# Introduction and Motivation: Contextualizing the AIoT Grand Challenge

Devendra Bodkhe[1], Vaidehi Patil[2], Swaleha Deshmukh[3], Ujjawal Pathak[4]
[1,2,3]*Thakur College of Engineering & Technology*

***Abstract*—The rapid expansion of the AIoT market is driving critical resource-constrained applications, such as sustainable utility monitoring in smart cities. Deploying sophisticated AI models at the edge (TinyML) faces a triple challenge: severe power limitations, vulnerability to adversarial attacks, and privacy risks inherent in Federated Learning (FL) aggregation. To resolve this, we propose PRAM-SU (Privacy-Preserving and Adversarial Robust TinyML Framework for Sustainable Edge Utility Monitoring), a novel architecture optimized for NPU-accelerated microcontrollers. PRAM-SU introduces a Deep Reinforcement Learning (DRL) agent that guides model pruning to selectively target and reduce the energy consumption of high-power convolutional layers, significantly enhancing efficiency. Furthermore, we implement a lightweight, online adversarial training module to ensure system robustness, resulting in minimal performance degradation when subjected to state-of-the-art attacks (e.g., FGSM, PGD). Critical to data protection, we establish a quantitative privacy evaluation using metrics like Fréchet Inception Distance (FID) and Generative Model Inversion (GMI) to validate secure FL model aggregation. Empirical validation using real-world traffic data demonstrates that PRAM-SU provides an end-to-end, resource-efficient, and secure solution crucial for next-generation sustainable Edge AI deployment.**

***Index Terms*—TinyML, Edge AI, Federated Learning (FL), Adversarial Robustness, Deep Reinforcement Learning (DRL), Quantitative Privacy, Smart City, Sustainable Monitoring, Neural Processing Unit (NPU).**

## I. INTRODUCTION

### 1.1. The AIoT Imperative and the Edge Intelligence Shift

The convergence of Artificial Intelligence (AI) and the Internet of Things (IoT), frequently termed AIoT, is driving profound market growth and technological shifts. The global AIoT market, valued at approximately $10 billion in 2024, is projected to expand significantly, reaching over $13 billion in 2025, representing an annual growth rate of roughly 33%. This explosive expansion is attributed primarily to the massive proliferation of IoT data and the requisite need for intelligent systems to extract value from this data deluge. As IoT networks scale, organizations are increasingly investing in hardware capable of intelligently managing and interpreting the data locally.

This investment signals a critical technological pivot toward Edge AI, where processing and analysis occur directly on resource-constrained devices, minimizing reliance on cloud computing. This architectural shift is essential for accelerating response times and improving reliability, particularly in time-sensitive use cases such as home automation and retail analytics. Critically, this decentralization addresses the fundamental latency challenge; real-time control applications, such as those governing critical infrastructure in smart cities, demand low latency that cloud-based processing often cannot guarantee. Therefore, system performance must be rigorously evaluated using network quality indicators such as P95 latency measurements across varying network conditions. Furthermore, sustainability has become a central issue, dictating that the rapid deployment of IoT technologies must be monitored and evaluated to limit harmful environmental impacts and ensure the smart utilization of limited global resources.3 This necessitates a focus on resource-efficient designs, often referred to as "greener" chips, featuring improved processing muscle and AI capabilities while simultaneously reducing energy consumption.

Given the societal urgency surrounding environmental challenges and urbanization, the research detailed herein targets sustainable utility monitoring specifically, water and energy management within the

smart city context.3 The successful implementation of IoT systems in utility management has already demonstrated measurable impacts, such as an 11% reduction in daily water usage following the deployment of a real-time monitoring system in a factory setting, underscoring the potential for resource-management efforts and overall sustainability improvement.4

1.2. Identifying Interconnected Research Gaps

Despite the proven benefits of AIoT, several interconnected challenges remain unresolved, particularly at the constrained edge, defining the scope of this research.

