The Nucleus ***63***, *No. 1 (2026) 48-54*

https://doi.org/10.71330/thenucleus.2026.1517

www.thenucleuspak.org.pk

The Nucleus

ISSN 0029-5698 (Print)

ISSN 2306-6539 (Online)

AI-Driven ML and DL Approaches for Internet of Things Security in Developing Countries: A Systematic Review

Naseeb Sayadan, Muhammed Junaid Arshad, Syed Zafar Ali Shah

Department of Data Science, University of Engineering and Technology, Lahore

ABSTRACT

The Internet of Things (IoT) has significantly expanded the attack surface of modern digital infrastructure. In developing regions, low-cost devices are deployed across fragile networks for precision agriculture, public utilities, clinical settings, and city governance, often with minimal cybersecurity oversight. These conditions substantially elevate the risk of cyberattacks. Traditional intrusion detection systems impose computational and memory requirements that microcontroller-grade IoT hardware cannot satisfy. In this study, we systematically surveyed the use of lightweight AI, ML, and DL methods to protect resource-constrained IoT (RC-IoT) devices. Benchmarked solution classes range from shallow classifiers and compressed neural networks to edge-gateway models and federated learning frameworks, which are evaluated based on their accuracy, memory footprint, inference latency and energy consumption. Well-tuned models can maintain high detection rates without exceeding the hardware limitations. A fundamental trade-off exists between detection accuracy and energy consumption: deeper models that improve detection fidelity also increase power demand, a critical concern for battery-operated nodes in settings where reliable grid power cannot be assumed. Techniques such as federated learning and compact cryptographic primitives enable privacy-preserving coordination of distributed nodes. Grounded in a layered defense architecture, this review concludes with actionable guidance for engineers, institutions, and policymakers in resource-constrained environments.

Keywords: *IoT security, intrusion detection, resource-constrained devices, lightweight machine learning, edge computing, anomaly detection, federated learning, developing countries.*

1. Introduction

The Internet of Things (IoT) has emerged as one of the most prominent technologies for enabling smart cities, industrial monitoring, agriculture, healthcare, and a wide range of interconnected services [1]. IoT systems enable real-time data collection and automated decision-making by connecting sensors, actuators, and communication modules to cloud or edge platforms. However, the level of deployment has made security challenging. Millions of devices share data across diverse networks, and many are not designed with robust security in mind.

Rapid deployment, low power consumption, and small form factors are important areas of focus for manufacturers of these devices. Consequently, many IoT products remain outfitted with substandard firmware management, lack encryption, have limited update routines, and use default credentials [1]. These vulnerabilities enable attacks, such as creating botnets, launching distributed denial-of-service (DDoS) attacks, obtaining sensitive information, and gaining access to physical systems. The Mirai and Bashlite incidents underscore that IoT devices can be used as part of a large-scale attack infrastructure [2].

The situation is worse in developing nations in South Asia, Sub-Saharan Africa, and Latin America [3]. The IoT is a key asset for managing water-related systems, monitoring the environment, supporting telehealth, distributing energy, and enabling smart agricultural systems. Deployments frequently utilize low-cost microcontrollers that require minimal bandwidth and operate on tight public-sector budgets [4]. Robust security solutions in the cloud and enterprise firewalls are often out of reach because of reliability constraints, high processing power requirements, and human resource limitations.

These local constraints make AI-based intrusion detection an attractive alternative to cloud-dependent solutions for edge computing. Lightweight models can be used to analyze traffic at the gateway or near the device and detect unusual traffic patterns, thereby reducing the reliance on remote cloud processing [5]. This review can be distilled to a single question: What are the feasible avenues for AI-supported IoT security solutions in developing countries when a major design constraint limits hardware availability, energy consumption, and deployment costs?

2. Socio-Economic and Cybersecurity Context in Developing Countries

Developing areas cannot simply have technical issues related to IoT security. It is influenced by digital capacity (with the ESPO in mind), funding shortages, regulatory gaps, and a shortage of skilled cybersecurity personnel [6]. These limitations affect the types of security architectures that can be implemented and maintained, with long-term implications. Security solutions developed for enterprises in high-income countries are generally unsuitable for these contexts for several interconnected reasons. Enterprise-grade intrusion detection systems assume reliable cloud connectivity, dedicated IT personnel, and a budget for proprietary licenses, none of which can be guaranteed in low-resource settings. Hardware in developing-country deployments is often lower-specification (e.g., 8-bit or 32-bit microcontrollers with limited RAM), the power supply may be intermittent, and over-the-air update infrastructure is frequently absent. Regulatory and procurement frameworks may also lag, leading to insecure devices remaining in the

*Corresponding author: 2025msds10@student.uet.edu.pk

field for extended periods. Consequently, security solutions for these regions must be lightweight enough to run locally on constrained hardware, energy-efficient enough to operate on intermittent power, and maintainable by personnel with generalist rather than specialist cyber-security expertise.

