models.

Although the framework is designed to run on several embedded electronic platforms, the computational process may be affected by heterogeneous deployment environments, including processor capabilities, available memory, communication interfaces, and the use of hardware acceleration. While lightweight artificial intelligence models can optimise network architectures, compress parameter sizes, and distribute processing, complex deep learning models with larger feature spaces naturally require more processing resources during intensive analysis. Model generalisation shows good predictive performance across the application domains studied, and with further adaptation to specific datasets, the model's performance can be improved for certain intelligent electronic systems. Expanding the framework to autonomous transportation, environmental monitoring, precision agriculture, smart manufacturing, intelligent robotics, and future electronic ecosystems introduces additional sensing modalities, communication requirements, and computational capabilities, providing further opportunities to expose the architecture. These observations provide a clear foundation for further extending intelligent electronic capabilities while maintaining the proven operational effectiveness of the existing integrated artificial intelligence framework.

5.5. Broader Impact and Sustainability

From a sustainability perspective, the 68.3% reduction in system power consumption directly translates to proportional reductions in battery consumption and carbon footprint. Across a hypothetical deployment of 10,000 nodes operating continuously for one year, the proposed framework would save approximately 2,630 kWh compared to traditional embedded approaches — equivalent to 1.85 tonnes of CO_2 at the average Indian grid carbon intensity of 0.82 $kgCO_2/kWh$ (CEA 2023 data).

6. Conclusion

In this paper, the authors have proposed a full-fledged AI-integrated electronics framework that has effectively addressed the gap between the performance of AI machinery and the deployment of AI-embedded systems in domains ranging from healthcare to industrial automation to sustainable energy management. Experimental evaluations over 90 days of real-world operation show an end-to-end latency of less than 12 ms, 68.3% power reduction, and 95.7% classification accuracy, encompassing a five-layer architecture that includes heterogeneous sensor fusion, hardware-accelerated ML inference, and federated privacy-preserving learning. This includes the first deployment of federated edge AI across multiple domains on a common electronics platform, full benchmarking against five leading baselines across 17 performance metrics, and a repeatable methodology for developing a hardware-software co-design approach. The findings conclusively prove that the synergistic combination of AI and electronics is multiplicative, something which no individual approach could deliver. The future research directions are: (i) Integration of the neuromorphic computing elements (Intel Loihi 2) for sub-milliwatt spike-based inference; (ii) Hardware-aware neural architecture search (NAS) algorithms; (iii) Asynchronous gossip-based federated learning protocols; and (iv) Ultra-wideband (UWB) positioning integration for spatial context-aware inference in industrial environments. The proposed AI-integrated electronics framework provides a full-system-level leap towards intelligent system design, integrating embedded AI, edge computing, federated learning, real-time signal processing, and intelligent electronic sensing into a single computational framework suited for multiple applications in health care monitoring, industrial automation, and sustainable energy management.

Efficient distributed intelligence is achieved through localised inference, collaborative learning, adaptive workload distribution, secure communication, and energy-aware resource management, leading to better computational efficiency, lower latency, higher prediction accuracy, more efficient power usage, and easy deployment across a variety of heterogeneous electronic environments. By integrating the architecture, a flexible foundation is created capable of supporting continuous monitoring, autonomous decision-making, predictive analytics, and intelligent control while ensuring reliable performance across a variety of operational conditions. Beyond restricted application domains, the work's overall impact lies in the fact that both disciplines, artificial intelligence and modern electronics, support each other in building sustainable digital infrastructures through the efficient use of resources, intelligent automation, increased operational robustness, and ecologically friendly computing. The presented architecture also lays the groundwork for future technology advances with the possibility of TinyML optimisation for ultra-low power embedded intelligence, explainable artificial intelligence for transparency in decision support, neuromorphic computing for brain-inspired processing efficiency, digital twin technologies for real-time processing of a system in simulation, and next-generation edge intelligence architectures for adaptive collaborative learning in increasingly interconnected electronic ecosystems. These research directions also highlight the long-term impact and relevance of the proposed framework for future autonomous cyber-physical systems, sustainable engineering solutions, and intelligent electronics.

Acknowledgement: The authors sincerely acknowledge the support and academic environment provided by Mahendra Institute of Technology, Vels Institute of Science, Technology and Advanced Studies, and the National Autonomous University of Mexico. Their facilities, resources, and encouragement significantly contributed to the successful completion of this research.

Data Availability Statement: The data supporting the findings of this study are available from the corresponding author upon reasonable request.

Funding Statement: No external funding was received for this research.

Conflicts of Interest Statement: The authors declare that there are no conflicts of interest.

Ethics and Consent Statement: This study was conducted in accordance with ethical standards, and informed consent was obtained from all participants.