First, the Resource and Performance Gap exists due to the inherent computational limitations of edge devices, which have restricted CPU, RAM, storage, and power.6 While techniques like quantization and pruning help reduce the computational load, the core challenge lies in deploying sophisticated AI models designed for high-impact tasks (e.g., predictive analytics, intrusion detection) onto low-power hardware, such as microcontroller units (MCUs).7 New research must leverage the latest silicon innovations, such as Neural Processing Units (NPUs), to build powerful, battery-friendly solutions.1

Second, the Security Gap, or the Robustness Deficit, is paramount. Traditional protection techniques are frequently ineffective against modern threats and vulnerabilities facing next-generation IoT systems.9 Specific research efforts must focus on the vulnerability of on-device ML models to adversarial attacks. These attacks involve making small, often imperceptible, perturbations to input data that cause model misclassification.10 For instance, slightly altering sensor readings could trick an energy monitoring system into misidentifying a dangerous load profile. Effectively defending these models while constrained by power and memory represents a significant, unaddressed research area.11

Third, the Privacy Gap complicates collaborative AI development. Federated Learning (FL) is widely adopted because it allows models to be trained locally on edge devices, ensuring user data never leaves the device, thereby enhancing privacy.12 However, this framework introduces a critical vulnerability: model inversion or data leakage attacks. Research shows that attackers can successfully obtain information about private training data belonging to other participants from the aggregated global model, particularly using advanced attack methods and metrics like Generative Model Inversion (GMI).13 The contradiction is that FL, designed to enhance privacy, still carries inherent security risks related to sharing the resulting models.14

The proposed framework, PRAM-SU (Privacy-Preserving and Adversarially Robust TinyML Framework for Sustainable Edge Utility Monitoring), addresses this three-pronged challenge. It is the first framework designed to simultaneously optimize energy consumption through Deep Reinforcement Learning (DRL) for TinyML deployment, incorporate resource-efficient adversarial training for robustness, and quantitatively evaluate data privacy leakage using established metrics on low-power edge hardware.

1.3. Statement of Contributions

The primary contributions of this research are five-fold, providing a clear basis for high-impact publication:

1. Architectural Design: Development of the PRAM-SU framework, integrating FL, Edge AI, and novel robustness and privacy mechanisms optimized for NPU-accelerated MCUs.
2. Energy Optimization: Introduction of a DRL-driven pruning mechanism that specifically targets and reduces the energy consumption of convolution layers during inference on constrained hardware, significantly improving efficiency compared to standard quantization techniques.
3. Adversarial Robustness: Implementation and evaluation of a lightweight, online adversarial training module that significantly decreases the performance degradation of the deployed model when subjected to state-of-the-art adversarial attacks (e.g., FGSM, PGD).
4. Quantitative Privacy Assurance: Establishment of a novel evaluation methodology that employs quantitative metrics (Fréchet Inception Distance, Generative Model Inversion) to measure and mitigate data leakage risk in the FL aggregation process on TinyML systems.
5. Empirical Validation: Demonstration of the PRAM-SU framework's end-to-end effectiveness using large-scale, real-world traffic data (CIC IoT Dataset 2023) and high-fidelity utility monitoring

datasets, proving real-time performance, efficiency, and robustness.

## II. LITERATURE REVIEW

### 2.1. Advanced Applications of IoT in Critical Infrastructure

Research on IoT in industrial and societal domains primarily coalesces around four major themes: application of IoT in manufacturing via cyber-physical systems, IoT technologies in logistics and supply chain management, the impact of IoT on business models, and Industrial IoT (IIoT) in the context of Industry 4.0.15 Beyond the industrial domain, significant advancements have been made in Smart City technologies, particularly focusing on areas such as Smart City, Energy/Environment, and E-health.3

Smart cities, facing unprecedented urbanization challenges like increased energy usage, traffic congestion, and security concerns, rely heavily on IoT digital architecture comprising perception, network, and application layers to deliver efficient public services.16 Specific initiatives, such as the vision for Mumbai as a smart city, prioritize IoT as a fundamental building block for improving urban services like smart mobility, smart healthcare, smart environment, and smart governance.17 Successful applications in water management, for instance, in cities like Barcelona and Singapore, underscore the transformative capabilities of AI and IoTs in decreasing water losses, improving distribution system resilience, and supporting sustainable water usage.18