Cyber resilience is not universal across all countries. A survey conducted worldwide as part of the Global Cybersecurity Outlook 2025 showed a clear lack of confidence between the developed and developing worlds, with Africa and Latin America expressing greater concern about NSIs' incident response capabilities [7]. In addition to high Internet penetration, Latin America also faces significant exposure to large-scale attacks; however, some countries are still working to establish basic national capabilities [8]. Insecurely shipped devices and a lack of localized IoT security frameworks could increase the exposure of public utilities in parts of Sub-Saharan Africa [9]. Meanwhile, South Asian countries are expanding their household and municipal IoT systems and are experiencing an increasing number of cyber incidents and a lack of local experience [3].

Design challenges stemming from infrastructure limitations also complicate security design. Continuous cloud offloading can be costly and unpredictable in rural/semi-urban deployments due to variable connectivity and bandwidth [10]. Most of the time, the use of basic packet traces can introduce privacy risks in transit and increase latency by sending the packet traces to remote servers [11]. Security intelligence must be closer to the network edge, where it can be near the data source and respond rapidly.

The same result holds true due to a lack of financial and human resources. Many public authorities and SMEs cannot justify the cost of enterprise-grade appliances or a dedicated team to operate them. A few countries have benefited from capacity-enhancing international programs, such as those of the World Bank, to build cyber resilience and from national systems, such as the national CSIRTs. However, local systems must remain affordable, accessible, and sustainable for the country's technical personnel.

3. Systematic Review Methodology

The review was conducted in accordance with the Preferred Reporting Items for Systematic Reviews and Meta-Analyses (PRISMA) guidelines [12]. The inclusion criteria required studies to (1) propose AI, ML, or DL approaches for IoT security; (2) report empirical hardware metrics such as accuracy, memory consumption, inference latency, or energy consumption on resource-constrained platforms; and (3) be published in peer-reviewed venues between 2013 and 2026. The exclusion criteria eliminated papers that focused solely on cloud-based IDS architectures, contributed only policy or vision statements without empirical data, relied exclusively on simulations without hardware grounding, or used datasets now considered outdated (e.g., the KDD Cup 99 dataset). Quality assessment was conducted using a structured scoring rubric that prioritized (a) direct hardware profiling on real IoT platforms (e.g., ESP32, ARM Cortex-M, Raspberry Pi), (b) reporting of at least three hardware-relevant metrics (accuracy/F1, memory footprint, and latency or energy), and (c) reproducible experimental setups with clearly described datasets and model configurations. The studies were independently evaluated, and borderline cases were resolved by consensus among the authors.

IEEE Xplore, ScienceDirect, ACM Digital Library, SpringerLink, and Google Scholar were searched. To build the core search string, a set of terms were combined for IoT, AI/ML/DL, intrusion detection, cybersecurity, resource constraints, and developing countries: ("IoT" OR "IoT") AND ("machine learning" OR "deep learning" OR "AI") AND ("intrusion detection" OR "cybersecurity" OR "anomaly detection") AND ("resource-constrained" OR "developing countries" OR "edge computing" OR "microcontroller").

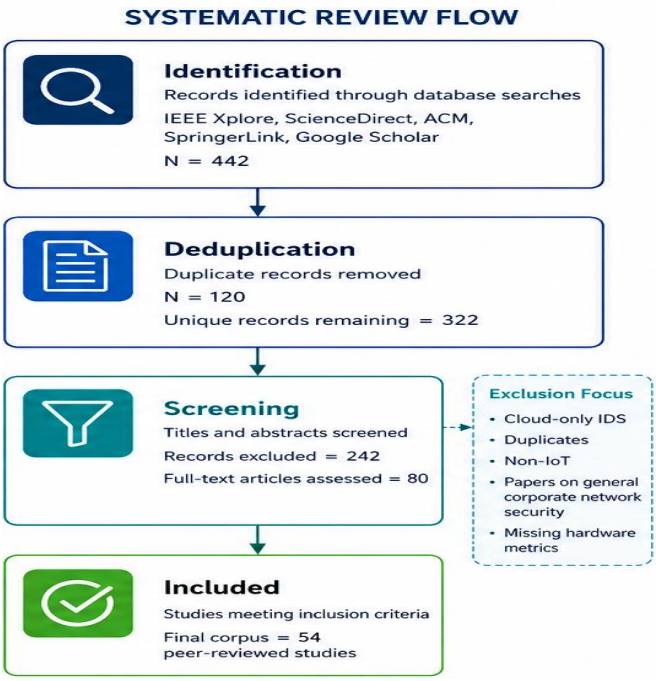

Fig. 1: PRISMA flow diagram showing the study selection process. Stage 1 (deduplication): 442 records retrieved; 120 duplicates removed, resulting in 322 unique records being included. Stage 2 (title/abstract screening): 242 papers were excluded for the following reasons: cloud-only IDS focus (n=89), no IoT relevance (n=74), and policy/vision papers without empirical data (n=79). Stage 3 (full-text review): 26 further exclusions: hardware metrics absent (n=14), purely simulated setups (n=7), and outdated datasets (n=5). Final corpus: 54 peer-reviewed studies forming the basis of the technical synthesis.