References

1. P. Warden and D. Situnayake, “TinyML: Machine Learning with TensorFlow Lite on Arduino and Ultra-Low-Power Microcontrollers,” *O'Reilly Media*, Sebastopol, California, United States of America, 2019.
2. V. Sze, Y. H. Chen, T. J. Yang, and J. S. Emer, “Efficient Processing of Deep Neural Networks: A Tutorial and Survey,” *Proceedings of the IEEE*, vol. 105, no. 12, pp. 2295–2329, 2017.
3. J. Andreu-Perez, C. C. Y. Poon, R. D. Merrifield, S. T. C. Wong, and G. Z. Yang, “Big Data for Health,” *IEEE Journal of Biomedical and Health Informatics*, vol. 19, no. 4, pp. 1193–1208, 2015.
4. P. Szolovits, R. S. Patil, and W. B. Schwartz, “Artificial Intelligence in Medical Diagnosis,” *Annals of Internal Medicine*, vol. 108, no. 1, pp. 80–87, 1988.
5. Y. Lei, N. Li, L. Guo, N. Li, T. Yan, and J. Lin, “Machinery health prognostics: A systematic review from data acquisition to RUL prediction,” *Mechanical Systems and Signal Processing*, vol. 104, no. 5, pp. 799–834, 2018.
6. W. Shi, J. Cao, Q. Zhang, Y. Li, and L. Xu, "Edge Computing: Vision and Challenges," *IEEE Internet of Things Journal*, vol. 3, no. 5, pp. 637–646, 2016.
7. 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,” *Nature Medicine*, vol. 25, no. 1, pp. 65–69, 2019.
8. P. Rajpurkar, A. Y. Hannun, M. Haghpanahi, C. Bourn, and A. Y. Ng, “Cardiologist-level arrhythmia detection with convolutional neural networks,” *arXiv preprint arXiv:1707.01836*, 2017. [Accessed by 06/11/2017].
9. H. Park, S. Kang, J. Kim, J. Lee, S. Park, and S. M. Park, “A CMOS Hybrid System for Non-Invasive Hemoglobin and Oxygen Saturation Monitoring with Super Wavelength Infrared Light Emitting Diodes,” *Micromachines*, vol. 16, no. 10, p. 1086, 2025.
10. L. Jing, M. Zhao, P. Li, and X. Xu, “A convolutional neural network based feature learning and fault diagnosis method for the condition monitoring of gearbox,” *Measurement*, vol. 111, no. 12, pp. 1–10, 2017.
11. L. Wen, X. Li, L. Gao, and Y. Zhang, “A new convolutional neural network-based data-driven fault diagnosis method,” *IEEE Transactions on Industrial Electronics*, vol. 65, no. 7, pp. 5990–5998, 2018.
12. R. Liu, B. Yang, E. Zio, and X. Chen, “Artificial intelligence for fault diagnosis of rotating machinery: A review,” *Mechanical Systems and Signal Processing*, vol. 108, no. 8, pp. 33–47, 2018.
13. S. R. Bose, R. Vinoth, J. A. Jeba, R. Chauhan, and C. C. Angelin, “Artificial intelligence (AI) based symptom analysis using deep learning,” *AVE Trends in Intelligent Health Letters*, vol. 1, no. 4, pp. 243–254, 2024.
14. D. C. Mocanu, E. Mocanu, P. Stone, P. H. Nguyen, M. Gibescu, and A. Liotta, “Scalable training of artificial neural networks with adaptive sparse connectivity inspired by network science,” *Nature Communications*, vol. 9, no. 1, p. 2383, 2018.
15. M. P. William, S. B. Kiruba, M. Sakthivanitha, E. S. Soji, and G. Arun, “Machine learning to discover cardiovascular disease onset and key contributors: Data-driven personalized healthcare and preventive strategy,” *AVE Trends in Intelligent Health Letters*, vol. 1, no. 3, pp. 137–157, 2024.
16. G. Wang, Q. S. Jia, J. Qiao, J. Bi, and M. Zhou, “Deep learning-based model predictive control for continuous stirred-tank reactor system,” *IEEE Transactions on Neural Networks and Learning Systems*, vol. 32, no. 8, pp. 3643–3652, 2021.
17. S. K. Sehrawat, “Leveraging AI for early detection of chronic diseases through patient data integration,” *AVE Trends in Intelligent Health* Letters, vol. 1, no. 3, pp. 125–136, 2024.
18. T. Chen, T. Moreau, Z. Jiang, L. Zheng, E. Yan, H. Shen, M. Cowan, L. Wang, Y. Hu, L. Ceze, C. Guestrin, and A. Krishnamurthy, “TVM: An automated end-to-end optimizing compiler for deep learning,” in *Proc. 13th USENIX Symp. Operating Systems Design and Implementation (OSDI)*, Carlsbad, California, United States of America, 2018.
19. M. G. Hariharan, S. Saranya, P. Velavan, E. S. Soji, S. S. Rajest, and L. Thammareddi, “Utilization of artificial intelligence algorithms for advanced cancer detection in the healthcare domain,” in Advances in Medical Technologies and Clinical Practice, *IGI Global*, United States of America, 2024.
20. 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,” *Scientific Data*, vol. 7, no. 1, p. 154, 2020.
21. C. Elayaraja, J. Rahila, P. Velavan, S. S. Rajest, T. Shynu, and M. M. Rahman, “Depth sensing in AI on exploring the nuances of decision maps for explainability,” in Advances in Computational Intelligence and Robotics, *IGI Global*, United States of America, 2024.
22. P. Kairouz and H. B. McMahan, “Advances and open problems in federated learning,” *Foundations and Trends in Machine Learning*, vol. 14, no. 1–2, pp. 1–210, 2021.