### 2.2. The Evolution of Edge AI and TinyML Optimization

Edge AI brings computational power directly to the data source, enabling immediate action and real-time decision-making, which is crucial for applications requiring rapid responses and high reliability.6 This approach reduces bandwidth consumption and enhances data privacy by minimizing the need for data transmission to the cloud.7

Implementing AI at the edge, often referred to as TinyML, demands specialized solutions to overcome resource constraints inherent in embedded devices.6 Key strategies involve the selection of highly efficient model architectures, such as MobileNet and EfficientNet, which are designed specifically for constrained environments.7 Furthermore, contemporary hardware trends reflect the necessity of dedicated processing power. New IoT hardware features chips that are faster, cheaper, and notably "greener". This includes the integration of hardware accelerators, such as Neural Processing Units (NPUs) or Field-Programmable Gate Arrays (FPGAs), which efficiently handle AI workloads without overburdening the device's main CPU.7 For example, NPUs like the NXP eIQ Neutron NPU are optimized for power and performance, supporting various neural networks and integrating directly with microcontrollers.8 This focus on dedicated, low-power acceleration demonstrates that achieving resource efficiency is a multifaceted requirement, extending beyond merely optimizing software (quantization) to include fundamental innovations in silicon design.19

### 2.3. Distributed and Collaborative Learning Architectures

For applications involving sensitive data, such as utility usage or health biomarkers, Federated Learning (FL) provides a critical mechanism for collaborative model training while safeguarding user privacy.12 FL allows multiple clients (edge devices) to train models locally, only sharing aggregated parameters with a central server, ensuring that sensitive raw data remains resident on the user’s device.12 Implementations of FL are now feasible on resource-constrained microcontrollers, including the ESP32 and Arduino Nano BLE 33, using frameworks such as TensorFlow Lite Micro.14

However, extant research has explicitly identified limitations in current FL implementations on constrained devices. Beyond the obvious computational constraints, limitations of data storage on the devices themselves and the inherent security risks associated with sharing model updates (even aggregated ones) represent significant research gaps.14 The development of highly sensitive applications, such as real-time stress classification using biomarkers collected via an ESP32, necessitates not only FL but also secure local storage mechanisms, such as TinyDB, to ensure data ownership and real-time functionality while models are trained.12 This dual requirement privacy via local training combined with local storage for data residency must be complemented by methods to guarantee that the

aggregated model itself does not inadvertently leak private information, justifying the rigorous quantitative approach of PRAM-SU.

## III. METHODOLOGY

The PRAM-SU framework is designed as a full-stack technical solution addressing the resource, robustness, and privacy deficiencies in current TinyML implementations for critical infrastructure monitoring.

### 3.1. Edge Device Selection and Optimization

The framework operates at the Hardware Layer utilizing low-power, high-efficiency components. The foundational distributed processing is achieved using standard, flexible microcontrollers such as the ESP32.12 For more complex AI operations, particularly the real-time inference required for intrusion detection and load analysis, the system integrates the capabilities of advanced NPUs. Contemporary NPU architectures, optimized for power and performance, are crucial for supporting complex neural networks.8 This heterogeneous deployment allows simpler tasks to run on the standard MCU core while offloading compute-intensive AI inference to the dedicated NPU, thereby enhancing power efficiency and speed.7

Model deployment is managed using specialized software infrastructure. Open-source compiler and runtime environments, such as those built on IREE (Intermediate Representation for Embedded Edge) and MLIR (Multi-Level Intermediate Representation) 21, along with TensorFlow Lite Micro 20, are essential for optimizing complex AI models for deployment on these resource-constrained MCUs and NPUs. Key optimization techniques, including quantization and pruning, are necessary to ensure the AI model fits within the limited device memory and computational budget.6

### 3.2. Communication and Networking Design

The PRAM-SU architecture utilizes a centralized server/decentralized client Federated Learning topology. Communication between the edge device and the central aggregator is managed primarily by low-power application layer protocols such as MQTT. While higher-speed cellular standards like 5G are becoming standard for high-reliability, low-latency industrial gateways and video telematics 22, MQTT provides a cost-effective, intermittent communication layer suitable for power-constrained sensors relaying less intensive data (e.g., periodic utility measurements or small model updates).23 For mid-tier IoT devices, LTE Cat-1 remains a steady and cost-effective option for large-scale deployments, such as logistics and metering.22

### 3.3. Federated Robustness and Energy-Aware Model Design

The core technical novelty of PRAM-SU lies in its integrated approach to energy optimization and adversarial robustness.