The selection process was conducted in four stages as follows. An initial search returned 442 records; after deduplication, 322 unique items remained in the review. Title and abstract screening eliminated cloud-focused IDS papers and non-technical contributions, such as policy statements and vision papers, reducing the pool to 80 studies. The full-text review assessed eligibility based on the

presence of hardware-relevant metrics, exclusion of purely simulated or outdated experimental setups, and explicit consideration of resource constraints. The final corpus comprised 54 peer-reviewed studies published between 2013 and 2026 and served as the basis for the technical synthesis.

4. Layered IoT Vulnerabilities and Security Placement

IoT security needs must be implemented at the same level as the architecture's ability in each layer of the IoT ecosystem. A single model used across the network is not always feasible, particularly if the network includes battery-powered sensors and more powerful local gateways.

Physical data are gathered by sensors and actuators, RFID tags, ESP32 devices, and ARM Cortex-M processors at the perception layer [13]. All of these have finite capacities for flash, RAM, and energy. Without optimizations, public-key operations such as RSA encryption or elliptic-curve cryptography can incur high costs on low-end nodes [14]. Threats include physical tampering, spoofing, device cloning, and unauthorized firmware changes.

The network layer includes routing and communication protocols, such as 6LoWPAN, RPL, MQTT, and CoAP, for data exchange [1]. Low-power routing protocols are susceptible to sinkhole, wormhole, Sybil, Selective forwarding, and replay attacks [2]. Therefore, this is significant for lightweight authentication, traffic profiling, and routing anomaly detection.

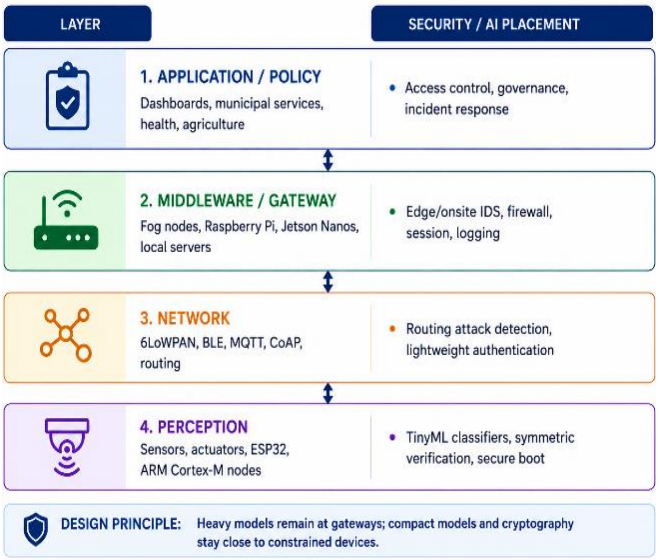

Fig. 2: Layered IoT security architecture with suitable AI and security placement for constrained deployments.

The computing power available at the middleware, gateway, and application layers is greater. Onboard heavier feature extraction, deep learning, and ensemble-based intrusion detection can be implemented using local fog servers, Raspberry Pi, and Jetson Nano. These layers are also simple to log, update, manage, and enforce policies and/or dashboards. Based on the literature reviewed, an appropriate design is stratification, with compact models and symmetric verification in the constrained nodes, robust DL, and hybrid classifiers in gateways.

5. AI-Driven Security Models for Edge and Gateway Deployment

The reviewed studies indicate that detection performance and resource costs are important considerations when adopting IoT intrusion-detection systems. Table I lists some hardware-validated and hardware-aware approaches.

5.1 Tree-Based Models

Tree-based models remain appealing in constrained environments, where the likelihood of intrusion can be captured by a small set of rules. Decision Trees and Random Forests are generally simpler, require fewer floating-point operations than Deep Networks, and are easier to understand. These models can be based on the most informative traffic features using feature selection methods such as mutual information, Gini index, Pearson correlation, and sequential forward selection [15]. Edge-IIoTset [9]. A feature-extracted Decision Tree was applied to forecast the state of an Edge-IIoTset deployment, achieving high accuracy on a Jetson Nano with sub-millisecond inference. Models based on hardware-aware LightGBM systems also performed well in terms of accuracy while requiring less flash memory [17].