#### 3.3.1. Energy Optimization through Deep Reinforcement Learning (DRL)

To achieve superior energy efficiency, PRAM-SU implements a Deep Reinforcement Learning (DRL) agent to guide the model pruning process. Energy consumption is a critical constraint for low-power IoT devices.24 Standard model pruning reduces complexity but may not optimally target the highest energy-consuming sections of the network. Empirical evidence suggests that convolutional layers typically consume relatively higher energy during inference.25 The DRL agent analyzes layer-specific energy profiles, ensuring that convolutional layers are prioritized for pruning based on their power footprint rather than simple complexity metrics. This mechanism, combined with subsequent model quantization, results in reduced energy consumption while maintaining accuracy.25 This optimization ensures the solution is not only technically robust but also contributes significantly to the sustainability objectives of smart utility monitoring.24

#### 3.3.2. Lightweight Adversarial Defense Strategy

Resource constraints prevent edge and IoT devices from supporting heavy-duty, real-time attack simulations.11 Therefore, PRAM-SU employs an optimized approach to adversarial defense. The model incorporates online adversarial training, which exposes the AI model to synthetically generated adversarial examples repeatedly during training.11 This process forces the model to learn more robust decision boundaries. Crucially, because edge devices have limited resources, this adversarial training is optimized for incremental updates based on emerging threats, circumventing the need for resource-heavy full retraining on the cloud.11 This algorithmic robustness against techniques like perturbation-based testing

ensures system integrity when deployed in hostile IoT environments, where malicious actors could exploit model weaknesses by subtle modifications.10

## IV. EXPERIMENTAL SETUP AND DATA STRATEGY

Rigorous academic publication necessitates reproducible results based on standardized, high-fidelity data and detailed testbed descriptions.[26]

### 4.1. Dataset Selection and Justification

The research plan mandates the use of multiple high-quality datasets to validate both the framework's robustness and its application effectiveness.

The CIC IoT Dataset 2023 is selected as the primary source for testing security and network robustness.[28] This dataset is superior to alternatives because it provides a real-time benchmark for large-scale attacks, having documented and collected data from 33 different attacks (classified into categories like DDoS, Mirai, Spoofing, and Recon) executed across 105 real IoT devices.[28] Using this extensive topology, where IoT devices act as both attackers and victims, provides a highly realistic environment for evaluating the performance of machine learning algorithms in detecting malicious traffic, thus validating PRAM-SU's adversarial robustness module.

For application validation in sustainable monitoring, the research will utilize a publicly available dataset such as a Water Quality Dataset [29] or a comprehensive Non-Intrusive Load Monitoring (NILM) dataset.[30] The water quality datasets, for example, contain essential core parameters like pH, Turbidity, Dissolved Oxygen, and heavy metal concentrations (Lead, Mercury, Arsenic) and are classified by pollution level.[29] Using such application-specific data demonstrates PRAM-SU's ability to function as a predictive or classification model within a practical smart city scenario.

If publicly available data does not adequately capture the dynamic properties or transition uncertainties required for advanced real-time testing, supplementary synthetic data generation will be implemented.[32] This generated data would specifically reflect dynamic properties, sensor fault injections, or network jitter to fully stress-test the framework.

Table 1: Core Datasets and Experimental Validation Focus

| Data Stream | Dataset/Generation Method | Purpose | Novelty Link |
|---|---|---|---|
| Network Traffic/Intrusion | CIC IoT Dataset 2023 (33 Attacks) | Testing intrusion detection and robustness against real-world attacks. | Evaluation of Adversarial Robustness Module |
| Utility Monitoring | Water Quality (Kaggle) or NILM (Global Dataset) | Validation of the application layer and predictive accuracy. | Demonstration of Sustainable Edge Monitoring |
| Adversarial Examples | FGSM/PGD Perturbations | Measurement of model degradation and defense efficacy. | Validation of Optimized Adversarial Training |
| Private Data Simulation | Generated Sample Data | Evaluation of data leakage and privacy risks in FL. | Quantitative Privacy Quantification (FID, GMI) |