5.2 Compressed Neural Networks

The second method uses a compressed network. Although standard deep neural networks (DNNs) tend to be large, pruning, quantization, and TinyML design can significantly reduce their size to fit within a microcontroller. An FNN was tested on an ARM Cortex-M microcontroller, achieving an F1 score of 0.976 for the ToN IoT dataset [18]. High accuracy has also been achieved with a few quantized 1D CNN models and a limited amount of flash memory and operations [17]. These findings indicate that leveraging hardware knowledge from the initial design is essential for efficient use of limited hardware resources in deep learning solutions.

5.3 Gateway-Level Deployments

Two larger models, sequential and hybrid, are possible for deployment at the gateway level. The autoencoder-LSTM network can reduce dimensions and capture temporal attack patterns. In contrast, CNN-LSTM and temporal convolutional networks can capture the temporal characteristics of traffic flows across the spatial dimension [19, 20]. These models require more processing power but can be used in systems running Raspberry Pis or edge coprocessors, where multiple downstream endpoints can be shielded using a single, locally configured gate.

5.4 Summary: Model Selection for Constrained Environments

Across the reviewed studies, three deployment tiers emerged as the most suitable for constrained environments. At the sensor node tier (e.g., ARM Cortex-M, ESP32), tree-

based models, such as decision trees, light gradient boosting machines (LightGBM), and TinyML feedforward networks, offer the best balance of accuracy and resource frugality, typically fitting within tens of kilobytes of flash and achieving sub-millisecond inference times. At the local gateway tier (e.g., Raspberry Pi 4, Jetson Nano), autoencoder-LSTM hybrids and quantized CNNs are viable and can protect multiple downstream nodes from a single edge platform. For distributed or multi-site deployments, where raw data sharing is infeasible, federated learning frameworks secured with lightweight symmetric ciphers, such as SPECK-32, represent the recommended approach, accepting a modest accuracy trade-off in exchange for privacy and bandwidth efficiency. Practitioners are advised to select a model tier consistent with their target hardware platform, validate the inference latency and energy consumption directly on the device, and apply feature selection to reduce the computational overhead before deployment.

Table 1: Hardware-Validated Performance Comparison of AI Paradigms for IoT Security

Algorithm / Model	Dataset	Hardware	Accuracy / F1	Resource Footprint	Latency / Cost	Ref.
Decision Tree (DT)	Edge-IIoTset	Jetson Nano	99.94% / 0.999	< 1 MB	0.185 ms	[16]
LightGBM (HW-NAS)	Edge-IIoTset	Microcontroller	95.3% / 0.951	75 KB flash	Low	[17]
TinyML FNN	ToN_IoT	ARM Cortex-M	F1 = 0.976	31 KB	120K pps	[16]
Quantized 1D CNN	Edge-IIoTset	ARM Cortex-M4	97.2% / 0.970	190 KB flash	840K FLOPs	[17]
Autoencoder-LSTM	BoT-IoT / UNSW	Raspberry Pi 4	99.4% / 0.994	Low–Med	Low latency	[19]
TabNet-L (FP16)	BoT-IoT / TON	Edge coprocessor	92.1% / 0.921	98 MB peak	18.4 ms	[21]
SPECK-secured FL	CICIoT2023	ESP32-S3 clients	85.43%	504 B RAM	138.3 mJ / round	[13]
TFNN (k-NN FL)	Custom IoT	Low-power nodes	Competitive	Minimal	Label-only TX	[11]

Abbreviations: DT: Decision Tree; LightGBM: Light Gradient Boosting Machine; FNN: Feedforward Neural Network; CNN: Convolutional Neural Network; LSTM: Long Short-Term Memory; FL: Federated Learning; pps: packets per second; FLOPs: floating-point operations; TX: transmission.

6. Hardware Validated Deployment Profiling

Physical testing is mandatory because simulations cannot reveal memory bottlenecks, bus latency, thermal behavior, power consumption, or the implementation footprint on the embedded device, all of which are necessary to detect potential faults in the system. Several of the aforementioned studies went beyond desktop assessments and tested the models or their encrypted components on actual IoT devices [13].

A testbed was built to investigate the exchange of secure federated learning updates among three microcontroller-based clients connected to a Raspberry Pi 4-based aggregator [13]. The study included measuring voltage and current traces in the laboratory to contrast lightweight symmetric block ciphers with HMAC-SHA-256 for integrity and ECDH-HKDF for session-key generation.

The best software efficiency was achieved with the ESP32-S3 and the SPECK-32. It achieved a throughput of 8.68 MB/s, a mean memory of ~504 bytes, and consumed approximately 138.3 mJ per update round [13]. SIMON-32 was also moderately inefficient in terms of overhead, and PRESENT-64 performed poorly in software but showed efficiency in some hardware scenarios. For deployment in developing countries, these findings warrant a cautionary lesson: security considerations and ASIC results are not the only criteria for choosing the cryptographic implementation of the target microcontroller.

Edge-first deployment is supported by gateway-level profiling. Gateway profiling also enables a gateway-first deployment. IDS models can be deployed on more robust platforms with reasonable latency and power consumption requirements, such as Raspberry Pi 4 and Jetson Nano, which can serve as local fog nodes in schools, municipal systems, farms, clinics, and utility systems.

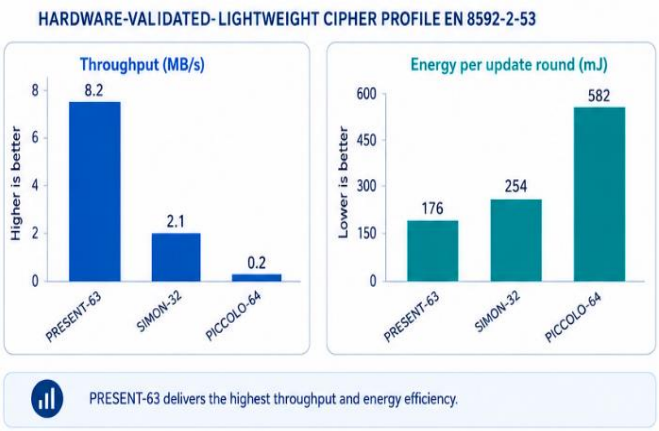

Fig. 3: Cipher profiling comparing the throughput and energy consumption of the ESP32-S3 microcontroller across different cryptographic algorithms.

7. Collaborative and Federated Learning Frameworks

The typical deployment of IoT devices in developing countries should be limited to a few sites with inadequate bandwidth and data. In particular, it is desirable to upload only information about the malicious flow learned by distributed nodes or gateways to train an intrusion detection model, rather than raw flow traffic, thereby making federated learning (FL) more valuable. In the usual FedAvg/federated learning architecture, the global model is sent from the server to clients, which train locally, then return protected updates, and the server averages the returned updates and generates the next global model to send back to the clients.

After 20 communication rounds, the secure federated IDS achieved 85.43% accuracy with ESP32 clients and the CICIoT2023 dataset, preserving the raw data at each client [13]. Although it is not as high as some centralized models, the privacy and bandwidth affordances, along with deployment capabilities, come into play when connectivity is spotty, or institutional data sharing is sensitive.

A less collaborative option is the Tiny Federated Nearest Neighbors (TFNN). Each node has its own k-nearest-neighbor classifier and sends only the label to the server. The labels are aggregated by majority voting on the server [11]. This minimizes communication overhead and is beneficial for sensor networks with limited power; however, it requires that the data quality from each sensor be acceptable and that each sensor participate adequately in the data-gathering process.

Another required component is lightweight mutual authentication. Symmetric primitives for CoAP authentication systems are sufficiently efficient for implementation in devices with stringent handshake costs (e.g., municipal devices) while avoiding the use of DTLS [22].

Despite its promise, federated learning for IoT security faces several practical challenges that must be addressed in real deployments. First, the communication overhead can be substantial; even transmitting compressed model updates across low-bandwidth links in rural or semi-urban deployments adds latency and energy costs per training cycle. Second, model poisoning poses a serious threat, in which a compromised client deliberately uploads manipulated gradients to degrade or backdoor the global model; defenses such as robust aggregation (e.g., Krum, Median aggregation) and anomaly detection of update distributions are active areas of research. Third, client heterogeneity, arising from differences in hardware capabilities, local dataset sizes, and class distributions across sites, can cause the global model to converge slowly or become biased toward more capable or data-rich nodes. Techniques such as FedProx, clustered federations, and personalized local adaptations have been explored to mitigate these effects. Practitioners deploying FL in developing countries should explicitly evaluate these challenges and select aggregation strategies and communication schedules that are consistent with local connectivity constraints and security threat models.

8. Benchmark Datasets for IOT Security Evaluation

The choice of dataset significantly influences the measured IDS performance. Although modern IoT-specific benchmark datasets improve ecological validity over older corpora, they are often limited in their generalizability. Common issues include a severe class imbalance between benign and attack traffic, limited diversity in attack types, lack of real-world operational variability, and biases introduced by synthetic traffic generation. Many datasets are collected in controlled laboratory environments that do not fully reflect the heterogeneity, noise, and protocol mix observed in real-world deployments, particularly in developing nations. Moreover, publicly available datasets are static snapshots that do not capture concept drift as device firmware, user behavior, and attacker tactics evolve. These limitations should be considered when interpreting the reported accuracy figures; high benchmark accuracy does not guarantee equivalent performance in production. Legacy datasets, such as KDD Cup 99 and NSL-KDD, are no longer relevant because they encode obsolete traffic patterns, exhibit high redundancy, and lack IoT-specific protocol coverage. Thus, modern IoT security research has focused on datasets of smart home, industrial, and edge device traffic [14].