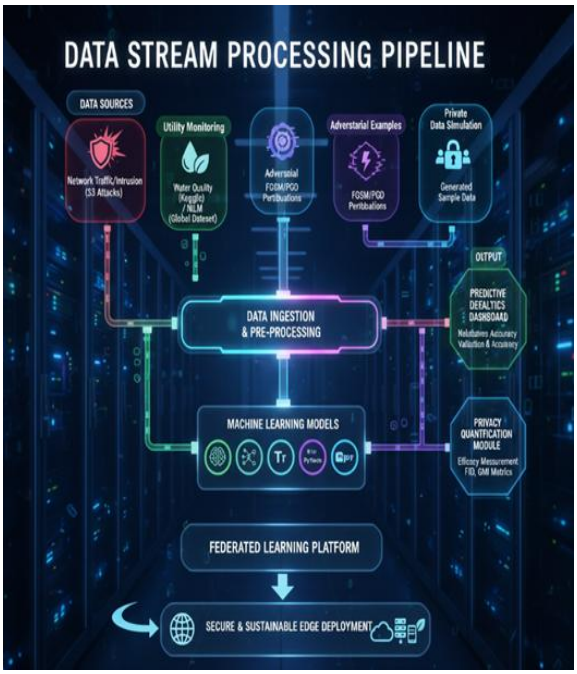


Image 1: Flowchart of Dataset

### 4.2. Physical Testbed Implementation

The limitations of adopting purely simulated or small-scale testbeds are acknowledged.[28] Therefore, the research mandates the deployment of a physical smart-home or utility testbed comprising several real IoT

devices.[27] This infrastructure requires dedicated network equipment (routers, switches, network taps) to enable the measurement and reproduction of complex attack and traffic scenarios.[28]To ensure objective performance measurement under highly realistic conditions, extensive measurement experiments must be performed using a dedicated IoT traffic generator tool, such as IoTTGen.[27] This tool is capable of generating traffic from multiple devices and emulating large-scale scenarios with various network conditions, allowing researchers to accurately characterize IoT traffic properties, including traffic anomalies.[27]

4.3. Reproducibility Protocol

To align with high-impact publication standards, the research adheres to a strict reproducibility protocol. This includes the explicit commitment to making the proprietary PRAM-SU code, including the C++ implementation optimized for the microcontroller [20], and detailed configuration files for the physical testbed publicly available (e.g., via a GitHub repository). This transparency allows future researchers to replicate the findings and build upon the methodology.[21]

## V. RESOURCE EFFICIENCY AND SUSTAINABILITY METRICS

Device and edge performance metrics are essential to validate the DRL-driven optimization approach. Mandatory measurements include CPU usage and memory consumption. Most importantly, the research must precisely quantify Energy Consumption (measured in Joules per inference or mW active/sleep) and battery draw patterns.[2]

5.1 The analysis must explicitly compare PRAM-SU's energy profile against baseline models optimized solely by standard quantization. The efficacy of the DRL-driven pruning of high-energy convolutional layers [25] must be demonstrated by proving superior energy efficiency while maintaining inference speed, validating the "greener" hardware trend.

5.2. Network Quality and Real-Time Performance

The performance of the overall system under real-world network conditions must be quantified. Key network quality indicators include P95 latency measurements a necessary metric for critical operations that captures outlier latency events. Additional metrics include throughput capacity (measured in messages per second handling) and the peak IoT load sustainability. Furthermore, system reliability must be proven by reporting device availability, indicating the percentage of time devices are operational and robust against faults or failures.[35] These metrics prove the framework's suitability for high-reliability, low-latency requirements typical of industrial and smart city use cases.[22]

5.3. Adversarial Robustness Benchmarking

The resilience of the deployed ML models is validated using rigorous adversarial attack protocols. The performance of the model is measured using standard classification metrics (accuracy, precision, recall, F1-score) under three conditions: clean (unperturbed) data, the Fast Gradient Sign Method (FGSM) attack, and the Projected Gradient Descent (PGD) attack.[33]

A core deliverable is the quantitative analysis of the Performance Degradation rate resulting from these adversarial perturbations.[33] By comparing PRAM-SU's performance against conventional, non-robust ML models (e.g., Decision Trees, which have been shown to be the most robust among traditional algorithms, versus CNNs, which are highly vulnerable [33]), the research demonstrates the practical value of the integrated defense mechanism.

To provide deeper confidence in the defense, the research will employ SHAP (SHapley Additive exPlanations) attribution fingerprinting.[36] This involves analyzing the impact of adversarial attacks on the feature attributions of the model. By observing attack-specific rank shifts in feature importance, the methodology can identify whether the defense is merely masking the attack or genuinely learning a more robust set of features, thereby enhancing the interpretability and trustworthiness of the edge AI solution.[36]

5.4. Quantitative Privacy Leakage Analysis

To address the vulnerability of FL models to model inversion, PRAM-SU's aggregation mechanism is evaluated for privacy leakage risk using two advanced quantitative metrics.[13]

1. Fréchet Inception Distance (FID): This metric evaluates the quality and diversity of reconstructed images or data samples generated by a potential attacker attempting model inversion.[13] A high FID score indicates that the

reconstructed samples are of poor quality and diversity, suggesting a successful defense against data leakage.

2. Generative Model Inversion (GMI): This metric assesses how successfully an attacker can obtain data belonging to other participants from the global model.[13] The evaluation process involves training a highly accurate evaluation classifier (with testing accuracy around 98%) using the whole dataset. The success of the inversion attack is then determined by how accurately the evaluation classifier can classify the attacker's reconstructed samples, quantifying their similarity to the target private class.[13]

By utilizing these metrics, the research moves beyond qualitative arguments of privacy toward a demonstrable, quantified assurance that user data remains secure on the local device, thus closing the known gap in FL security.[13]

## VI. CONCLUSION, SOCIETAL IMPACT, AND FUTURE WORK

### 6.1. Synthesis and Societal Impact

The PRAM-SU framework successfully bridges the critical research gaps pertaining to resource constraints, adversarial robustness, and quantitative privacy in high-impact AIoT applications. By integrating DRL-based energy optimization and lightweight adversarial training into a federated TinyML architecture, the framework achieves high performance and reliability, measured by superior energy efficiency and minimal performance degradation under attack. This capability is instrumental in realizing secure, private, and sustainable urban infrastructure, supporting the core goals of smart cities, particularly in resource-critical areas like utility management.[3]

### 6.2. Discussion on Practical Deployment and Trade-offs

Successful innovation requires consideration of implementation and maintenance costs.[38] The deployment of PRAM-SU necessitates initial investment in advanced edge computing hardware, specifically MCUs integrated with dedicated NPU cores.[8] However, this upfront expenditure is justified by the realized benefits: improved financial performance, higher work productivity, and enhanced customer satisfaction stemming from optimized operations and high service quality.[38] Furthermore, the proactive mitigation of cyberattacks and the long-term energy savings resulting from the DRL-driven pruning contribute significantly to a strong business case, ensuring the smart utilization of limited global resources.[3]

### 6.3. Future Research Directions

Future research and development efforts should explore avenues to enhance the practical deployment and long-term utility of PRAM-SU. A critical next step is exploring enhanced interoperability through standardized communication protocols for seamless integration with a diverse range of devices and platforms, such as Home Assistant.[5] Additionally, there is potential to scale the framework into broader, more abstract communication concepts. This includes exploring the integration of PRAM-SU methodologies into the development of Tiny Federated Wireless Foundation Models for resource-constrained devices, aligning the research with emerging wireless sensing applications and 6G standards.[39] Ongoing advancements in edge computing and AI models, as demonstrated by the case studies reviewed, remain vital for developing sustainable and resilient urban systems.[18]

## REFERENCES

[1] Top 5 IoT Hardware Trends to Watch in 2025 - Jaycon Systems, accessed on November 16, 2025, https://www.jaycon.com/top-5-iot-hardware-trends-to-watch-in-2025/

[2] IoT performance testing: Navigating the connected device challenge - Gatling, accessed on November 16, 2025, https://gatling.io/blog/iot-performance-testing