Table 2: Comparative Taxonomy of IoT Security Benchmark Datasets

Dataset	Topology / Protocols	Primary Attacks	IoT Representativeness	Typical Best Use	Source
UNSW-NB15	Synthetic; standard IP	Analysis, backdoors, DoS, exploits	Moderate	Baseline comparison	[19]
BoT-IoT	Node-RED; MQTT, CoAP, HTTP, UDP	DDoS, DoS, fingerprinting, keylogging	Moderate; class skew	Botnet and DoS testing	[2]
ToN_IoT	Industry 4.0; Modbus, MQTT	Backdoor, DDoS, injection, ransomware	High	Industrial telemetry IDS	[18]
Edge-IIoTset	13 physical sensors; MQTT, BACnet, CoAP	DDoS, SQL injection, MitM, malware	Very high	Multi-domain IoT IDS	[14, 17]
CICIoT2023	105 smart-home devices	33 attack types incl. Mirai, web attacks	Exceptional	Large physical smart-home testing	[13, 14]
FL-MU	30 devices; AMQP, CoAP, MQTT	Slowite, sinkhole, Sybil, forwarding	High; FL-specific	Federated IDS benchmarking	[12]

In addition, there is a moderate representativeness of IoT devices in UNSW-NB15 and BoT-IoT. ToN_IoT enhances capabilities such as industrial telemetry and multi-attack detection, thereby improving efficiency in Industry 4.0 scenarios [4]. This is particularly pertinent given the various IoT domains [17] and physical sensors assigned to the Edge-IIoT set. CICIoT2023 also offers a large testbed of smart home devices, including 105 devices of various types, and can be readily used for various types of cyberattacks. FL-MU is specifically tailored for federated intrusion detection benchmarking [23].

9. Discussion and Open Challenges

9.1 Concept Drift

As IoT traffic evolves with firmware updates, service growth, user behavior, or attackers' tactics, it is important to understand how these changes influence traffic patterns. Thus, static models can become less accurate over time [2]. Continuous learning and adaptive retraining must be planned to be energy-limited for the system's memory and energy.

9.2 Hardware Validation Gap

Many models have not been tested on hardware, even when software is the target and the (typically low-cost) hardware platform is a microcontroller. Underreporting may appear acceptable when measured against published application bits, whereas the actual flash use, latency, or energy drain is unacceptable [24]. A set of standard metrics must be reported: accuracy or F1 score, false alarm rate, RAM and flash usage, inference latency, power consumption, and target hardware specifications.

9.3 Class Imbalance and Data Scarcity

Real IoT networks generate far more malicious traffic than benign traffic, and rare but critical attacks can be missed due to this imbalance. Oversampling methods, such as SMOTE, can be employed during training but may be excessively burdensome for use on devices. It is promising to work on lightweight cost-sensitive learning tasks, one-class anomaly detection tasks, and local calibrations.

10. Strategic Action Plan for Developing Regions

A helpful roadmap should encompass technical deployment, regulation, and capacity building. For technical teams, this implies an edge-first architecture that avoids sending large amounts of raw traffic to remote clouds. Simpler classifiers, such as TinyML or rule-based classifiers, can be deployed at a gateway to defend limited nodes, or more complicated models, such as quantized CNNs, pruned FNN, or LightGBM, can be used to protect more powerful machines [20, 25]. To consider Federated Learning for distributed deployments, the sharing of data must be considered, either because it is sensitive or because of bandwidth constraints [13].

Governments must establish minimum requirements for the cybersecurity of imported Internet of Things (IoT) products. Measures to prevent the use of universal default passwords can be incorporated into standards inspired by ETSI EN 303 645, which can mandate secure methods for applying updates and specify the actions to be taken when vulnerabilities are found. Low-cost devices are essential for developing markets, as many do not guarantee security.

Capacity building is also important. To ensure that deployed AI defense systems can be maintained by engineers, national CSIRTs, cybersecurity agencies, universities, and local technology sectors should collaborate to train the engineers. If the models in a technically strong solution cannot be changed locally by these teams, incidents cannot be responded to, or the behavior of these devices cannot be validated following deployment, then it will fail [26].

This review demonstrates that AI/IoT security can be viable across diverse settings with properly designed models that account for hardware constraints in developing countries. Lightweight, tree-based, quantized neural networks, and TinyML models can deliver effective solutions that do not consume excessive computational resources, local memory, latency, or energy at the gateway level.