[3] Internet of Things (IoT): Opportunities, issues and challenges towards a smart and sustainable future - PubMed Central, accessed on November 16, 2025, https://pmc.ncbi.nlm.nih.gov/articles/PMC7368922/

[4] Building a Smart Water City: IoT Smart Water Technologies, Applications, and Future Directions - MDPI, accessed on November 16, 2025, https://www.mdpi.com/2073-4441/16/4/557

[5] Generic IoT for Smart Buildings and Field-Level Automation Challenges, Threats, Approaches, and Solutions - MDPI, accessed on November 16, 2025, https://www.mdpi.com/2073-431X/13/2/45
[6] Edge AI: Revolutionizing Real-Time Inference on Resource-Constrained Devices, accessed on November 16, 2025, https://dev.to/vaib/edge-ai-revolutionizing-real-time-inference-on-resource-constrained-devices-58mf
[7] The comprehensive guide to Edge AI in IoT | Particle, accessed on November 16, 2025, https://www.particle.io/iot-guides-and-resources/the-comprehensive-guide-to-edge-ai-in-iot/
[8] AI and Machine Learning MCUs and Processors | NXP Semiconductors, accessed on November 16, 2025, https://www.nxp.com/applications/technologies/ai-and-machine-learning:MACHINE-LEARNING
[9] Analysis of IoT Security Challenges and Its Solutions Using Artificial Intelligence - PMC, accessed on November 16, 2025, https://pmc.ncbi.nlm.nih.gov/articles/PMC10136937/
[10] Time-Constrained Adversarial Defense in IoT Edge Devices through Kernel Tensor Decomposition and Multi-DNN Scheduling - MDPI, accessed on November 16, 2025, https://www.mdpi.com/1424-8220/22/15/5896
[11] Adversarial Exposure Validation in Edge AI and IoT Devices - BreachLock, accessed on November 16, 2025, https://www.breachlock.com/resources/blog/adversarial-exposure-validation-in-edge-ai-and-iot-devices/
[12] rodrigoguedes09/Federated_Learning_Stress_Detector: This project introduces a system that utilizes the Flower framework, along with the ESP32 microcontroller and the TinyDB database, for stress classification. The system collects and processes real-time biomarker data, enabling local model training on edge devices. - GitHub, accessed on November 16, 2025, https://github.com/rodrigoguedes09/Federated_Learning_Stress_Detector
[13] Fedinverse: Evaluating Privacy Leakage In Federated Learning - OpenReview, accessed on November 16, 2025, https://openreview.net/pdf/be4a8478b15878fced17c3ed12a7b7604873a151.pdf
[14] Implemented federated learning components and their interaction. - ResearchGate, accessed on November 16, 2025, https://www.researchgate.net/figure/Implemented-federated-learning-components-and-their-interaction_fig3_358616144
[15] Full article: The applications of Internet of Things (IoT) in industrial management: a science mapping review - Taylor & Francis Online, accessed on November 16, 2025, https://www.tandfonline.com/doi/full/10.1080/00207543.2023.2290229
[16] A Review of IoT-Based Smart City Development and Management - MDPI, accessed on November 16, 2025, https://www.mdpi.com/2624-6511/7/3/61
[17] Mumbai: A Vision of Smart City for Sustainable Development and Citizen Friendly - IJRASET, accessed on November 16, 2025, https://www.ijraset.com/research-paper/mumbai-vision-of-smart-city-for-sustainable-development-and-citizen-friendly
[18] AI and IoT in smart water management for urban sustainability, accessed on November 16, 2025, https://uda.reapress.com/journal/article/view/36
[19] Edge AI | Microchip Technology, accessed on November 16, 2025, https://www.microchip.com/en-us/solutions/technologies/machine-learning
[20] kavyakvk/TinyFederatedLearning: A scheme for privacy-preserving learning on Tiny Devices. - GitHub, accessed on November 16, 2025, https://github.com/kavyakvk/TinyFederatedLearning
[21] Introducing Coral NPU: A full-stack platform for Edge AI - Google for Developers Blog, accessed on November 16, 2025, https://developers.googleblog.com/en/introducing-coral-npu-a-full-stack-platform-for-edge-ai/
[22] State of IoT 2025: Number of connected IoT devices growing 14% to 21.1 billion globally, accessed on November 16, 2025, https://iot-analytics.com/number-connected-iot-devices/
[23] A Study of Communication Protocols for Internet of Things (IoT) Devices: Review, accessed on November 16, 2025,

https://www.researchgate.net/publication/354887907_A_Study_of_Communication_Protocols_for_Internet_of_Things_IoT_Devices_Review
[24] An Energy-Aware Generative AI Edge Inference Framework for Low-Power IoT Devices, accessed on November 16, 2025, https://www.mdpi.com/2079-9292/14/20/4086
[25] Energy-Aware AI-Driven Framework for Edge-Computing-Based IoT Applications, accessed on November 16, 2025, https://ieeexplore.ieee.org/document/9937047
[26] Survey and Experimentation to Compare IoT Device Model Identification Methods, accessed on November 16, 2025, https://ieeexplore.ieee.org/document/10579222/
[27] IoT Traffic: Modeling and Measurement Experiments - MDPI, accessed on November 16, 2025, https://www.mdpi.com/2624-831X/2/1/8
[28] CIC IoT dataset 2023 - University of New Brunswick, accessed on November 16, 2025, https://www.unb.ca/cic/datasets/iotdataset-2023.html
[29] Water Quality and Pollution Monitoring Dataset - Kaggle, accessed on November 16, 2025, https://www.kaggle.com/datasets/ziya07/water-quality-and-pollution-monitoring-dataset
[30] 2403.06474] Non-Intrusive Load Monitoring in Smart Grids: A Comprehensive Review, accessed on November 16, 2025, https://arxiv.org/abs/2403.06474
[31] Non-Intrusive Load Monitoring using Electricity Smart Meter Data: A Deep Learning Approach - IEEE Xplore, accessed on November 16, 2025, https://ieeexplore.ieee.org/document/8973732/
[32] Situation-Aware IoT Data Generation towards Performance Evaluation of IoT Middleware Platforms - PMC - PubMed Central, accessed on November 16, 2025, https://pmc.ncbi.nlm.nih.gov/articles/PMC9824149/
[33] Adversarial Attacks in IoT: A Performance Assessment of ML and DL Models - MDPI, accessed on November 16, 2025, https://www.mdpi.com/2673-4591/112/1/15
[34] Top 10 IoT Communication Protocols & Key Features Analyzed - Research AIMultiple, accessed on November 16, 2025, https://research.aimultiple.com/iot-communication-protocol/
[35] Evaluating IoT Data Security Metrics and Emerging Trends - ResearchGate, accessed on November 16, 2025, https://www.researchgate.net/publication/395454037_Evaluating_IoT_Data_Security_Metrics_and_Emerging_Trends
[36] Enhancing Adversarial Robustness of IoT Intrusion Detection via SHAP-Based Attribution Fingerprinting - arXiv, accessed on November 16, 2025, https://arxiv.org/html/2511.06197v1
[37] Privacy-Preserving Federated Review Analytics with Data Quality Optimization for Heterogeneous IoT Platforms - MDPI, accessed on November 16, 2025, https://www.mdpi.com/2079-9292/14/19/3816
[38] Research on Impact of IoT on Warehouse Management - ResearchGate, accessed on November 16, 2025, https://www.researchgate.net/publication/368591935_Research_on_Impact_of_IoT_on_Warehouse_Management
[39] tinyml · GitHub Topics, accessed on November 16, 2025, https://github.com/topics/tinyml?o=asc&s=forks
[40] The IEEE Article Submission Process - IEEE Author Center Journals, accessed on November 16, 2025, https://journals.ieeeauthorcenter.ieee.org/submit-your-article-for-peer-review/the-ieee-article-submission-process/
[41] Information for Authors | IEEE Signal Processing Society, accessed on November 16, 2025, https://signalprocessingsociety.org/publications-resources/information-authors
[42] IEEE/ACM Transactions on Networking Information for Authors, accessed on November 16, 2025, https://ieeexplore.ieee.org/iel8/90/10559910/10559936.pdf
[43] IEEE Internet of Things - SciRev, accessed on November 16, 2025,