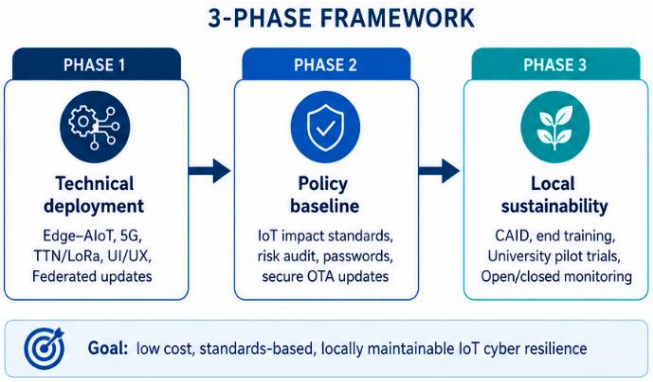

Fig. 4: Three-phase roadmap for AI-driven IoT security adoption in developing nations

11. Conclusion

It is also apparent from the evidence that hardware awareness is key to making deployment decisions. The model may be sufficiently accurate in simulations but is not a good fit when considering the memory and/or power requirements. Likewise, cryptographic protocols for federated learning need to be characterized using real microcontrollers placed in the field. The empirical profiling results for the ESP32-S3 using SPECK-32, SIMON-32, and PRESENT-64 highlight the importance of empirical profiling [13].

For developing countries, the recommended path forward is to deploy efficient, lightweight protection at the device layer, reinforce it at the gateway, extend it through privacy-preserving collaborative frameworks, and underpin the entire stack with national regulatory policies that restrict the importation of insecure devices. Specific future research directions for lightweight AI-based IoT security in developing countries include the following:

i. Adaptive and continual learning models can handle concept drift under strict memory and energy budgets.
ii. Standardized hardware benchmarking conventions enable fair cross-study comparisons on low-cost microcontroller platforms that are commonly used in these regions.
iii. Cost-sensitive and one-class learning techniques for class-imbalanced datasets that reflect real operational traffic patterns in smart agriculture, telehealth, and municipal deployments.
iv. Integration with low-power hardware accelerators, such as FPGA-based or RISC-V edge chips, which are increasingly accessible in emerging markets.

v. Open, regionally representative IoT traffic datasets capturing the protocols, device types, and attack profiles relevant to South Asia, Sub-Saharan Africa, and Latin America are required to address this issue. Pursuing these directions in parallel with institutional capacity building can protect the integrity of critical services while enabling sustainable digital transformation across the developing world.

References

[1] R. Ahmad and I. Alsmadi, "Machine learning approaches to IoT security: A systematic literature review," *IoT*, vol. 14, p. 100365, Jun. 2021.

[2] E. C. P. Neto, S. Dadkhah, R. Ferreira, A. Zohourian, R. Lu, and A. A. Ghorbani, "CICIoT2023: A real-time dataset and benchmark for large-scale attacks in IoT environment," *Sensors*, vol. 23, no. 13, p. 5941, 2023.

[3] L. H. Mahdi and A. A. Abdullah, "Fortifying future IoT security: A comprehensive review on lightweight post-quantum cryptography," *Engineering, Technology & Applied Science Research*, vol. 15, no. 2, pp. 21812–21821, Apr. 2025.

[4] A. Alsaedi, N. Moustafa, Z. Tari, A. Mahmood, and A. Anwar, "TON_IoT telemetry dataset: A new generation dataset of IoT and IIoT for data-driven intrusion detection systems," *IEEE Access*, vol. 8, pp. 165130–165150, 2020.

[5] A. Pasam, A. Sinha, and A. Mailewa, "Lightweight AI-based intrusion detection models for IoT devices: A comparative review," in *Proc. IEEE Int. Carnahan Conf. Security Technology (ICCST)*, Oct. 2025, pp. 1–7.

[6] N. Moustafa and J. Slay, "UNSW-NB15: A comprehensive data set for network intrusion detection systems (UNSW-NB15 network data set)," in *Proc. Military Communications and Information Systems Conf. (MilCIS)*, Nov. 2015, pp. 1–6.

[7] N. Moustafa and J. Slay, "The evaluation of network anomaly detection systems: Statistical analysis of the UNSW-NB15 data set and the comparison with the KDD99 data set," *Information Security Journal: A Global Perspective*, vol. 25, nos. 1–3, pp. 18–31, Apr. 2016.

[8] Y. Yigit, C. Chrysoulas, G. Yurdakul, L. Maglaras, and B. Canberk, "Digital twin-empowered smart attack detection system for 6G edge of things networks," in *Proc. IEEE Globecom Workshops (GC Wkshps)*, Dec. 2023, pp. 178–183.

[9] M. A. Ferrag, O. Friha, D. Hamouda, L. Maglaras, and H. Janicke, "Edge-IIoTset: A new comprehensive realistic cyber security dataset of IoT and IIoT applications for centralized and federated learning," *IEEE Access*, vol. 10, pp. 40281–40306, 2022.

[10] L. Takhellambam, U. Boby, N. Hoque, K. R. Singh, and M. Bhuyan, "FL-MU: A benchmark dataset for federated intrusion detection in IoT networks," *IEEE Access*, vol. 13, pp. 191037–191064, 2025.

[11] A. Imteaj, U. Thakker, S. Wang, J. Li, and M. H. Amini, "A survey on federated learning for resource-constrained IoT devices," *IEEE Internet Things J.*, vol. 9, no. 1, pp. 1–24, Jan. 2022.

[12] P. Fusco, A. Montefusco, G. P. Rimoli, F. Palmieri, and M. Ficco, "TinyML-based intrusion detection system for handling class imbalance in IoT-edge domain using Siamese neural network on MCU," in *Advanced Information Networking and Applications*, L. Barolli, Ed. Cham, Switzerland: Springer, 2025, pp. 389–402.

[13] A. Chatterjee and B. S. Ahmed, "IoT anomaly detection methods and applications: A survey," *IoT*, vol. 19, p. 100568, Aug. 2022.

[14] R. Srikanth, G. Vadlakonda, G. Simuni, and M. Sinha, "Federated learning for IoT: A decentralized approach to enhance privacy and efficiency in cyber-physical systems," SSRN, May 2025.

[15] P. K. Yemmanuru, J. Yeboah, and K. E. N. G, "Systematic literature review of machine learning for IoT security," in *Proc. Int. Conf. Computational Science and Computational Intelligence (CSCI)*, Dec. 2023, pp. 227–233.

[16] P. Agustin, G. Sebastian, and M. J. Erquiaga, "IoT-23: A labeled dataset with malicious and benign IoT network traffic," Zenodo, Jan. 2020.

[17] N. A. Hamad, K. A. A. Bakar, F. Qamar, A. M. Jubair, R. R. Mohamed, and M. A. Mohamed, "Systematic analysis of federated learning approaches for intrusion detection in the IoT environment," *IEEE Access*, vol. 13, pp. 95410–95444, 2025.

[18] M. Baqer, "Lightweight federated learning approach for resource-constrained IoT," *Sensors*, vol. 25, no. 18, p. 5633, 2025.

[19] A. I. Jony and A. K. B. Arnob, "A long short-term memory-based approach for detecting cyber-attacks in IoT using CIC-IoT2023 dataset," *Journal of Edge Computing*, vol. 3, no. 1, pp. 28–42, May 2024.

[20] A. B. Varlık, "Global cybersecurity disparities and the strategic vulnerability of high-income nations," *Güvenlik Bilimleri Dergisi*, vol. 15, no. 1, pp. 91–112, May 2026.

[21] A. Villafranca, I. Tasic, and M.-D. Cano, "TRUSTLab dataset: A real-world CICFlowMeter dataset for IoT/edge intrusion detection," *Frontiers in Computer Science*, vol. 8, Art. no. 1803271, 2026.

[22] N. Koroniotis, N. Moustafa, E. Sitnikova, and B. Turnbull, "Towards the development of realistic botnet dataset in the Internet of Things for network forensic analytics: Bot-IoT dataset," *Future Generation Computer Systems*, vol. 100, pp. 779–796, Nov. 2019.

[23] P. Fusco, G. P. Rimoli, and M. Ficco, "TinyML and Federated Learning for Resource-Constrained Medical Devices," in *Artificial Intelligence Techniques for Analysing Sensitive Data in Medical Cyber-Physical Systems*, M. Ficco and G. D'Angelo, Eds. Cham, Switzerland: Springer Nature, 2025, pp. 113–126.

[24] I. Brass, L. Tanczer, M. Carr, M. Elsden, and J. J. Blackstock, "Standardising a Moving Target: The Development and Evolution of IoT Security Standards," in *Living in the IoT: Cybersecurity of the IoT*, 2018, IET, London, U.K., doi: 10.1049/cp.2018.0024.

[25] R. Pietrantuono, M. Ficco, and F. Palmieri, "Survivability analysis of IoT systems under resource exhausting attacks," *IEEE Trans. Inf. Forensics Security*, vol. 18, pp. 3277–3288, 2023.

[26] A. Gerodimos, L. Maglaras, M. A. Ferrag, N. Ayres, and I. Kantzavelou, "IoT: Communication protocols and security threats," *IoT and Cyber-Physical Systems*, vol. 3, pp. 1–13, Jan. 2023.